Class Notes Semantics I & II, UT, Austin, Spring 2015-Spring 2018

Hans Kamp

September 7, 2018

Contents

Preface to the Current State of this Document

The present document consists of notes that were compiled over a number of years during which I taught first the second and then the first Formal Semantics course for graduate students in Linguistics at the University of Texas, Austin. The notes were at first produced as Word documents. What follows here started as an attempt to convert those files into Latex. But soon the conversion turned into something more ambitious, with many additions that were not contained in the Word documents as I had them.

At the time – September 2018 – that I am preparing these Notes for the website with published an unpublished work of mine (which is probably the site from which the reader will probably have got these Notes) the present document remains unfinished in several ways. First, only a comparatively small part of the original Word document describing the presuppositional treatment of tense and aspect has been incorporated in the current latex version. (It was because my original plans were otherwise that Chapter 4 bears the title 'Tense and Aspect II', in spite of the fact that of its current 230 pages only about 40 actually deal with matters of tense and aspect.) I decided in the end on the compromise of breaking off the present document at the point where it currently ends, since the stopping point can be seen as a natural conclusion to the task outlined at the beginning of PART II (see Section 3.1) and time for adding more right now is lacking.

This isn't the only sense in which the document is unfinished. The central and ultimate goal of these Notes is to provide new guidelines for the construction of logical forms of natural language sentences, discourses and texts and for the model-theoretic evaluations of those logical forms – and that for a substantial part of English. This goal is evidently an open-ended one; an enterprise pursuing such a goal will probably never reach the point where it can claim that its work is done. But apart from all else that is missing, there are some particular topics that I would have liked to add to the Notes and that I had originally planned to incorporate into them: Chapters about information structure, about propositional attitudes and about modality. For the first two of these detailed accounts within a DRT format exist, and it shouldn't be too hard to integrate these accounts within the general DRT-based architecture developed in PART II. But unfortunately I haven't yet found the time to do that.

The Notes are also unfinished in yet a third sense. They start with a comparatively brief discussion of the syntax-semantics interface of English in the spirit of Montague Grammar, in a form that closely follows the presentation in (Heim & Kratzer 1998), the book that has been used for many years in the UT Linguistics Department as textbook for first level graduate courses in formal semantics (just as it has in many other universities around the world). Ch. 1 of the Notes, which goes back to the times when I taught the Semantics II course at UT, was meant as a kind of rehearsal of material that I assumed the students, who had previously taken Semantics I, already knew. At the same time this discussion was used as the basis for a discussion in Ch.. 2 of the reasons that around 1980 led to Discourse Representation Theory and a presentation of DRT's original formulation, in which DRSs ('Discourse Representation Structures', the formulas of DRT's logical form languages) are constructed 'top-down' from syntactic sentence trees. All of this, Ch. 2 as well as Ch. 1, is preliminary to the version of DRT that is developed in PART II, in which DRSs are constructed 'bottom-up', in a manner that is much closer than the original top-down construction method to the compositionality principles of Montague Grammar.

In my own view, PART I is on the one hand, as an introduction to formal semantics in the spirit of Montague Grammar, too brief and too eclectic, and on the other hand, as a preamble to PART II it is unreasonably and discouragingly long. When teaching the Semantics II Seminar at UT in the years after I taught Semantics I I came to realize how little incoming cohorts of graduate students in Linguistics Programs can be relied upon these days to bring along even the most basic preliminaries for a serious course in formal semantics. So I decided to adopt a curriculum that starts from zero, beginning with a presentation of predicate logic, then the Lambda calculus, then Montague Grammar, then a brief transition from Montague Grammar to old style DRT and finally a (somewhat impressionistic) presentation of the bottom-up DRT developed in PART II of the present Notes. So far, most of this teaching material only consists in the form of slides. But the slides were put together with the aim of bing self-contained, and I think they should serve well enough to give someone who already knows this material in some form get a good idea of what the students were being offered. This material, consisting of the documents 'Semantics I, UT: Predicate Logic, Lambda Calculus, Montague Grammar' (slides) and 'Semantics I, UT: From Montague Grammar to DRT' (paper-like text), can be found elsewhere on the website that contains these Notes. (Website: http://www.ims.uni-stuttgart.de/archiv/kamp/). Eventually PART I of the Notes ought to be replaced by some of this material.

Even when considered on its own the orgnization of PART II leaves much to be desired. Over the years sections have been added to it in various places, in some instances with the tentative intention to make them into independent publications at some later point. Since these sections have some connections with topics that were part of the Notes already, parking them in places adjacent to where those topics are discussed does make a certain sort of sense. Among these later inserted sections, which may strike the reader as studies, or studies *in spe*, in their own right, are (i) those about the polymorphism of *and* and *or*, about definitions in natural language (as special forms and uses of biconditionals), and about a variety of different types of temporal adverbials in Ch. 3 and (ii) the discussion of interpretation strategies for coreference anaphora and bridging in Ch. 4. Given the often exploratory nature of these sections, as well as their mere size, these sections will probably impress the reader as excrescences (which is what they are). But I have left them in even so, in the perhaps vein hope that some of those who look at these Notes will look at some of these sections too, and may give me some advice about how to improve them and what to do with them.

Most importantly, if this document may look like a book, that is only because in the end I adopted Latex's 'book' format, as that is the one formatting option which automatically generates a Table of Contents and updates it whenever changes are made anywhere in the document. I struggled for years with the 'article' format. That not only makes for absurdly long sections in a document of this size – calling a part of several hundred pages of text a 'section' seemed after a while intolerably coy – but it also makes it necessary to update a handmade table of contents in the only way such a table can be updated: by hand. I have had to do that so often that in the end I got fed up and gave up and resolved to switch to the 'book' format. But the switch was made only for this reason. So, please don't think I think this is a book.

One respect in which this document does not live up to the standards for books are the serious gaps in its bibliography. Originally I had resolved not to include references at all. When I started on the notes, they were meant to be just an ad hoc aid to the students and references were to be supplied by different channels. Then, when the project grew, I began to add references. But since that was done in a rather haphazard and desultory manner, that has made the absence of references that are clearly missing all the more glaring. I hope to fix this problem – as well I can – at some later time. But for now I want to apologize in advance to all who will look at some of this material and who may be unable to find the citations that ought to be there.

PART I

Chapter 1

Introduction: the What and the How of these Notes

This is going to be a course on formal semantics of natural language, and that means that we will be making use of formal methods in dealing with issues about linguistic meaning.

One thing should be stressed from the start. What ultimately counts in formal semantics is the 'pre-formalization' analysis. The informal insight into how a construction works or what a word really means (in the sense of what semantic contributions it makes to the complex expressions that contain it) is crucial. If the informal analysis is no good, then no formalization will be able to save it or improve on it. All that a formalization can do for you in such a situation is to help you see how bad your pre-formalization analysis really was. (But of course, that can be quite useful too.)

When your pre-formalization insights are right, however, and your formalization tools are of the right sort, then formalization may often enhance the result, by sharpening the formulation of your analysis and by presenting it with a cogency and transparency that could not have been achieved without it.

Connected with this is a second benefit. When a formalization tool is right for the tasks in which it is employed, then it will not only provide good formalizations of good analyses; it will often also suggest new questions about the analyzed phenomena and others related to them; and sometimes it will also suggests answers to those questions, advancing inquiry in novel and often surprising ways.

Finally, formalization also serves another purpose. In the last instance, theories of natural language semantics should be theories of the semantics of entire *languages*, or at least of large, naturally circumscribable fragments of them. They should not be limited to the analysis of some collection of local problems – problems about particular grammatical constructions –, no matter how hard or challenging those problems may be or how ingenious the solutions. For what we ultimately want is an understanding of how the language works in its entirety. Understanding how a particular construction is built from its morpho-syntactic constituents, and how that morpho-syntactic blueprint determines the way in which the different constituents make their contributions to the meaning of the construction, is important in its own right. But even more important is often the way it augments our understanding of the compositional repertoire of the language as a whole.

The study of individual constructions can contribute to our understanding of the compositional system of a language in two complementary respects. On the one hand it can show us some of the compositional principles that are at work in the language: that the language is one that includes these principles within its repertoire – that its repertoire is at least this rich. But we may also find that when looked at the right way the principles involved in the particular construction we are looking at are the same that we also encountered when studying other constructions of the language, confirming a budding hypothesis that the repertoire we have already charted is all or nearly all that the language needs and uses. Such hypotheses are appealing on the one hand because they promise to make the linguist's task over-all an easier one: If the repertoire of basic operations is small, she can hope to reach the completion of her inventory of that repertoire sooner. But they are appealing also insofar as they confirm an intuition that many linguists share and that is in large part what drives their explorations: at a certain level the compositional resources of the language are really quite simple, notwithstanding the apparent diversity that meets the eye of the innocent observer; looked at in the right way that diversity is the result of ingenious redeployment, in ever new combinations, of the same small basic set of operations. A language with such a small repertoire of basic operations ought to have an advantage over languages with larger repertoires in that it is easier to acquire and perhaps also easier to apply in use: there is less to learn, and there are fewer tools to have to choose from both in the production of utterances and in their interpretation.

To what extent this intuition is justified is a point for debate. Much of formal semantics over the past decades has taken it for granted that the repertoire of basic operations is very small indeed, consisting of the operation of function application and perhaps one or two others; and a good deal of energy has gone into formalizing the analyses of individual constructions in a way that is compatible with this assumption. That endows the work in formal semantics with a certain uniformity which is agreeable and comforting to its practitioners. But we should not become oblivious to the question whether this underlying assumption is realistic – whether it is really true that these are all the operations that human languages have in their repertoire. In this course we will proceed with an open mind on this point.

While comprehensive formalizations of the semantics of languages should be the aim in the back of our minds, for many theoretical purposes such formalizations are neither needed not useful. Human languages are very big and complex systems, and that remains true if the hypothesis about a very small set of basic semantic operations should prove to be true. And this is equally true of comprehensive formalizations of their syntax and semantics: they become hard or impossible to survey for the human theorist, and the particular points that a given study wants to make about particular words or constructions would be drowned in a mass of information that is not directly relevant. For this reason, large formalizations have been carried out, to the extent that they have been, almost exclusively within Computational Linguistics. There they serve as the basis of machine-implemented algorithms that are designed for tasks such as automated summarization and machine translation. (An example of such an implemented comprehensive formalization is the CCG-DRS system Boxer of J. Bos[1]. In contrast, formal accounts that are useful to linguists because of the insights they provide into particular linguistic problems won't usually be like this. They will deal with just the few phenomena that are at issue, and the more perspicuous way of presenting them is by leaving most of the surrounding system out, or in the background. But those who present their results in such a localized and streamlined fashion should never forget that it ought to be possible to integrate those results into more comprehensive accounts. This is all the more important because it cannot be taken for granted that such an integration will be straightforward. Often new and unexpected problems arise in connection with the interactions between the phenomena to which the given results pertain and other phenomena which are treated in the more comprehensive account. And when that happens, some further careful analysis will be needed before the integration can be successful. This is a possibility for which we all have to keep

[1] See (Bos 2008)

an open eye, all the time.

One reason why reading the formal semantics literature can be a challenge is that different researchers employ different formalization tools. This means that a reader who wants access to a representative part of the literature will have to be equipped with a basic knowledge of more than one tool set. Among the things that a good training in formal semantics should provide is therefore something of an overview of the tool kits that are most widely used. But in addition to that a formal semanticist must have active, operative knowledge of at least one set of tools, so that she can apply those tools effectively in her own research. Ideally it should be left to the individual researcher which tools she wants to use, and her training should put her in a position to make an educated choice. But applying one set of tools consistently and effectively in dealing with a coherent set of research questions is something that has to be learned too, and learning to use just one set of tools is already quite a handful – more than enough for a single seminar. We too will therefore, after an introductory phase during the first couple of weeks of this seminar, concentrate on the use of just one tool set, which in our case will be a current incarnation of Discourse Representation Theory. One reason for this choice is that Temporal Reference and Aspect will play a prominent part among the phenomena with which we will deal, and that DRT seems a good tool for dealing with that topic. (It was in response to previously unexplored properties of tense in discourse that DRT was originally conceived, and some of the features that distinguish it from other formalization tools were introduced specifically for the sake of dealing with those properties.)

I have been speaking of formalization 'tools' so far. But in formal semantics those tools largely take the form of what might better be called a 'formalism' – a formal language which is used to specify the semantics of expressions of the 'Object Language' – the natural language or natural language fragment or natural language construction or set of phenomena that is being investigated). The traditional view is that there are two ways in which formalisms can specify the semantics of Object Language expressions: (i) by identifying the *semantic values* of those expressions, or (ii) by providing *logical forms* for them. However, in practice this is more a difference in conceptualization of what the semantic theory is or should be doing than a difference in what the theory is actually doing in practice. First, in either case the formalism will come with its own semantics (usually stated in model-theoretic terms; I am assuming the reader knows what that means, but if not, don't worry), which determines semantic values for its terms. If these terms are assigned as logical Forms to expressions of OL, then this confers their semantic values upon those OL expressions. Conversely, a theory whose purported goal is to assign semantic values to OL expressions will make use of its formalism by assigning terms to OL expressions as a way of identifying the semantic values of those expressions. But even when that is the philosophy behind the theory, it is in practice the form of the terms of the formalism that the theory uses to identify the semantic values it assigns to OL expressions which reveals most clearly what it has to say about the semantics of the Object Language. In this sense the terms used in semantic value theories play the part of logical forms no less than the logical forms of theories which present themselves openly as logical form theories, just as those logical form theories cannot help assigning semantic values to the expressions for which they provide logical forms. In view of this we will refer to such formalisms indiscriminately as 'Logical Form Formalisms' (LFFs), irrespective of whether the theory that uses them presents itself as logical forms or as semantic value theory.[2]

1.1 Interlude: Logical Form Formalisms and the formalization of linguistic theory

The use that is made of LFFs in formal semantics must be distinguished from what is normally understood by formalization in the Philosophy of Science. In Philosophy of Science formalization of a theory means that everything the theory says is presented in formalized form. In this way every prediction the theory makes takes the form of a theorem of the formalized theory – a statement that logically follows from its axioms. In the 20-th Century, formalization of this sort was considered by some as the ultimate goal of science. But in spite of that, even today there are not all that many examples of such formalizations, and most of them are found within pure mathematics. (The first systematic, remarkably comprehensive effort in this direction was Russell and Whitehead's *Principia Mathematica*.) Mostly, scientific theories are so complex that formalization is very difficult and that it doesn't promise much in the way of theoretical benefits, even were it carried out successfully. Furthermore, with theories in the empirical sciences, beginning with physics and chemistry, there is the problem of how theory is related to the practice of experimentation. One should like a formalization to make predictions – in the form of statements that are logically entailed by the axioms of the the-

[2] A quite detailed (if not fully up-to-date) précis of the formalisms that are most commonly used in formal semantics can be found on the website for this course, under 'Course Documents'.

ory – about factual observations and the outcomes of experiments that have been done or or are still to be performed. But the complexities of how the 'theoretical' predictions that the theory makes relate to the results obtained through actual observation or experimentation are such that this part of the theory tends to escape formalization. That is no conclusive reason for staying away from formalization but it takes away some of the original motivation.

In principle, formalization in this full sense of the word is a possible option for theories of natural language semantics no less – but also no more – than it is for theories about other scientific subjects. And the de facto situation is much the same. Few if any in the profession are concerned with formalization of linguistic theory in this sense, and we won't be either. But let us reflect briefly on what would be involved in such a formalization and on the requirements for the over-all formalism in which such a formalization could be carried out, if only to make fully explicit how such an over-all formalism differs from LFF, both in role and, necessarily in form and logical power. Here is a summary of what the formalization of a model-theoretic approach to the semantics of a natural language or natural language fragment OL would have to include:

(i) a complete formalization of the syntax of the 'object language' OL;
(ii) a formalization of the formalism of which the theory makes use in the way described above;
(iii) a formalization of the notion of a 'model' for LFF – recall the parenthetical remark about a model-theoretic account for LFF in the previous subsection – and (thereby) also for OL.
(iv) a formalization of the semantic value definition, which assigns to each combination of a well-formed term of LFF and a model the semantic value that the term determines in that model.

That is a lot of formalizing, and a powerful formalism is needed within which all this can be carried out. (Expressive power is needed in particular for the formalization of the notion of a model.) Most importantly, the over-all formalism needed in this enterprise has to be distinguished from the LFF that is used in the theory. In fact, formalization in this thoroughgoing sense involves *formalizing LFF within the over-all formalism.* It is this last requirement – the formalizations of the syntax and semantics of LFF within the over-all formalism – that deserves emphasis here. In order to be able to do its job as provider of logical forms for systems as expressively powerful as human languages LFFs will have to be expressively powerful systems too. And that entails (for reasons that we cannot go into here, but which ultimately relate to the incompleteness results of Gödel and Tarski) that the over-all formalism in which the formalization of LFF is to be carried out, must have even greater expressive power. So it is not only that the over-all formalism has a more pervasive role to play than LFF in a complete formalization of the given linguistic theory; the two formalisms will have to be formally distinct, with LFF corresponding to a proper part of the over-all formalism.[3]

As said, we won't engage in complete formalizations of the kind alluded to here any more than anyone else does in Formal Semantics. But, once more, it is important to realize that there is a crucial difference between such thorough formalization and the use that we will be making of LFFs. And it seemed right that this was stated explicitly at least once. To summarize:

The short moral of this section is: One must distinguish 'formal semantics', which makes use of LFFs, from formalized theories in the sense of the Philosophy os Science.

1.2 A bit of a historical and a contemporary perspective

The roots of formal semantics as the discipline exists today can be traced back at least as far as the work of Frege and Peirce in the 19-th Century that led to the formulation of Predicate Logic. (It must be possible to trace it back even farther, but I won't try to do that.)

Predicate Logic offered logicians, mathematicians and philosophers for the first time in history a truly sophisticated and powerful tool for the expression of complex meanings. Now, something became possible that had never been possible before: to formalize intricate mathematical propositions in a form that made it easier to track their logical implications and to detect errors in mathematical or logical reasoning which it is hard or impossible to spot in an unformalized setting. As the practice of putting statements from the mathematical sciences into predicate logic form proved its usefulness, it gained currency within these communities, though much less so, at this point in time, among linguists. In hindsight it is easier to think of the reason why

[3]When complete formalizations are carried out, as we sometimes do in mathematics, the over-all formalism used is typically some version of Set Theory, such as Zermelo-Fraenkel (ZF) or Gödel-Bernays (GB)

the method of ‘formalization into predicate logic form’ did not have much of an impact on the linguistic community. In a nutshell – we will come back to this point below, but for now we state the point as succinctly as possible – what was missing was any systematic account of how linguistic meaning is determined by linguistic form. In this regard the method used by logicians and philosophers had nothing to offer. This is how one would operate. Someone would propose a formalization of some natural language sentence or sentences in predicate logic form and the merits of that proposal would then be judged on the basis of (a) one’s understanding of the given natural language (typically this would be one’s mother tongue) and (b) a proper command of Predicate Logic. Once the formalization has been authenticated on the basis of this kind of information, the formalizing formula can then be investigated, as a stand-in for the sentence it formalizes, for its logical and semantic properties and that can be very useful. But this doesn’t tell us anything about what it is about the grammatical form of the formalized natural language sentence that accounts for why it has the semantic and logical properties it has. And that is precisely what is important to the linguist.

This changed with the work on natural language semantics that was done by Richard Montague in the second half of the nineteen sixties and early seventies. Montague succeeded in doing for natural language – more precisely: for certain fragments of English – what the founders of modern mathematical logic – Frege, Gödel, Tarski and others – had accomplished for formal languages like the First Order Predicate Calculus: Define the meanings of sentences and their syntactic consituents in a systematic way on the basis of their syntactic form. The syntax that Montague used in his model-theoretic treatments of fragments of English was perhaps not exactly what a linguist would want; but he was acutely aware that at least in rough outline syntax should conform to age-old and uncontroversial intuitions about the grammatical structure of complete sentences of languages like English (and, originally, of Latin and Ancient Greek). Subsequent work by Barbara Partee and others turned Montague’s work into a form that made it more attractive and useful to linguists. Since then formal semantics has been mostly pursued within linguistics.

As time went on, it became gradually clearer to formal semanticists how different natural languages really are from the formal languages of symbolic logic. Montague’s famous pronouncement that ’there is no fundamental difference between natural languages and the artificial languages of formal logic’ embodied an important truth when he put it forward in the late sixties: that natural languages depend on a systematic connection between form and meaning no less than logical languages and therefore that they can be described using the same methods that had already been successfully applied in the description of those. Over the intervening years that point has become almost embarrassingly commonplace (largely because of the impact of Montague's contributions); in fast, it now seems so obvious and inevitable that it is hard for us to imagine now that there was a time when it wasn't. But the parallelisms between natural and logic languages only goes so far. There are al sorts of ways in which natural languages differ from the artificial languages that originated on the logicians' drawing boards. This is directly connected with the circumstance that natural languages are used for a much richer spectrum of purposes, and in a much wider range of different contexts, than those for which logical languages have been designed and used. True, mathematics, science and philosophical analysis are among the things you can do with natural language. But these are *very special uses* of human languages. They are uses that for the most part are concerned with the making of statements, the formulation of hypotheses and conjectures and the drawing of inferences. They very heavily involve writing. And they came very late in the historical development of human languages, an have been part of a development that has de facto been restricted to those languages that satisfy, minimally, the precondition of having a writing system and an extensive literate and, ultimately, a scientific tradition.

One aspect of writing, and scientific writing in particular, is that with it a need develops for getting away from *context dependence*. Communication by writing differs from communication through speech in that the production and the interpretation of a verbal message need not occur in the same place and at the same time. The advantages of this are too obvious to need extolling. But there is also a certain disadvantage (if disadvantage it is) in that temporally and spatially distant communications can't rely on the kind of information that is typically shared by people who are talking to each other and who, by necessity, are doing that at the same time and (invariably before the advent of the telephone) in the same place. Written communication cannot take advantage of shared contextual information in the way that spoken conversation can, and does.

To compensate for this absence, uses of language that involve writing have had to develop ways of structuring verbal messages that can stand on their own to an extent that spoken messages typically do not have to. This is true already to some extent when written communication is still personal, as when people write to each other who share a lot of common information and know a lot about each other; but it is true to a much greater extent for written texts that are intended for publication and that are composed by an author who is writing for a wide and diversified audience, about which she often knows very little, and about which she is in fact not entitled to make many assumptions for the simple reason that there isn't all that much in common between the various readers for all of whom the text is meant. Thus writing, and scientific writing perhaps more than any other kind, have developed sophisticated strategies for becoming context-independent, in all sorts of ways that are dispensable in spoken communication. The result has been a highly sophisticated mode of using language, which is just one from a remarkably rich repertoire of different ways we use the languages we speak, and one that from the perspective of the theory of meaning is exceptional in that linguistic form is required to determine meaning pretty without the assistance of non-verbal context, and thus pretty much on its own.

Looked at from this angle the logically motivated and logically oriented study of language that dominated formal semantics when that enterprise first got going, and which still dominates it to a considerable extent today, is, you might say, a kind of perversion. It focuses, originally to the point of excluding all else, on one way of using human languages that is highly specialized and that is something of late addition in the development of human language. If we really want to understand how language works, then this surely is an odd end at which to start.

Here are a few of the aspects and uses of natural language that are negligible in its written scientific uses (or even completely absent) but which have come to be recognized as crucial for an understanding of how human languages work in general. Note well, these are just some of those aspects. The list could be made a good deal longer.

(i) non-indicative speech acts: making statements is only one of the many different purposes to which language can be put. It is also used for doing many other things: making promises, giving advice, ordering people, or giving them permission, to do things, making requests, asking questions, christening, naming, defining, greeting, apologizing, producing exclamation of joy, approval, shame or abhorrence, as well as many things that require a highly conventionalized setting, such as arraigning, sentencing or acquitting in court, appointing and dismissing, marrying and annulling marriages, excommunicating (this last one a curious favorite of speech act theorists). Some of these (but not all of them) involve proposition contents. For instance, when I say to you 'I promise to read your piece by tomorrow', I am making a promise and that promise involves the propositional content that

'I will read your piece by tomorrow'. The established methods of formal semantics are capable of dealing with this propositional dimension of those speech acts that have such a dimension; but as they stand, they are not suitable for dealing with the 'speech act dimension' (i.e. with that which makes a promise a promise, rather than, for instance, a command or a piece of advice).

(ii) Vagueness. Almost all predicate words we find in natural languages are vague to a lesser or greater extent. Up to a point that is an inevitable consequence of the way human cognition conceptualizes the world in which we live: Almost without exception our concepts come with fuzzy edges, and the words that we use to denote those concepts inherit the vagueness of their denotata. But in fact the vagueness of language goes further than that. Often we use expressions with a certain degree of vagueness or looseness not because we are condemned to this by the limitations of our cognition, but because it is convenient to speak loosely or vaguely; speaking loosely or vaguely is a way of conveying enough information for the purpose at hand without conveying too much. The literature on the semantics and logic of vagueness is very large by now, but for the most part it has been curiously divorced from the rest of the semantics literature. In this seminar vagueness won't be a topic either.

(iii) As indicated above, in face to face communication we can rely on all sorts of contextual information that that s not available when we write. Human languages are designed to exploit contextual information by enabling their users to take them for granted and to phrase their utterances in such a way that they can build on it. This practice of building on information that can be glanced from the context in which the utterance takes place can take all sorts of forms. The one that is perhaps best familiar is the way in which a speaker can rely on non-verbal context when using demonstrative phrases like *that bird*, or *that roof* etc. – for instance in an utterance like 'Look at that bird/at the bird on that roof!' Often utterances of demonstrative phrases are accompanied by some kind of gesture that gives the audience a hint of the direction in which the referent of the phrase can be found, and the descriptive content of the phrase will then do what else is needed to make the reference unambiguous. But there are other ways too in which non-verbal contextual information can complete the information provided by the uttered words - various kinds of *presuppositions* that these words rest on and that the context enables the recipient to resolve.

In written communication (and in spoken communication too), the effect of context can to some extent be simulated verbally, by exploiting the 'context' set up in the earlier parts of a discourse or text in the phrasing of the following utterances or sentences. With regard to referring noun phrases this exploitation takes the form of what is usually called*anaphora*. An example is the use of *this book* in

(1.1) The first novel by Irving that I read, was superb, I thought. So then I decided to read another book by him. But this book wasn't nearly as interesting.

By the time the reader of this passage gets to the sentence 'But this book wasn't nearly as interesting.', the preceding sentences have set up a 'discourse context' in which there are two books by Irving, the first of which the author of the passage found superb. The phrase *this book* in the last sentence picks up the last introduced book, which plays the same sort of role here as the bird pointed out by the speaker in the example of the last paragraph. There is also another anaphoric element in this last sentence, viz. the 'comparative construction' *as interesting.* This construction needs a second argument (as interesting as what?), but the wording doesn't give us such an argument. What is evidently intended as second argument is the only other book introduced into the discourse context (the one that the author did like). The interpreter of *as interesting* must and can recover this argument from the context.

Demonstrative reference and anaphora are perhaps the best understood of the phenomena we have mentioned. Anaphora, moreover, has a special significance for formal semantics in that it can be said to have divided the semanticist camp. Some semanticists see the kind of sentence boundary crossing anaphora illustrated in (1.1) as belonging to the domain of semantics, and pursue a *dynamic* approach to semantics, in which the establishment and exploitation of discourse contexts is treated as an integral part of the way in which the semantics of natural languages works. (The approach we will adopt through most of this course is also in this spirit.) However, the majority of semanticists prefer to regard anaphora as one of the many *pragmatic* aspects of the theory of linguistic meaning, which ought to be treated at some 'post-semantic' level where the semantics has already done all its work.

Except for anaphora and the closely related phenomenon of presupposition none of these aspects of language will be discussed in this course (or not at any rate in the principal part of it which will be documented in those notes). In fact, the semantics that we will be concerned with is of the 'old-fashioned' kind which reflects the historically rooted prejudices that I have made an effort to expose in this section. But that isn't quite as bad a thing as I have just made it sound. For it remains true that a very large part of what we can do with language involves, in one way or another, the propositional contents expressed by its clauses and sentences. It is of the first importance to be fully aware that this is only part of the story. Even when this part of the theory of meaning could be brought to a happy conclusion – and we are still very far from that for any language, even English, which has been studied far more closely and extensively than any other language spoken on earth – much would still be left of what the theory of meaning needs to accomplish in toto. The phenomena alluded to in this section, and quite a few others besides, will have to be dealt with as well. (And whether we classify them as 'semantic' or 'pragmatic' makes of course no difference to that.)

To end this on a happy note: There will be lots and lots for all of us to do, and for a very long time to come!

1.3 Truth Conditions and Compositionality. Comparison of some formal approaches

This section is meant to serve two purposes at once:

(i) To illustrate some of the semantic problems that arise even when we look at reference, predication and quantifcation in an atemporal setting; and

(ii) to give a sample of a number of different formal methods that have been used to deal with these problems.

(For someone who has had no exposure to formal semantics before this will be a lot all at once. But we will soon restrict attention to one of these methods, and we will get plenty of practice with that. There is no reason for panic.)

We start with some examples whose analysis allows us to focus on the basic concepts of all systematic semantics of natural language: Reference, Predication and Quantification.

Before the advent of formal natural language semantics as we know it today – that is, before the mid-sixties of the last century – logicians, some philosophers and the occasional linguist would make use of Predicate Logic to spell out the truth conditions of natural language sentences. And indeed, Predicate Logic is a good formalism for stating the truth conditions that are based on just these three mechanisms.

Here is a pair of examples:

(1.2) a. Two languages are spoken by everybody in this room.
b. Everybody in this room speaks two languages.

(Example originally due to N. Chomsky)

Questions relating to (1.2):

(i) What is the intuitive difference between (1.2.a) and (1.2.b) as far as their truth conditions are concerned?
(ii) How do we represent the truth conditions of these sentences in first order predicate logic?
(iii) How can we account for these truth conditions in a systematic, compositional way? Can we use predicate logic for this purpose?

Here is an intuitive judgment of what these sentences can mean: (1.2.a) is genuinely ambiguous between (i) a reading in which the subject has wide scope over the *by*-phrase – there are two particular languages such that everyone speaks both of these languages – (ii) a reading according to which everyone in the room speaks two languages but these languages need not be the same for everyone. For (1.2.b) this second reading seems the dominant one. But it seems hard to exclude the first reading completely. One reason why it doesn't seem a good idea to exclude the first reading as a possible reading for (1.2.b) is that a continuation like that in (1.3) seems to be perfectly good for (1.2.b) (and no less than for (1.2.a)).

(1.3) a. Two languages are spoken by everybody in this room. They are Chinese and English.
b. Everybody in this room speaks two languages. They are Chinese and English.

Representing the different readings of (1.2.a) and (1.2.b) in Predicate Logic is quite straightforward. (In fact, these readings are just the sort of thing in Natural Language Semantics for which the Predicate Calculus is well-suited.) Predicate Calculus formulas identifying the two readings are given in (1.4.a,b).

(1.4) a. $(\forall x)((\mathrm{P}(x) \ \&\ \mathrm{R}(x)) \rightarrow (\exists y)(\exists z)(\ \mathrm{L}(y) \ \&\ \mathrm{L}(z) \ \&\ y \neq z \ \&\ \mathrm{S}(x,y) \ \&\ \mathrm{S}(x,z)))$

b. $(\exists y)(\exists z)(\mathrm{L}(y)\ \&\ \mathrm{L}(z)\ \&\ y \neq z\ \&\ (\forall x)((\mathrm{P}(x)\ \&\ \mathrm{R}(x)) \rightarrow (\mathrm{S}(x,y)\ \&\ \mathrm{S}(x,z))))$

(With 'P' for 'person','R' for 'in the room', 'L' for 'language' and 'S' for 'speaks'.)

Someone with a reasonable amount of practice in using the Predicate calculus for such purposes and an acute enough sense of English to understand what the two readings of the sentences in (1.2) are will need no further argument that the formulas in (1.4) capture those readings correctly. But further support for this could be given by going through an analysis of the truth conditions of these formulas that consists in applying the *truth definition* for Predicate Logic to them. (This will be familiar to all or most of you. But it may serve as a little refresher of some of the fundamental features of model-theoretic semantics that will be part of the foundations on which we will be building in this class.)

The model-theoretic method in semantics consists in (i) specifying for the given formal or natural language for which a semantics is being sought a class of *models*- abstract structures that play the part of possible situations or worlds that the sentences of the language can be thought to be about and in relation to which they will be true (when they describe the modeled situation correctly) or false (when they do not). The models specify all the information that is needed to determine whether any given sentence of the language is true or false in the worlds or situations they represent. If the language is Predicate Logic, then this information will consist of (i) a *Universe* U consisting of all the individuals in the modeled situation or wold: (ii) extensions relative to U for all the predicates of the language (subsets of U for the 1-place predicates, sets of ordered pairs of elements of U for the 2-place predicates and so on; often one makes use of the *characteristic functions* of these extensions).

The *Truth Definition* relates sentences of the language to models in determining which sentences are true in any given model M and which are false. The information on which the Truth definition relies is on the one hand the information in the model M and on the other the recursive syntactic structure of the sentence. The dependence on the form of the sentence is *compositional*: the relation in which the sentence or syntactic part of it to M is directly determined by the relations to M of its immediate syntactic constituents. In the case of the Predicate Calculus this compositional dependency (of the relations of compound expressions on those of their constituents) is comparatively simple since all syntactic constituents of a formula a rev also formulas. But there is nevertheless a complication, which has to do with free variables. The immediate constituent of a sentence beginning with a quantifier is normally a formula which has free occurrences of the variable which that quantifier binds; and if we go down farther into syntactic structure of a complex sentence we may encounter sub formulas with any number of free variables. But a formula with free variables isn't simply true or false in a model M; it can be regarded as true or false only, given some assignment of objects from the universe U of M to those variables. To deal with this complication it is customary to state the Truth Definition as dependent on assignments of objects from U to all the variables, and to state the dependence of the truth value in M of a quantified formula as one that involves quantification over assignments. Thus the Truth Definition will contain as clauses for existentially and universally quantified formulas clauses that can be stated as in (1.5).

(1.5) a. $[(\exists x)\phi(x)]^{M,\mathbf{a}} = 1$ iff there is a d in U such that $[\phi(x)]^{M,\mathbf{a}[d/x]} = 1$;

b. $[(\forall x)\phi(x)]^{M,\mathbf{a}} = 1$ iff for all d in U, $[\phi(x)]^{M,\mathbf{a}[d/x]} = 1$

(Here $\mathbf{a}[d/x]$ is the assignment that is just like $\mathbf{a}$, except that d is the value that it assigns to the variable x.)

Using these 'truth clauses' for the existentially and universally quantified sub-formulas of (1.4.a) (where (1.4.a) itself is also counted as one of its own sub-formulas) and clauses for formulas whose main connective is & or $\rightarrow$ (which we do not state separately here) we get, through repeated application of the various clauses, to the following truth conditions for (1.4.a):

$[[(\forall x)((\mathrm{P}(x)\ \&\ \mathrm{R}(x)) \rightarrow (\exists y)(\exists z)(\ \mathrm{L}(y)\ \&\ \mathrm{L}(z)\ \&\ y \neq z\ \&\ \mathrm{S}(x,y)\ \&\ \mathrm{S}(x,z)))]]^{M,\mathbf{a}} = 1$ iff

for all d in U, $[[(\mathrm{P}(x)\ \&\ \mathrm{R}(x)) \rightarrow (\exists y)(\exists z)(\ \mathrm{L}(y)\ \&\ \mathrm{L}(z)\ \&\ y \neq z\ \&\ \mathrm{S}(x,y)\ \&\ \mathrm{S}(x,z))]]^{M,\mathbf{a}[d/x]} = 1$ iff

for all d in U, either $[[(\mathrm{P}(x)\ \&\ \mathrm{R}(x))]]^{M,\mathbf{a}[d/x]} = 0$ or $[[(\exists y)(\exists z)(\ \mathrm{L}(y)\ \&\ \mathrm{L}(z)\ \&\ y \neq z\ \&\ \mathrm{S}(x,y)\ \&\ \mathrm{S}(x,z))]]^{M,\mathbf{a}[d/x]} = 1$ iff

for all d in U, either $[[(\mathrm{P}(x)\ \&\ \mathrm{R}(x))]]^{M,\mathbf{a}[d/x]} = 0$ or there are e, f in U such that $[[(\mathrm{L}(y)\ \&\ \mathrm{L}(z)\ \&\ y \neq z\ \&\ \mathrm{S}(x,y)\ \&\ \mathrm{S}(x,z))]]^{M,\mathbf{a}[d/x][e/y][f/z]} = 1$ iff for all d in U, either $[[\mathrm{P}(x)]]^{M,\mathbf{a}[d/x]} = 0$ or $[[\ \mathrm{R}(x))]]^{M,\mathbf{a}[d/x]} = 0$ or there are e, f in U such that $[[\mathrm{L}(y)]]^{M,\mathbf{a}[d/x][e/y][f/z]} = 1$ and $[[\ \mathrm{L}(z)\ ]]^{M,\mathbf{a}[d/x][e/y][f/z]} = 1$ and $e \neq f$ and $[[\mathrm{S}(x,y)\ ]]^{M,\mathbf{a}[d/x][e/y][f/z]} = 1$ and $[[\ \mathrm{S}(x,z)]]^{M,\mathbf{a}[d/x][e/y][f/z]} = 1$.

Whether this analysis can do much to bolster anyone's confidence that (1.4.a) correctly captures one of the possible readings of the sentences in (1.2) may be debatable. For one thing, the final outcome of the analysis – the last line starting with 'for all d in U' – seems to do no more than repeat, in a rather prolix and opaque fashion what (1.4.a) expresses in much more concise and surveyable form. And that is because the notation of Predicate Logic has a kind of 'optimal semantic transparency' – the syntax of its formulas is a perfect image of their meaning. With expressions of natural languages this is in general not so. Here a formal model-theoretic analysis can often be revealing in a way that it can never be for perfectly transparent formal languages like Predicate Logic.

So if the point of this exercise is just to strengthen our conviction that (1.4.a) correctly represents one of the readings of the sentences in (1.2), the exercise can hardly be called a success. But the point was not so much that, but rather to put is in mind again of how the model theory for Predicate Logic works and in particular what role is played by assignments in the Truth Definition. Note well in this connection that the assignments used in the set-up summarized above are all *total* assignments – they always assign values to *all* the variables of the calculus. But this is not the only way in which the Truth Definition can be set up. Note that no formula of the calculus will have more than a finite number of free variables, and that in the semantic evaluation of a formula only assignments are needed to the free variables that it contains. So it ought to be possible to make the Truth definition do its work while only considering *partial finite assignments* each of which is defined only for some finite set of variables (although the size of the set may vary in the course of a truth evaluation along the lines of the reduction of (1.4.a) above). In natural language semantics the use of finite assignments has gradually gained way. For instance, find it now both in the version of model-theoretic semantics presented in (Heim & Kratzer 1998) (henceforth 'H&K') and in the different versions of DRT, including the one that we will be using here.

To restate the Truth Definition sketched above in such a way that it only makes use of finite assignments is somewhat awkward (and that is the reason why the inventor of assignment-based truth definitions for languages with quantifier, Alfred Tarski, opted for total assignments). But both in the set-up of H&K and in DRT the use of finite assignments emerges naturally.

This is all that needs saying here about question (ii) and we now turn to question (iii). Note that in our exploration of the usefulness of Predicate Logic in connection with the sentences in (1.2) one consideration was completely side-lined: How do we get to the formulas in (1.4) from the syntactic form of the English sentences in (1.2)? This is a question that was never properly addressed – and, it seems, wasn't even recognized as a relevant scientific question, because those familiar with the potential of formal logic thought that natural languages were just too unsystematic; there wasn't a systematic connection to be discovered and described in the first place. The breakthrough occurred only two thirds down the 20th Century through the work of Richard Montague. Montague saw that the syntax of significant portions of English could be stated in exact terms, which, like the syntax of Predicate Logic and other artificial logical languages such as the typed lambda calculus, could serve as input to a model-theoretic semantics. His way laid the foundations of Formal Semantics of Natural Language. largely as we know it today.

It will take us a little time to find our way to the particular kind of formal semantics that we will make use of through most of this course. Along this path we will have to cope with two major issues:

(Ia) We will have to opt for a syntax for the parts of English that we will be dealing with (those which display the semantic phenomena on which we want to focus); it is from the syntactic analyses that this syntax delivers that our semantics will derive its logical forms.

(Ib) We will have to make certain decisions about the *ontology* of our semantics. By *ontology* we understand roughly the same as what is made of this term in Artificial Intelligence: A theory of the fundamental ingredients of the world, as seen from the perspective of human cognition, as it shows itself in the languages we speak; this theory must specify the range of different kinds of entities that populate this world and the general formal properties they and relations in which they stand to each other. In a model-theoretic approach ontology makes itself felt in two related but quite different ways:

(i) in the way the model theory defines its *model class*. Each model-theoretic semantics must specify the class of models which serve as abstract counterparts of the different possible situations or worlds that the sentences the semantics deals with are taken to be able to talk about. This requires on the one hand that these models specify all the information that is needed to evaluate those sentences for truth and falsity and on the other that only those models are included in the class in which the different kinds of entities have the properties and stand in the relations that the ontology assumes. The first requirement is satisfied by giving the models the appropriate formal structure: they must thane elements representing the different kinds of entities assumed by the ontology and representatives of the various properties and relations that the ontology takes those entities to have and stand in. And the second requirement – that models are excluded from the class in which the properties or relations do not satisfy the right general constraints that the ontology assumes – is secured by laying down certain postulates (often referred to as *Meaning Postulates*), which express those constraints: only models are admitted which satisfy all the Meaning Postulates.

(ii) in the design of the logical form formalism. We noted above that whether a model-theoretic treatment of a language sees its task to be the specification of semantic values for the expressions of its Object Language or as the specification of logical forms for those expressions, the use of a Logical Form Formalism is in practice unavoidable, which has its own syntax and semantics, like we have for the Predicate Calculus. But how should that formalism be chosen? Here too, ontology has its say. Part of the conceptions it tries to capture is what the sentences of OL really say about their subject matter, contrary perhaps to what meets the eye at their surface. (The possibility of this – of a certain logical opacity of natural language, a possibility that sets natural languages apart from the 'logically transparent' languages of formal logic – has become one of the driving concerns in natural language semantics over the years.) The design of the LFF for the given model theory mist of course dovetail with the definition of the theory's models. For given the way the theory is set up, it is the expressions of LFF that must get their values in those models; the expressions of OL get their values via the LFFF expressions that the model theory assigns to them.

Both these issues, the syntax that is to be assumed for OL and the ontology as it manifests itself in the choice of model class, are fundamental;. But nevertheless they won't be all that prominent in what we will be doing in this class. As far as ontology is concerned, thous is because enough of it will be explicit through the LFF we will be using. As regards OL syntax the matter is more problematic. One of the difficulties that formal semantics has had to cope with ever since it developed, through the contributions and championing of Barbara Partee, into a syntax-semantics interface that linguists with a concern for and detailed knowledge of syntax could see the importance of and could take seriously is that there have been too many syntactic theories on the market. And often these theories haven't reached the degree of precision and explicates that a syntax must have if it can fulfill its side of the bargain in the kind of syntax-semantic interface that formal semantics pursues. This means that if the model-theoretic method is to proceed according to its own canons of explicitness, it should either adopt one of the syntactic theories on the market (and make its analyses fully explicit, or extend them to the constructions the given semantics wants to address, to the extent that that is necessary); but that leaves out the syntactic competitors and may create the unwanted and false impression that the semantics provided by the treatment is wedded to the chosen syntax in a way that it is not. An alternative approach that has been tried in the past is to provide a flexible, 'general purpose', syntax-semantics interface which can be, without too much additional effort, customized to the individual user's preferred syntax. But that is not without its problems, for one thing because the major syntactic alternatives that are available offer various different kinds of plugs or holes for the semantics to be slotted into. So quite a bit of work is left by these 'general purpose' approaches for the different customizers.

Fortunately this situation has improved somewhat over the past couple of decades through the advent of syntactic tree banks: careful syntactic annotations of large numbers of sentences from natural language corpora (such as the Wall Street Journal), which despite their flaws have set a widely accepted standard for what the syntactic structures of individual sentences ought to be like. One important advantage of this kind of approach is that it separates the question what syntactic structures should be assigned to the sentences of a language from questions that have to do with the deeper motivations and principles that, in some sense, explain why the syntactic structures are the way they are. the policy we will follow in this course is inspired, in a rough way, by this more recent development. We will adopt syntactic analyses for sentences as we go along, for the most part without reflecting too deeply on the underlying syntactic principles. By and large our syntactic structures will be those of the Chomskyan tradition. In that regard we will not differ from the approach taken in H&K. But the more systematic justification for those structures, in the form of an independently motivated syntactic theory, will be touched upon only incidentally.

We now turn in earnest to question (iii): How to derive Logical Forms, and the semantic values these determined in models from the model class of our semantics, from independently motivated/plausible syntactic structures. To that end it will help t begin by looking at some sentences that are simpler than those in (1.2), such as those in (1.6).

(1.6) a. Every man suffers.

b. John suffers.

c. John loves a dachshund.

d. Every man loves a dachshund.

e. John offended every linguist.

Predicate Logic is not a particularly good tool when it comes to constructing a systematic compositional semantics for a natural language. More effective is some form of the lambda calculus. (This is one of the decisive insights from the father of Formal Semantics, Richard Montague. Montague developed his own 'intensional' version of the lambda calculus, which has proved particularly useful for doing natural language semantics and he used it, among other things, for that purpose. For an influential contemporary approach to natural language semantics in which the lambda calculus is used, see in particular the use of this formalism in (Heim & Kratzer 1998).) This allows us to assign different syntactic categories logical types that reflect their semantics. We will look at the formal details as we go along.

We start with (1.6.a) and first look at this sentence from a perspective that follows Montague and, more specifically H&K. The ingredients we need are (a) a syntactic structure for this sentence and (b) something about the LFF that both Montague and H&K make use of. As regards (a), we assume the following (quite uncontroversial) syntactic structure for (1.6.a).

(1.7)

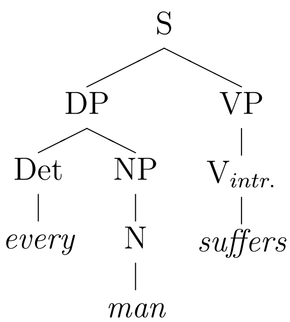

as regards (b), we start with a simplified version of the LFF that Montague addicted in his work on natural language semantics his Higher Order Intensional Logic (HOIL). The simplified version we will be using right now is the one also adopted in H&K. this is the extensional typed lambda-calculus. i will assume familiarity with this system, and in particular with its type declaration, which, starting from the basic types e and t, builds a space of function types by forming ordered pairs out of types already defined.

Montague's basic intuition about the denotations of quantifying noun phrases was that they must be of type <<e,t>,t> (or, in simplified notation, <et,t>). This works fine when such phrases occur in subject position, as in (1.6.a). For instance the subject DP of that sentence can be assigned the following semantics:

(1.8) [*every man*] $\rightsquigarrow \lambda Q.(\forall y)(\text{man'}(y) \rightarrow Q(y))$

and *every*, as an operator which acts on expressions of type <e,t>, such as the NP *man*, should then be given the obviously corresponding value by its semantic lexical entry:

(1.9) [*every*] $\rightsquigarrow \lambda P.\lambda Q.(\forall y)(\text{P}(y) \rightarrow Q(y))$

This will work out fine for a sentence like (1.6.a): the semantics of *every* is first applied to the semantic value man' of the noun *man* (which is of the right type, viz. <e,t>) so as to get the value $\lambda Q.(\forall y)(\text{man'}(y) \rightarrow Q(y))$, and this value is then applied to the semantics suffer' of the intransitive verb *suffer* to obtain $(\forall y)(\text{man'}(y) \rightarrow \text{suffer'}(y))$, which captures the truth conditions of (1.6.a) correctly.

But these semantic values for *every*-phrases do not work (or at least not in any straightforward way), when they occur in other positions. Consider for instance (1.6.e)

(1.6.e) John offended every linguist.

Here the object DP must combine with a two-place relation (an entity of type <e,<e,t>>, or, simplified, <e,et>) rather than with one of type <e,t>. With the DP value given above functional application isn't possible.
One way to try and deal with this problem is to assign to quantifying DPs in object position a different denotation. For instance, *every linguist* could be given the denotation:

(1.10)[*every linguist*] $\rightsquigarrow \lambda R.\lambda u.(\forall y)(\text{linguist'}(y) \rightarrow R(u, y))$

This denotation for *every linguist* can be obtained by providing *every* with the alternative lexical entry:

(1.11)[*every*] $\rightsquigarrow \lambda P.\lambda R.\lambda u.(\forall y)(\text{P}(y) \rightarrow R(u, y))$

<u>Exercise</u>: Show that this gives the right truth conditions for (1.6.e) if we assume the following syntactic tree for the sentence:

(1.12)

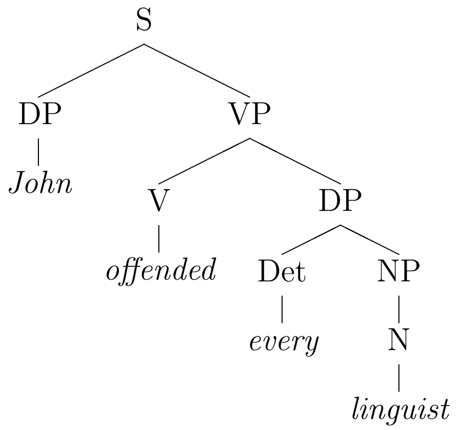

But what about yet other positions in which *every*-DPs occur? What, for instance when they occupy the innermost argument position of a ditransitive verb? In the light of what we have been saying neither of the two entries for *every* we have thus far proposed it seems clear that if we continue along the road on which we have embarked we will need yet another entry in which the second λ-abstract is over ternary relations (i.e. binds a variable of type <e,<e, <et>>>). If DPs can occur as arguments of words that function as predicates of even more than three arguments, then the *every* of an *every*-DP occupying the innermost argument position of such a predicate will yet another entry. This proliferation of different lexical entires, which may be needed only because to the particular strategy we have adopted here for dealing with quantifying DPs, seems suspect. perhaps we can find a way to avoid it.

The problem becomes worse once we accept that sentences with two or more quantifying DP can be 'scopally ambiguous'. Consider for instance, (1.13). It seems a plausible claim – in fact, this is now widely accepted among linguists, partly for reasons that we cannot go into here – that this sentence has both a $\exists\forall$-reading and a $\forall\exists$-reading.

(1.13)Some philosopher offended every linguist.

How can we get the $\forall\exists$'-reading for (1.13), if we insist, as we have been up to now, on treating *every linguist in situ*, so that its semantics must be combined with that of the verb before the result of that is combined with the semantics of the subject? Well, there is a way in which this can be done, so long as we are prepared to accept even more baroque alternative entries for *every* than we have already found it necessary to postulate. Here is the entry that does the trick:

(1.14)$[every] \rightsquigarrow \lambda P.\lambda R.\lambda \mathcal{Q}.(\forall y)(\mathrm{P}(y) \rightarrow \mathcal{Q}(\lambda u.R(u,y)))$

> (P is a variable of type <et>, R a variable of type <e,<et>> and $\mathcal{Q}$ a variable of type <<et>,t>.)

<u>Exercise</u> Compute the semantic value of (1.13) using the entry in (1.14) for *every*.

This succeeds in getting the truth conditions we wanted. But it feels like a Pyrrhic victory. One of the tricks it involves is reversing the function argument relation between subject DP and VP: contrary to what we saw in the previous examples involving quantifying DPs, the subject now plays the part of argument and the VP that of function applying to it. This not only feels rather ad hoc; it also seems to undo what has often been hailed as one of the signal achievements of Montague Grammar (one possible option of which we are pursuing in this discussion of *every*): Montague recuperated a uniformity in the treatment of nominal argument phrases – 'DPs' in the terminology we are using – that seems to have been lost with Frege's treatment of quantifier phrases like *every man* sand 'referential phrases like *Socrates* or *John*, which he saw as making a fundamentally different contributions to logical form, by treating all DP's as having type <<et>,t>. By doing so he could, it seemed maintain that at a deep level the semantic relation between subject and 'predicate' (i.e. the VP) is after all always the same - that of the subject acting as a function that is applied to the 'predicate'. But if we now admit analyses of what seem to be quite ordinary sentences in which this relation is reversed once more, it feels we are throwing an attractive and central feature of Montague Grammar to the wind. This doesn't look like the right way to go.[4]

[4] Perhaps this proliferation of lexical entries for *every* (and other quantifying determiners with it) isn't quite as objectionable as the discussion up to this point has it. Perhaps there is just one 'basic' entry for *every* – for instance the one given in (1.9) – from which all others that might be needed to deal with other sentences or sentence interpretations can be derived in some systematic way. Proposals along these lines have been made before and some of them seem appealing and have found wide approval. The first to come up with proposals of this kind were Partee and Rooth, who applied such a strategy to the apparently distinct semantic roles that it is necessary to the verb *to be*. Partee and Rooth cast the term *type shifting* for this method of obtaining derived lexical entries from basic ones. (See (Partee & Rooth 1983), (Partee 1972).)

An alternative method for dealing with this kind of problem is proposed by Heim and Kratzer. It involves a level of syntactic processing that is intermediate between the construction of the syntactic structure that directly captures the surface structure of the input expression as string of words and the level at which semantic value is assigned, and the assignment function operates on the output of the intermediate level of syntactic processing, and not on its input. The output of this syntactic processing level is usually referred to as 'LF' (to remind us of the notion of 'logical form' as it had long been used in a more or less informal mode, but not to be regarded as an abbreviation of this 'logicians' notion).[5] The part of this intermediate processing that is crucial for our present problem is a syntactic operation known as 'Quantifier Raising', in which a quantifying DP is moved from its position in surface structure to some higher position. It is this higher position which determines the scope of the quantifier that the quantifying DP contributes to the semantics.[6]

An example is the analysis H&K offer for (1.3.e).

(1.6.e) John offended every linguist.

The tree for the surface structure of this sentence

However, before endorsing any application of the type shifting strategy we should make sure (a) that there are no good solutions to the given problems which succeed without resorting to this strategy and (b) that there is an explicit statement of the type shifting application, which specifies besides the 'basic' entry (i) exactly what the shifting operations are – what there input is and what the corresponding output – and (ii) what triggers these operations are triggered (e.g. what feature of the syntactic configuration and the semantics of the constituents that prevents the application of the basic operation). Thus a lot more should be said and done before a type shifting approach should be adopted to the problems with *every* we have just been looking at. Especially in view of the last case, involving the entry (1.14), it is very hard to see how a general type shifting account could be formulated.

[5] In earlier days the syntactic level inputs of the inputs to the construction of LFs was referred to as 'Surface Structure' (SF). This is so in particular in Chomsky's *Theory of Government and Binding* (see in particular (Chomsky 1982), (?)). The term 'Surface Structure' was abandoned in the syntax model known as *Minimalism*, which Chomsky proposed as replacement for the *Theory of Government and Binding* (see (Chomsky 1995)) The term 'LF', though, was retained in this later model.

[6] The proposal to use Quantifier Raising in order to get scope relations involving quantifying noun phrases right was first made by Robert May in his dissertation (May 1977). See also the somewhat later (May 1985)

(1.15)

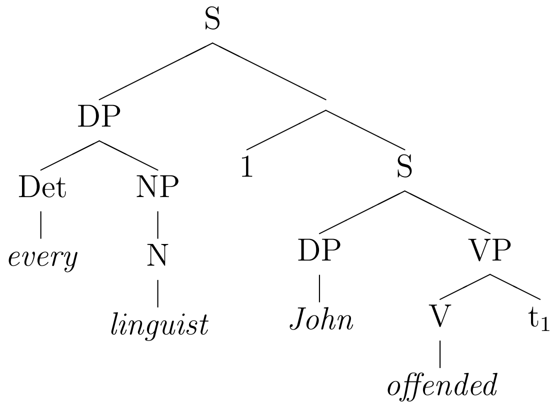

To make the semantics work for this syntactic analysis of (1.6.e) H&K add a notational device to the version of the lambda calculus they use. It looks like this: $[[\alpha]]^{\mathbf{a}}$.

Here α is the syntactic structure of an expression (at the level of LF) and **a** is a *partial assignment* of individuals to numbers.[7] For instance, **a** could be the assignment with Domain {1} (i.e. the assignment which is only defined for the number 1) and be such that **a**(1) is the individual b. (H&K use the special notation '$1 \rightarrow$ b' to denote this assignment.) Or **a** could be the extension with Domain {1,2} which assigns c to 1 and b to 2 and which in the special notation is denoted as '$1 \rightarrow$ c, $2 \rightarrow$ b' and so on. The 'empty function', the function with empty Domain $\emptyset$, is also counted as an assignment (and also denoted as $\emptyset$).

This notation becomes relevant only in cases where the syntactic tree contains *indices*; these are subscripts like the '1' in t_1 in (1.15). For instance, the term $[[t_1]]^{\mathbf{a}}$ is defined when, and only when, the index 1 belongs to the Domain of **a**, and then its value is the value that **a** assigns to 1. Thus $[[t_1]]^{1\rightarrow b} =$ b; and $[[\text{linguist}(t_1)]]^{1\rightarrow b}$ denotes the truth value 1 iff $[[\text{linguist}]](\text{b}) = 1$ (that is, iff b belongs to the extension of 'linguist' – the term we use to denote the extension of the English noun *linguist* – or, even more colloquially, iff b is a linguist).

Using this notation we can designate the semantic value of the VP of (1.15)

[7]H&K assume that there is a set D_e of 'individuals' (or 'objects'), the set of individuals existing in the world, that with this set D_e and the set {0,1} of truth values the world also contains all the 'higher type' domains D_β (where β is any higher type). (One such domain is the domain $D_{<e,t>}$ consisting of all functions from the set of individuals to the set of truth values.) These functions are the characteristic functions of potential candidates for extensions of 1-place predicates (given that D_e is the set of all individuals).

under the assignment $1 \rightarrow$ b as '[[λu. offended(u,t_1)]]$^{1\rightarrow b}$'. This expression is well-defined and denotes the function which to any individual c from D_1 assigns the truth value 1 if c offended b and the value 0 otherwise.

The notation also enables us to write down a term such as '[[λu. offended(u,t_1)]]$^{\emptyset}$'. But this term doesn't have a proper denotation, since $\emptyset$ doesn't assign a value to the trace t_1. This is different for '[[offended]]$^{\emptyset}$' or '[[linguist]]$^{\emptyset}$'. In these terms the expressions between the square brackets do not contain any indices. So they can be evaluated without value assignments to any indices. (It is convenient to permit omission of explicit reference to the empty assignment. So '[[linguist]]' is just short for '[[linguist]]$^{\emptyset}$'.)

In H&K's system Quantifier Raising always introduces a trigger for λ-abstraction at the level of interpretation. This trigger is the index that is adjoined just below the moved DP. To see how this works let us consider as simple a case as possible: that of sentence (1.3.a):

(1.6.a) Every man suffers.

Since DPs are always raised in the system, even in cases where this makes no difference to scope interpretation, the syntactic representation at LF is that in (1.16).

(1.16)

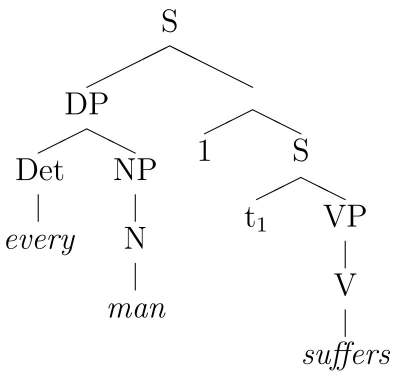

Consider the right constituent of the higher S node:

(1.17)

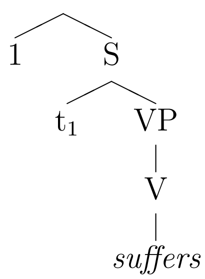

The value of this constituent under any assignment **a** can be informally described as follows: the left daughter of the tree indicates that we must, in order to obtain this value, form a λ-abstract over the term which denotes the value of the right constituent of (1.17) – that is: of the sentence structure with the trace t_1 in subject position and the VP *offended* – with respect to the value assigned to the index 1 (the left constituent of (1.17)). But this assignment will not be **a** itself (except in exceptional, 'accidental' cases, which are irrelevant to this discussion), but a modification or, more typically, an extension **a'** of **a** which has 1 in its Domain and assigns some individual u to 1. That is, the Domain of **a'** equals Dom(**a** $\cup$ $\{1\}$ and **a'**$(1) =$ u. We can use H&K's special notation to denote this assignment as '**a** $\cup$ $\{1 \rightarrow$ u$\}$'. But the crucial point about this last term is that in it 'u' is a variable (an individual variable of the 'metalanguage' – i.e. a variable ranging over the individuals in D_e, which belongs to the language in which we talk about the semantics of the Object Language' (here: some fragment of English that contains among others the sentences under consideration).[8]

This means that the term (1.18) is an 'open' term of the meta-language, which contains the free variable u. So this term itself will not denote a semantic value, since there is nothing that will assign a value to this variable.

(1.18) 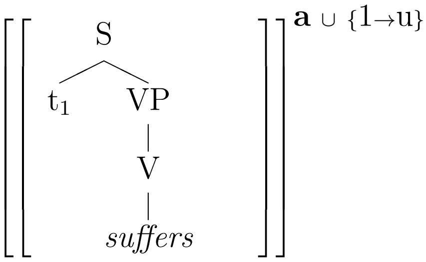

But this is no longer so when we λ-abstract over 'u'. The result of that is a function which maps each object u onto the value given by (1.18). We can denote this function as in (1.19).

[8]The LF structures of sentences and their constituents, to which the semantic theory assigns semantic values, do not contain variables in the logical sense of this term. Many of them contain indices and traces. But neither of these are variables in the sense of formal logic and mathematics. However, the adjoined indices trigger the introduction of variables into the terms of the meta-language that serve as vehicles for the determination of the semantic values and the theory is set up in such a way that each trace is bound by the adjoined trace that it bears, and thereby by the quantifying DP which left the trace behind as it was moved.

(1.19) λu. $\left[\!\!\left[\begin{array}{ccc} & \text{S} & \\ \text{t}_1 & & \text{VP} \\ & & | \\ & & \text{V} \\ & & | \\ & & \textit{suffers} \end{array} \right]\!\!\right]^{\mathbf{a} \,\cup\, \{1 \to \text{u}\}}$

(1.19) is the value of the tree in 1.17) under the assignment **a** – this is how in H&K the semantics of (1.17) is defined in terms of the semantics of (1.18).

Assuming that the lexical entry for the verb form *suffers* is as in (1.20) and making use of the rule that for any u, $[[t_1]]^{\mathbf{a} \,\cup\, \{1 \to \text{u}\}} = \text{u}$, it is not hard to see that the function term in (1.19) can be rewritten as in (1.21):

(1.20)*suffers* $\rightsquigarrow \lambda v.\text{suffers'}(v)$

(1.21) $\lambda u.\text{suffers'}(u)$

To get the value of the entire sentence (1.6.a) under assignment **a** we need to apply the semantic value of the raised subject DP *every man* under **a** to (1.21). The value of this DP under **a** is what we get when we apply the original entry (1.9) for *every'* to the value of *man'*. Here, both *every'* and *man'* are constants of the meta language which denote the semantic values of the English words *every* and *man*, so their denotations are assignment-independent; and the same is therefore true of the result that we get when we apply the denotation of the first to the denotation of the second. In other words, whatever **a** may be, we have the following identity:

(1.22)

$$\left[\!\!\left[\begin{array}{ccc} & \text{DP} & \\ \text{Det} & & \text{NP} \\ | & & | \\ \textit{every} & & \text{N} \\ & & | \\ & & \textit{linguist} \end{array} \right]\!\!\right]^{\mathbf{a}} \quad = \quad \lambda Q.(\forall y)(\text{man'}(y) \to Q(y))$$

The semantic value of sentence (1.6.a) (under assignment **a**) is now the result of applying the right hand side of (1.22) to (1.21). That is, we get:

$$(1.23)(\lambda Q.(\forall y)(\text{man'}(y) \rightarrow Q(y)))(\ \lambda u.\text{suffers'}(u)) = \\ (\forall y)(\text{man'}(y) \rightarrow \text{suffers'}(y))$$

Note that this value has nothing to do with the assignment **a** – **a** isn't mentioned anywhere in the specification of the value. This is a general feature of the values of LF structures without traces. All such LFs have values that are assignment-independent. This is so in particular for (the LFs of) complete sentences: these are either true under all assignments or else false under all assignments. Another way of putting this is that the values of such (trace-free) expressions are defined for the empty assignment $\emptyset$ (and therewith for all assignments that extend $\emptyset$, since every non-empty assignment is an extension of $\emptyset$).

Let us, for good measure, go through the computation of the semantic value for one simple sentence in detail, starting with its smallest constituents and working our way up step by step through the its syntactic tree to its highest S node. That will make it necessary and possible to so see in full detail what is involved in such calculations. The sentence I have picked for this purpose is (1.6.a). Among the things that this exercise will show is how small the assignments can be that are needed in such a calculation. In the case of (1.6.a) the only assignments needed are the empty assignment and assignments whose Domain consist just of the index 1.

Step 1.

The semantic value of the V node in the logical form (1.16) for (1.6.a) that immediately dominates the lexical item *suffers* is determined by the lexical entry for *suffers* given in (1.20). Put more exactly, this value is determined for all assignments, including the empty assignment $\emptyset$. In the representing term displayed in (1.24) the assignment chosen is the minimal assignment $\emptyset$.

$$(1.24)\left[\left[\begin{array}{c}\text{V}\\ | \\ \textit{suffers}\end{array}\right]\right]^{\emptyset} = \lambda u.\text{suffers'}(u)\ (= \text{suffers'})$$

(N.B. the last equality of (1.24) is justified by the fact that the meta-linguistic constant *suffers'* must denote a function from individuals to truth values. We can turn this function constant into a formula by combining it with the variable u as argument. If we then lambda-abstract over u we again get a function denoting term – the one between the two = signs in (1.24) – and the function it denotes is the same function as the one denoted by '*suffers'*'.)

Step 2.

Nothing happens when we pas from the V node in (1.16) to the VP node above it, which has V as its only daughter. Thgis configuration entails that the value in (1.24) is passed on unaltered as semantic value of the structure dominated by VP:

$$(1.25)\left[\!\!\left[\begin{array}{c}\text{VP}\\ |\\ \text{V}\\ |\\ \textit{suffers}\end{array}\right]\!\!\right]^{\emptyset} = \lambda u.\text{suffers'}(u)$$

Step 3.

The third step of our computation computes the value of the trace t_1. The trace can be evaluated only with respect to assignments that are defined for the index 1. Here again we proceed minimally, by considering assignments that are defined just for this index and no others. And we make use of the device explained above: of an assignment-denoting term of the metalanguage which contains a variable for the denoted value; that is, the superscript of our value term that we use is theassignment-denotting term '$1 \to v$', where v is a (meta-)variable for the value assigned to 1. (We could also have used the meta-variable u, like we did earlier. But in what follows a different variable will make things a little more perspicuous.) What we get – the value of t_1 under the assignment that assigns v to the index 1 – is just whatever object is denoted by v, in other words, v. (However, as v is a variable, the term doesn't actually tell us which object that is).

$$(1.26)[[\ t_1\ ]]^{1\to v} = v$$

Step 4.

The fourth step computes the value of the lower S-node in (1.6.a) from the values in (1.25) and (1.26). It may look like we have a problem here in that the terms (1.24) and (1.26) involve distinct assignments. But this can be easily overcome, viz. by making use of the following principle: When an expression has a semantic value for a given assignment **a**, then it also has that value for all assignments which extend **a**. This means that when we combine two terms involving different but compatible assignments, then we can rewrite each term as involving the union of the two assignments and then carry out the function application operation under this new assignment.

In the case at hand one of the assignments is the empty assignment and the union of this and the other assignment is simply that other assignment.

Proceeding in this way we get:

$$(1.27)\quad \left[\!\left[\begin{array}{ccc} & \mathrm{S} & \\ \mathrm{t}_1 & & \mathrm{VP} \\ & & | \\ & & \mathrm{V} \\ & & | \\ & & \textit{suffers} \end{array} \right]\!\right]^{1\to v} = \left[\!\left[\begin{array}{c} \mathrm{VP} \\ | \\ \mathrm{V} \\ | \\ \textit{suffers} \end{array} \right]\!\right]^{1\to v} \left([\![\mathrm{t}_1]\!]^{1\to v} \right) =$$

$$(\lambda u.\text{suffers'}(u))(v) = \text{suffers'}(v)$$

Step 5.

To compute from this the value of the next structure, in which the S node of (1.26) is the complement of the 1-labelled node, we must, as we have already seen, abstract over the relevant value-denoting variable (i.e. v) in the assignment-denoting term that occurs as superscript in (1.26). Recall that the effect of this operation is that the 1-indexed trace in the input structure to this abstraction operation gets bound by the occurrence of its index as separate node in the syntactic structure we are dealing with now. This means that the assignment function that is mentioned in the value term no longer needs to make an assignment to the index 1. We can capitalize on this by evaluating the present structure – in which t_1 is now bound – with respect to a slimmer assignment function, which we obtain by eliminating the assignment to 1 from the assignment referred to in the term that is subjected to lambda abstraction. (That is, we restrict the assignment function **a** of the input term to the function whose domain is $\mathrm{Dom}(\mathbf{a}) \setminus \{1\}$.) Since in the case before us the assignment of the input term, $1 \to v$, is defined only for the index 1, this restriction gets us back down to the empty assignment $\emptyset$.

(1.28) $$\left[\!\left[[\ 1\ [_{S}\ t_1\ [_{VP}\ [_{V}\ \textit{suffers}\ ]]]] \right]\!\right]^{\emptyset} = \lambda v. \left[\!\left[[_{S}\ t_1\ [_{VP}\ [_{V}\ \textit{suffers}\ ]]] \right]\!\right]^{1 \to v} = \lambda v.\text{suffers'}(v)$$

The remainder of the value computation – combining the value of the structure in (1.28) with that of the DP *every man* – goes just as shown in our more informal discussion above. First we evaluate, in Step 6, the phrase *every man*. Since neither its functor, the semantic value of *every*, nor its argument, the value of the noun *man*, depends on any indices, we can evaluate this expression relative to the empty assignment. This result, shown in Step 6, is then combined with the term on the left hand side in (1.28), with the result shown under 'Step 7'.
Step 6.

(1.29) $$\left[\!\left[[_{DP}\ [_{Det}\ \textit{every}\ ]\ [_{NP}\ [_{N}\ \textit{man}\ ]]] \right]\!\right]^{\emptyset} = \left[\!\left[[_{Det}\ \textit{every}\ ] \right]\!\right]^{\emptyset} \left(\left[\!\left[[_{NP}\ [_{N}\ \textit{man}\ ]] \right]\!\right]^{\emptyset} \right) =$$

$$(\lambda P.\lambda Q.(\forall y)(\mathrm{P}(y) \to Q(y)))(\text{man'}) = \lambda Q.(\forall y)(\text{man'}(y) \to Q(y))$$

(Here we make use of our first proposal for the semantics of *every* in (1.9). The semantic value for *man* is the function constant 'man' ', which maps all and only those individuals to the truth value 1 that are men. Compare the entry for *suffers* in (1.20).)

Step 7.

(1.30) 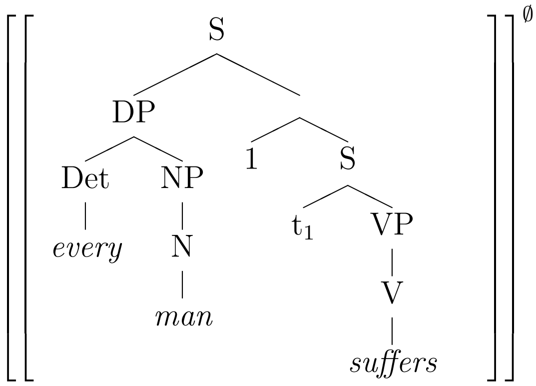 =

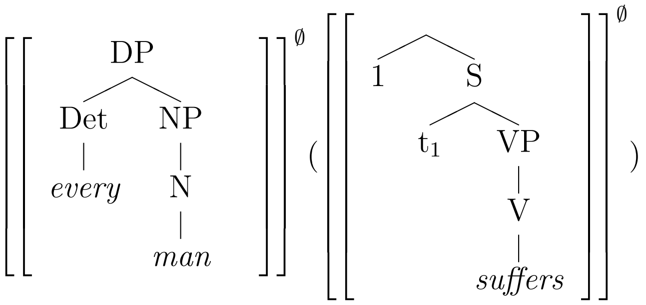 =

$$(\lambda Q.(\forall y)(\text{man'}(y) \to Q(y)))(\lambda v.\text{suffers'}(v)) =$$

$$(\forall y)(\text{man'}(y) \to \text{suffers'}(y))$$

This has been a rather pedantic exercise – lots of technical detail leading to a quite simple and modest result. The reason for adding a value computation for this very simple sentence is to make as clear as possible what the general principles are that govern the syntax-semantics interface in H&K's way of doing things. In particular, it is important to keep clear about what is syntax and what is semantics in these computations. To say it once more: all the terms occurring in the computation above (and in all value computations in this system) are terms belonging to the meta-language in which the theory is formulated. Furthermore, those terms that involve double square brackets, always contain a piece of syntax between the brackets. But what actually appears between the brackets are bits of the meta-language too: they are descriptions, in the meta-language, of the syntactic structures of object language expressions. And finally, as has been stressed already, the double bracket terms always contain sub-terms, appearing in their upper right corner, that denote assignments. Some such terms denote particular assignments which fully specify the individuals that they assign to indices. (An example would be the term '1 → b', where 'b' is an individual constant of the meta-language, which denotes some particular individual of the Domain of individuals that H&K assume as given.) But as we have seen, the assignment-denoting terms needed in the evaluation of traces that result from the raising of quantifying DPs must contain free variables for the values that they assign to indices, so that those variables can be bound by lambda-abstraction at some later stage of the computation.

It is important to appreciate these formal details, and going through the successive steps of one such a computation seems the right way to make clear what they are. But once should be enough. In the use web will be making of the H&K system in what follows we will proceed more speedily, knowing how the details could be sorted out should any one insist.

The semantics for (1.6.e) can be computed from the LF in (1.15) in much the same way as we have just computed that of (1.6.a). For easier reading sentence and logical form are repeated below.

(1.6.e) John offended every linguist.

(1.15)

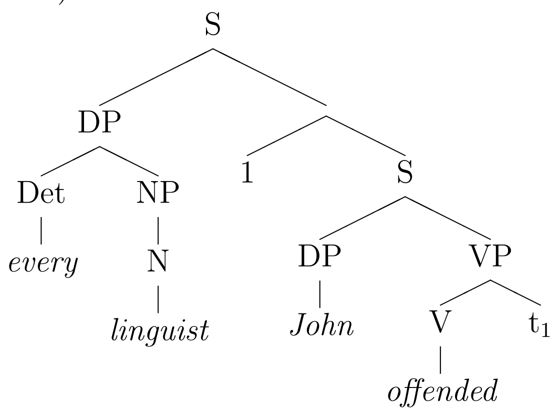

The final step of the value computation for (1.15) is the same as that for (1.6.a) except that the argument of the application is a different property-denoting term, and, uninterestingly, that the functor is built from the property of being a linguist rather than from that of being a man. One of the inputs to the computation of that term is the semantic value of the lexical item *offended*. *offend* is a transitive verb and thus denotes a 2-place relation.

In the functional setting of Montague Grammar such relations take the form of functions from individuals to function from individuals to truth values, with the assumption that the first argument is provided by what is syntactically the direct object and the second by the syntactic subject. Thus, the semantic value of *offended* can be represented as $\lambda u.\lambda w.$offended'(w,u). For the semantics of the structure dominated by the lower S node in (1.15) this works out as follows. Let us adopt for the trace t_1 the same semantic value that we used in the computation for (1.6.a). Then the value for the VP node is as in (1.31)

$$(1.31)\left[\!\!\left[\begin{array}{c} \text{VP} \\ \text{V} \quad \text{t}_1 \\ | \\ \textit{offended} \end{array}\right]\!\!\right]^{1\to v} = \lambda w.\text{offended'}(w,v)$$

To obtain the value of the structure dominated by the lower S node we make use of Montague's principle according to which all DPs are treated as denoting semantic values of type $<< e, t >, t >$. To denote this value we assume that the meta-language has an individual constant j which denotes the person John (the bearer of the name *John* as it is used in this sentence). Then the semantic value of the DP *John* can then be represented as: $\lambda Q.Q(j)$, where Q is a variable of type $< e, t >$. Applying this value for the DP to the semantics of the VP node we get the value in (1.32).

(1.32)$(\lambda Q.Q(j))(\lambda w.$offended'$(w,v)) =$ offended'(j,v)

The next step lambda-abstracts over the variable v, yielding $\lambda v.$offended'(j,v). Combining this with the semantic value of *every linguist* yields the what we want:

(1.33)$(\forall y)$(linguist'$(y) \to$ offended"$(j,y))$

There isn't much that is new in this second sample derivation, as compared with the first one. The main point, in the context of our discussions, is this: because the object phrase *every linguist* has been quantifier-raised in the syntactic structure from which the semantics is computed, we can make use of thee same lexical entry for *every* as for its occurrence in subject position in (1.6.a).

Quantifier Raising also makes it possible to deal with scopally ambiguous sentences like (1.13):

(1.13) Some philosopher offended every linguist.

In order to get the two interpretations for this sentence that it is generally agreed to have we must allow Raising to be indeterminate with regard to which quantifying DP ends up with scope over which. There are various ways in which we can define the principles for quantifier raising so that they produce this effect for sentences with this particular kind of scope ambiguity. For single clause sentences like (1.13) one option is to leave it open in which order the different DPs are to be raised, but to insist that when a DP is raised it s always raised to the highest possible position – that is: adjoined to what is that point the highest S-node. In this way the DP that is raised last gets the widest scope. An alternative is to insist on a fixed raising order but to allow raising to take the form of adjunction to any S-nodes identical with or dominating the original S-node of the surface structure tree. And a third possibility would be to allow for both kinds of freedom – order of raising and target position – at once. We take no stance on this matter but will, for definiteness' sake, assume that scope ambiguities of this kind are due to indeterminacy in the order of raising.[9]

On this assumption the ∀∃-reading of (1.13) is the result of first raising the subject DP and then the object DP, ass shown in (1.34.b,c)

(1.34)a.

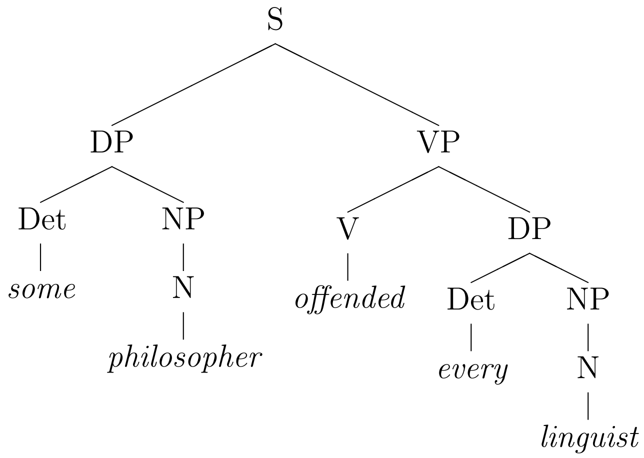

[9]For multi-clause sentences matters are more complicated, since some clauses boundaries are barriers to QR, in the sense that a quantifier DP occurring in the clause cannot be raised out of it – 'cannot' in the sense that otherwise truth conditions result that the sentence does not have on any of its possible readings. We do not pursue this complication here.

b. 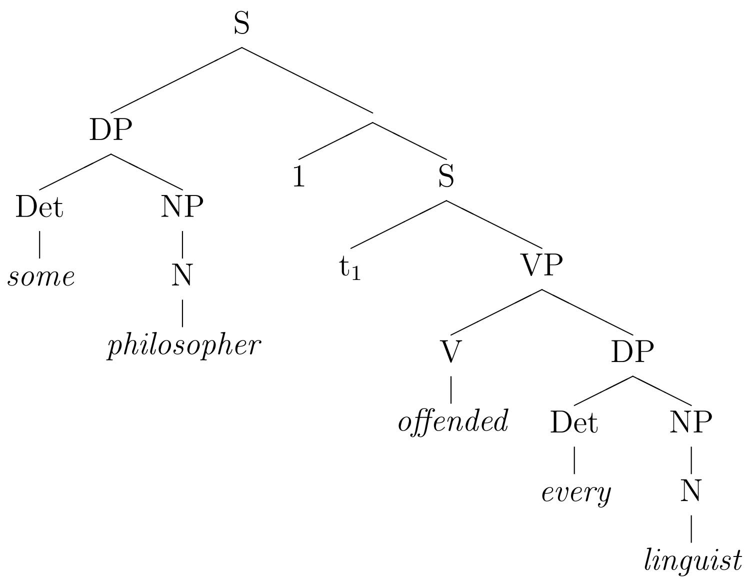

c. 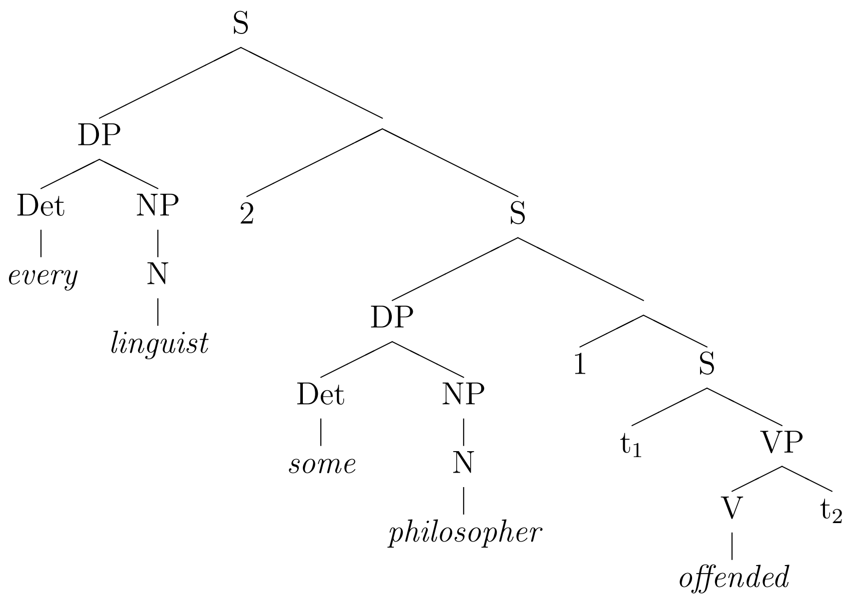

After the cases of value calculation that we have gone through it should not be difficult to carry out the calculation of the semantic value of (1.13) on the basis of (1.34.c) and show that this yields its ∀∃-reading. Further, first raising the direct object DP and then the subject DP and calculating the semantic value on the basis of this leads to the ∃∀-reading.

<u>Exercise</u>: Carry out these calculations.

<u>Remark</u>: Should Quantifier Raising be restricted to quantifying DPs or be

applied (obligatorily) to all DPs? It is hard to come up with a principled answer to this, not at least when we want to argue just on the basis of semantic considerations. For as H&K note, whether a non-quantifying DP is raised or left *in situ* is not going to make any difference to the truth conditions.[10] In the context of DRT, to which we turn below, there is a methodological consideration that favors not raising non-quantificational DPs out of their base positions. This is connected with the treatment in DRT (and other dynamic theories) of indefinite DPs, such as *a linguist. some philosopher* and so on. These are not treated as existentially quantifying phrases – the way they are in static semantics and have been in formal logic since Frege – but as 'indefinite terms'. Quite often indefinite DPs are used 'specifically', in the sense that the speaker uses them to talk about some particular individual she has in mind, but chooses an indefinite for this purpose in order to avoid suggesting to the addressee that she assumes that he will be able to identify who she is talking about, or that this individual is familiar to him. The logical effect of this is that the proposition expressed by the speaker involves maximal scope of the existential quantifying force of the indefinite; and that will be so irrespective of how deeply embedded the indefinite DP occurs within the sentence. For 'genuine quantifier phrases', such as those that begin with the determiner *every*, this is not so: they are subject to so-called 'island constraints', which put boundaries on how wide the scope of a quantifying DP can be given its syntactic position. (For instance, a genuinely quantifying DP occurring in a relative clause cannot raise to a position outside of that clause.) In the light of these considerations it may be preferable to treat the wide scope effects of specific indefinites not as the result of quantifier raising, but as entailed by the quasi-referential role that such indefinites play. Note well, however, that if we were to assume that indefinites can never be quantifier-raised, then the account just given for the scope ambiguity of (1.13) would be in jeopardy. To avoid this the simplest assumption is that indefinites can also function as genuine quantifier phrases and as such subject to quantifier raising, albeit with the restriction that their landing sites (i.e. the syntactic positions they end up in as a result of raising) can't be adjunction's to a node higher than the S node dominating their base positions (the positions from which they are moved). But the matter is complicated, for one thing because the semantics of more complex sentences, which ought to be testing ground for these assumptions about the form and

[10] Exercise Show this by computing the semantic value of (1.3.e) from (a) a syntactic structure in which *John* has been raised to a position above *every linguist* and (b) a syntactic structure in which *John* has been raised to a position below *every linguist*. In both cases you should get to the same truth conditions as we obtained on the basis of (1.15).

limits of quantifier raising, is difficult to judge, with diverging judgments from different speakers. I do not want to go into this more deeply, but also do not want to create the impression that the last word about scope ambiguity and quantifier raising has been spoken.

1.4 Inversely Linked Quantifiers

Quantifier Raising has additional benefits. Prominent among these are the cases of 'inversely linked' quantifiers.[11]
Consider the following sentences:

(1.35)a. Some apple in every basket is rotten.

b. No student from a foreign country was admitted.

(1.35.a) is a good example of inverse linking. Its natural reading is that for every basket there is some apple in it that is rotten. This interpretation is a case of 'inverse linking' because semantically the quantifier 'every basket' has scope over the quantifier 'some apple', while in the surface structure of the sentence the phrase *every basket* is, as part of the prepositional phrase *in every basket*, embedded within the larger phrase *some apple in every basket*, by adjunction to the noun *apple* (see the syntactic structure (1.36) below); and that would, if one were to compute the semantic value bottom up, lead to an interpretation in which 'every basket' has scope below 'some apple'.

As (1.35.b) shows, sentences of this general form are ambiguous between an interpretation in which its quantifying phrases are inversely linked and one in which they are linked as the surface structure would seem to predict. The arguably more natural reading of (1.35.b) is that which can be paraphrased as 'No foreign student was admitted.'. But the other, inversely linked reading, according to which there was one country no student from which was admitted, seems possible too. For (1.35.a) the non-inversely linked reading is hard to get because it is so implausible: How could one apple be in all baskets at the same time?

The task, then, is to explain how such sentences admit both readings. In particular, we must show that they have the inversely linked reading, and for this reading a mechanism like Quantifier Raising seems indispensable.

[11]It was the inverse linking phenomenon that originally motivated May to introduce Quantifier Raising. See footnote 6

(1.36.a) gives the structure of (1.35.a) before QR. (1.36.b,c) show how two successive raising operations, first of the entire subject DP out of its subject position and then of the embedded DP *every basket* out of the raised subject DP, can give us the structure from which the H&K interpretation procedure will give us the semantics of the inversely linked reading. (Note well that as we have reconstructed the procedure, the second raising moves the DP *every basket* out of the subject DP of which it is part to a position in which the raised DP is also adjoined to S.)

(1.36)a.

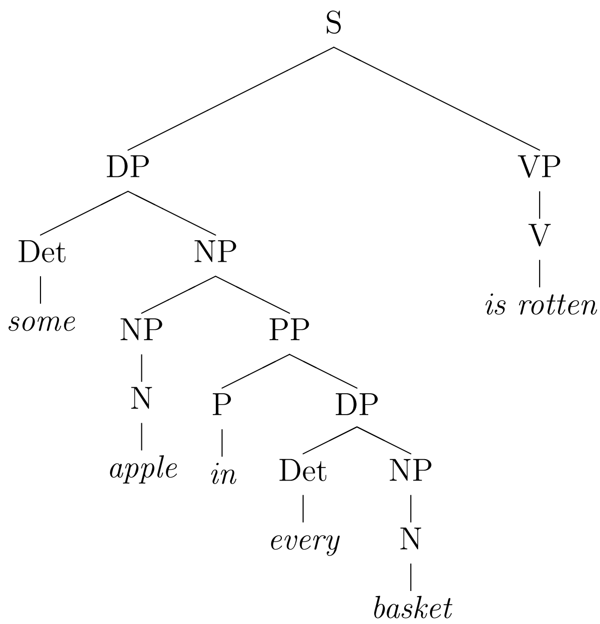

b. 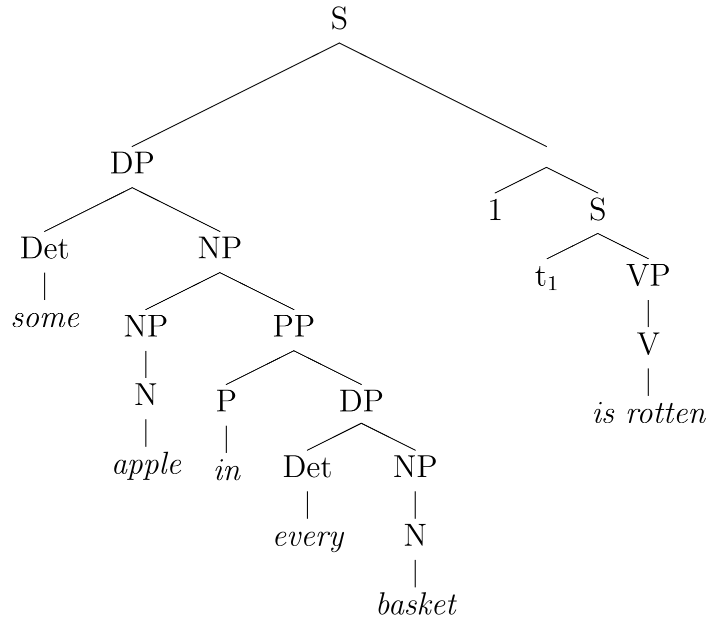
S
DP
Det
some
NP
NP
N
apple
PP
P
in
DP
Det
every
NP
N
basket
1
S
t1
VP
V
is rotten

c. 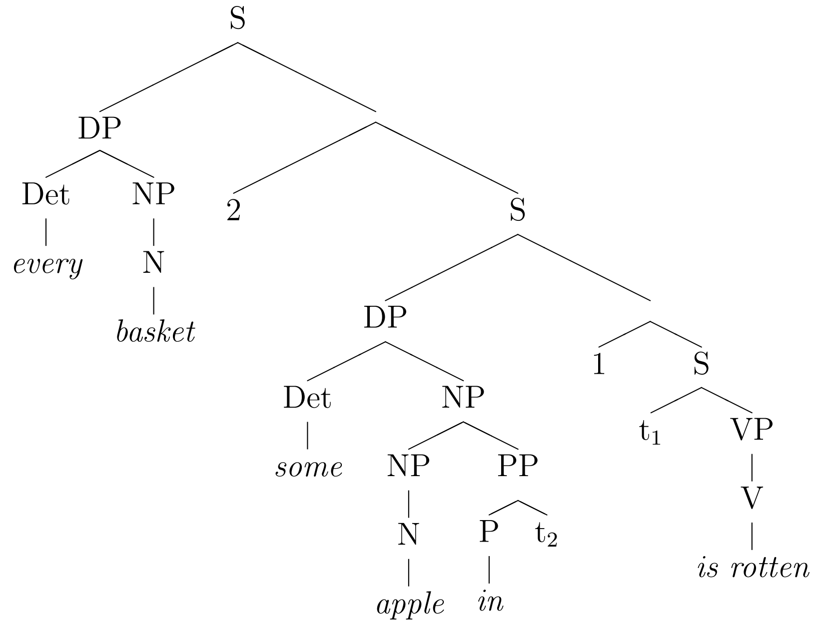
S
DP
Det
every
NP
N
basket
2
S
DP
Det
some
NP
NP
N
apple
PP
P
in
t2
1
S
t1
VP
V
is rotten

The syntactic structure in (1.36.c) leads to the reading of (1.35.a) that we were looking for, that according to which the sentence says that each basket contained one rotten apple.

Exercise: Show this!

Various other raising possibilities are prohibited, either because they violate certain constraints on movement (though that is an issue we will no further pursue here) or because the structure cannot be coherently interpreted using the H&K algorithm or perhaps for both reasons at once. Some examples of such illicit and/or abortive structures are shown in (1.37). The problem with (1.37.a), in which raising the DP *every basket* leads to its adjunction to the DP from within which it is moved rather than to adjunction to S, is that no proper interpretation of the complex DP is obtained. (Exercise: Show this!). (1.37.b) does yield the desired interpretation, but the structure may be ruled out on account of the fact that it contains a quantifying DP that hasn't undergone raising. (This is on the assumption that the construction of a proper syntactic structure as input to the computation of the semantics – a proper 'LF', in syntactic parlance – requires raising of *all* quantifying DPs.) If the unraised DP of (1.37.b) is raised, then the result will be either (1.36.c) or it will be (1.37.c). Whether (1.36.c) can be legitimately obtained from (1.37.b) depends on the exact constraints on movement. But in any case we have already assumed that this LF can be constructed and that it gives one of the readings of (1.35.a). (1.37.c), on the other hand, does not lead to a coherent interpretation because of the 'unbound' trace t_2.

(1.37)a.

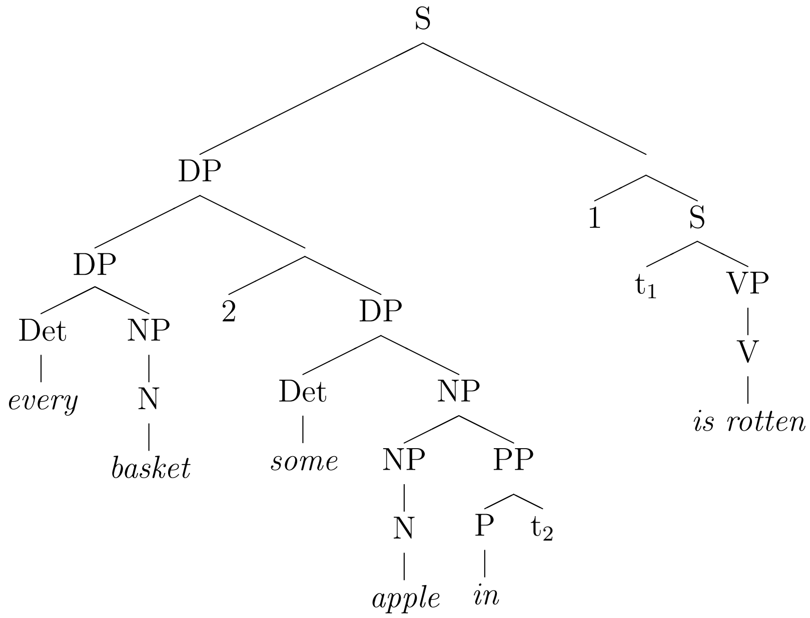
S
DP
1
S
DP
2
DP
t$_1$
VP
Det
NP
Det
NP
V
every
N
some
NP
PP
is rotten
basket
N
P
t$_2$
apple
in

b.

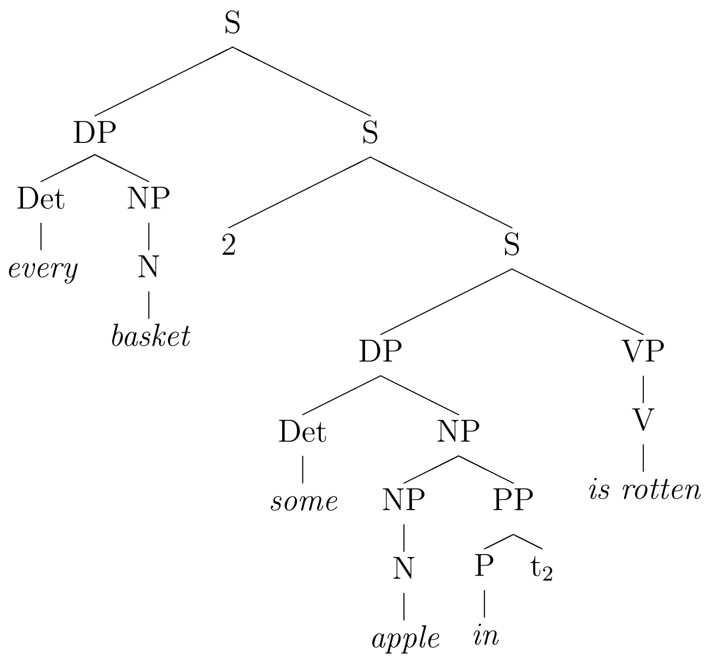
S
DP
S
Det
NP
2
S
every
N
DP
VP
basket
Det
NP
V
some
NP
PP
is rotten
N
P
t$_2$
apple
in

c.

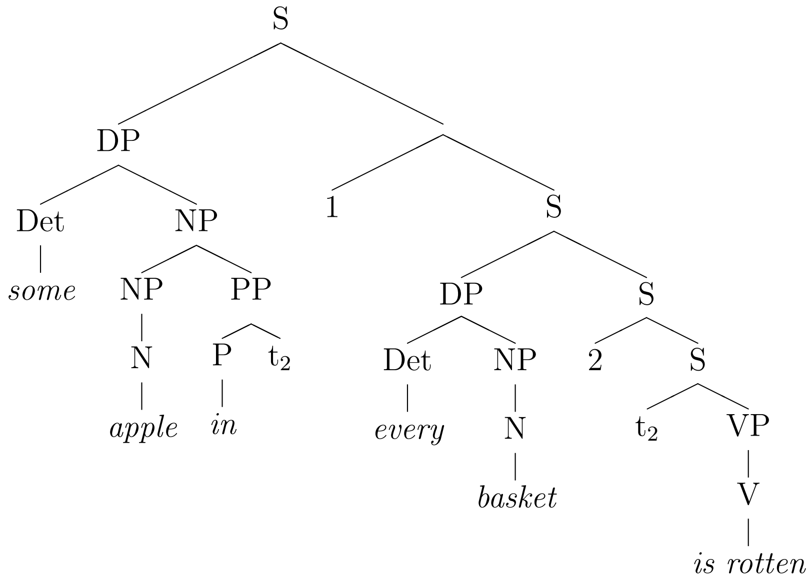

In order to get the reading in which the quantifier *every basket* gets narrow scope with respect to *some apple* we have to proceed differently: as shown in (1.36.d,e), first raising *every basket* out of the subject DP and then raising what remains of the subject DP to a higher position won't work since this leaves the trace left behind by *every basket* unbound. Rather, we must, as already noted by May, assume that *every basket* is raised to a position within the DP: there must be some kind of barrier within the subject DP that prevents *every basket* from moving beyond it, thereby keeping it within the subject DP and thus within the scope of the quantifier *some apple*. May assumed that the subject-internal raising is the effect of analyzing the PP *in every basket* as the Copula Complement of a reduced Relative Clause in which both the copula and the complementizer (normally a relative pronoun) are silent. RCs – whether full or reduced – are assumed to be barriers to QR: When a constituent of a RC is quantifier-raised, it must be adjoined to the S node below RC. The 'surface structure' of (1.35.a) in which the PP is analyzed as part of a reduced RC is as in (1.36.f). (1.36.g) gives the result of raising of *every basket* within the subject DP and (1.36.h) the result of subsequently raising the subject DP. (These last two raising operations can be executed in either order. There is no difference to the final outcome, either syntactically or semantically.)

(1.38)a.

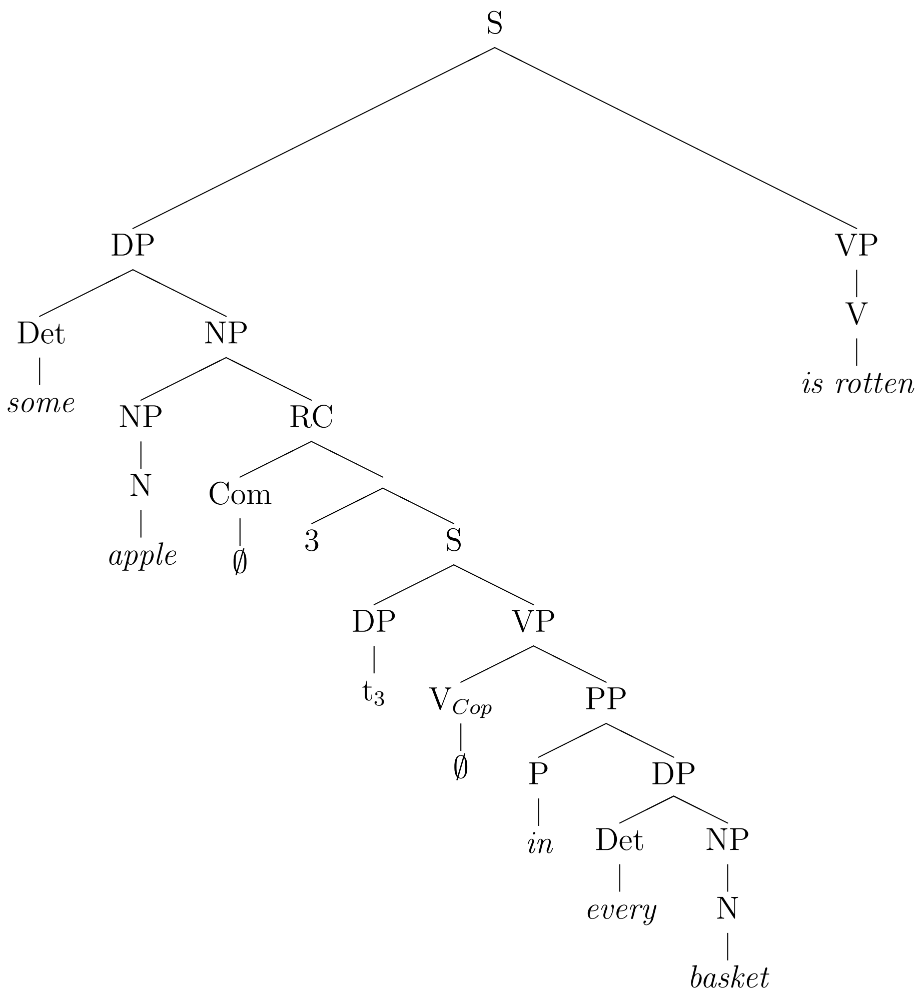
S
DP
VP
Det
NP
V
some
is rotten
NP
RC
N
Com
apple
∅
3
S
DP
VP
t3
VCop
PP
∅
P
DP
in
Det
NP
every
N
basket

b.

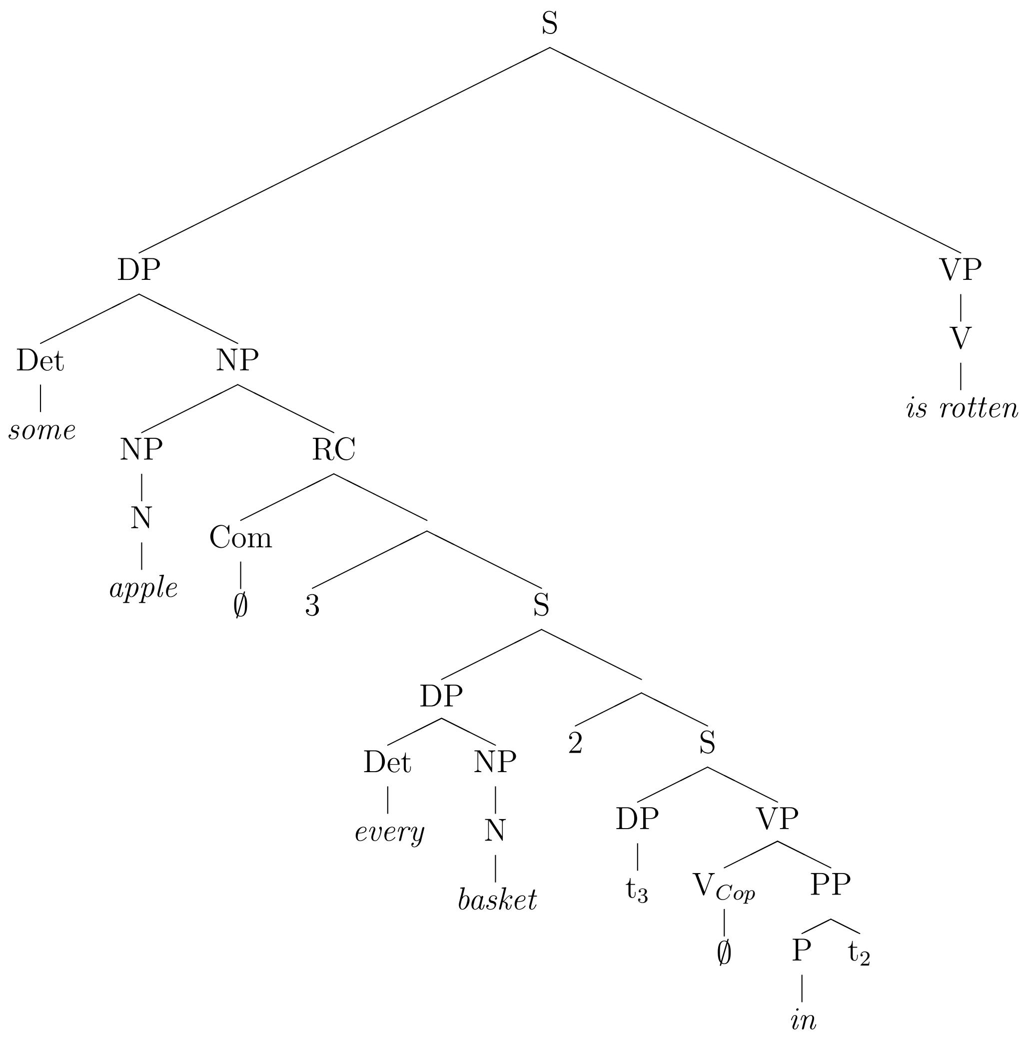
S
DP
VP
Det
NP
V
some
is rotten
NP
RC
N
Com
apple
∅
3
S
DP
2
S
Det
NP
every
N
basket
DP
VP
t3
VCop
PP
∅
P
t2
in

(1.39)

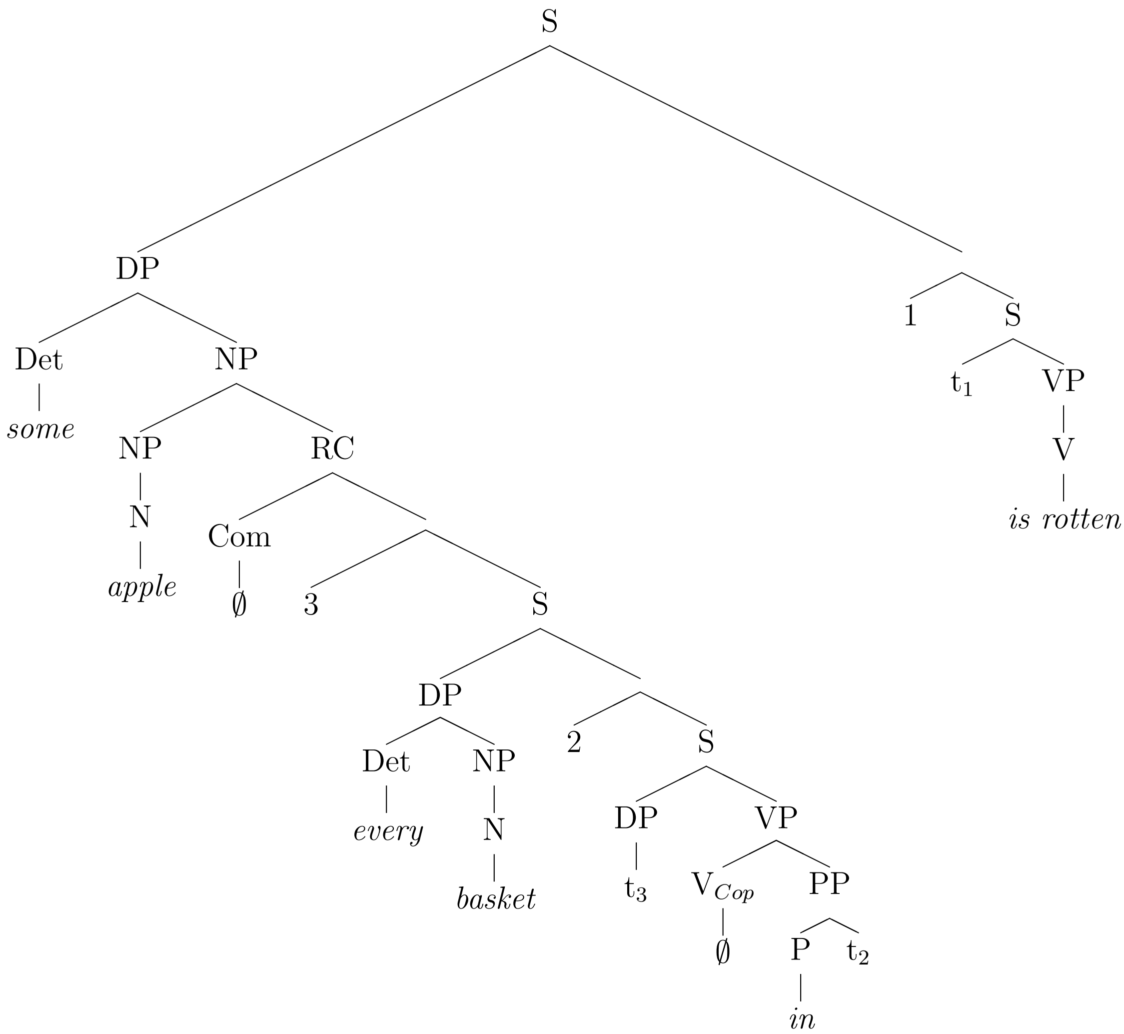

1.5 A Problem about Pronouns

There has been an important methodological controversy during the past three and a half decades over the treatment of third person pronouns. We can distinguish at least four different uses of the 3rd person singular pronouns of English:

(i) Deictic uses: the pronoun refers to an individual that is salient in the context, or which the speaker makes salient, for instance by pointing. Example:

(1.40)He is the one

(said by the victim of a mugging during a police line-up, perhaps pointing at the one she wants to refer to)

(ii) Discourse-anaphoric uses of pronouns that refer to particular (real world) individuals: There is a linguistic antecedent for the pronoun; but the pronoun ends up referring to a particular individual in the real world and it would be possible in principle to take the role of the linguistic antecedent to be that of a 1linguistic way of pointing' to the pronoun's referent.

(1.41)a. John works as a salesman. He is miserable.

b. If John works as a salesman, then he is miserable.

(iii) Uses involving discourse anaphoric to 'indefinite individuals' (so-called 'donkey pronouns')

(1.42)a. John has bought a goat. It is not allowed inside the house.

b. If John has bought a goat, then apparently it is not allowed inside the house.

(iv) 'True bound variable uses': the pronoun is syntactically dominated by its anaphoric antecedent.

(1.43)a. Every man who fancies a woman hopes she fancies him too.

b. Every man thinks he is cleverer than every other man.

H&K treat pronouns as ambiguous between bound pronouns – these are pronouns that get an index in the syntactic structure – and referential pronouns. The latter category includes both deictic uses of pronouns and pronouns that are anaphoric to antecedents which may be either inside our outside the sentence containing the pronoun, but where the antecedent is in a position from where it is unable to bind the pronoun. So the cut-off here is between use-type (iv) and the three other use-types (i)-(iii).

The occurrences of the pronoun *it* in the 2-sentence discourse (1.42.a) and the conditional sentence (1.42.b) are known as *donkey pronouns*. The phenomenon that these occurrences illustrate got its name from sentences like those in (1.44), mentioned by Geach in his (Geach 1962). Donkey pronouns have played a pivotal role in methodological discussions that started around 1980 and that led to a division within the formal semantic community. Discourse Representation Theory, the way of doing semantics that we turn to now, represents one approach to the problems donkey pronouns present. H&K can be seen as a major representative of the opposite perspective. [12]

(1.44)a. If Pedro owns a donkey he beats it.

b. If a farmer owns a donkey he beats it.

c. Every farmer who owns a donkey beats it.

[12]The approach to the donkey pronoun problems that is endorsed by Heim and Kratzer H&K treats donkey pronouns as so-called 'E-type pronouns'. The term 'E-type pronoun' goes back to G. Evans ((Evans 1977), (Evans 1980)). It is used to refer to a family of methods for dealing with the semantics of sentences like those in (1.44), which all come to this:

The pronoun is interpreted as a certain definite description.

There are three respects in regard of which different variants of this approach vary:

(i) How is this description obtained? Is it to be constructed from other linguistic material in the sentence or discourse and/or from the extra-linguistic context, and exactly what principles are involved in the reconstruction?

(ii) What is the exact relation – the relation alluded to by the phrase 'is interpreted as' – between the description and the pronoun? One subfamily, that of the so-called 'D-type treatments', treat the pronoun as the remnant of the description after the description has been subjected to a process of elision. (That is, the pronoun is regarded as elliptic for the description.) For details see ((Elbourne 2005)).
(iii) How do we deal with the contribution that is made by the definite description to the semantics of the sentence or discourse in which the pronoun occurs? (This is important. We should not take it for granted that the contribution made by the description is not problematic in the same way as the pronoun which is treated as its substitute.) In order to be able to interpret the description according to which it refer to the unique satisfier of its descriptive content, more recent versions of the E-type approach adopt a form of *Situation Semantics* to deal with this aspect of the problem: Situations can be assumed to be small enough to guarantee uniqueness. (Again, see ((Elbourne 2005)); the Situation Semantics used by Elbourne himself goes back to (Kratzer 1989), and was first used in the context of 'donkey pronouns' in (Berman 1991).)

Chapter 2

Another approach: Discourse Representation Theory

2.1 Discourse Semantics with Simple DRSs

A rather different way of dealing with donkey pronouns and related problems is proposed by *Discourse Representation Theory* (DRT).

We first have a look at the way in which anaphoric pronouns are dealt with by DRT in its original form. (See (Kamp & Reyle 1993).)

Preliminary remark.

DRT started out as an attempt to capture the systematic semantic connections between successive sentences in a discourse, including those which are established when a pronoun in one sentence of a discourse or text is interpreted as anaphoric to an antecedent that occurs in an earlier sentence. Furthermore, the aim was to show that these connections and the semantic effects they produce are essentially the same as those that are found in single sentences (i.e. when pronoun and anaphoric antecedent occur in the same sentence): inter-sentential and the intra-sentential cases of anaphora should be explained as instances of the same general mechanism.

Striking examples of such connections are those established by anaphoric pronouns whose antecedents occur in earlier sentences. So the behavior of such pronouns was one of DRT's first targets.

DRT's general strategy for text- and discourse-interpretation is based on the assumption – which is very largely true of the way in which multi-sentence discourse and texts are understood – that each initial part of the text or discourse, consisting of its first n sentences can be interpreted irrespective of what comes after it, and that the interpretation that is obtained for this initial segment can then be used as 'discourse contact' in the interpretation of the next sentence; and further that this interpretation of the next sentence will extend this discourse context, making it into the interpretation of the first n + 1 sentences, which can then in its turn serve as discourse context for the interpretation of sentence n+ 2. Thus each such interpretation does two things: it captures the content of the sentences processed and serves as context for the sentence that is to be interpreted next.

This architectural feature is common to all versions of DRT. But the older versions differ from the more recent ones in the ways they articulate the interpretation process for the individual sentences. The original method, of which we are going to give an illustration first, processes the syntactic trees for the sentences to which it is applied from the top down, decomposing the tree into parts that come to look more like formulas of predicate logic as one works his way down. By the time one is done a representation has been obtained that can be seen as a variant of first order predicate logic. (But'variant' is important here. It is the precise way in which the information in these completed representations – the completed *Discourse Representation Structures* or 'DRS's' – is organized that makes DRSs usable as discourse contexts.) This 'top-down' method contrasts with the more recent 'bottom-up' methods that are mostly used today, and which we will use during most of this course.

We now turn to some examples. We start with a few applications of the topdown method. Our first exercise of this kind will be the mini-discourse in (2.1).

(2.1) Pedro owns a donkey. He beats it.

As indicated, we begin by constructing a DRS for the first sentence of (2.1). Our input is the syntactic tree in (2.2). More exactly, it is the first step in the construction of a DRS, which is indicated by the box around the tree. The point of the box will become clear presently.

(2.2) 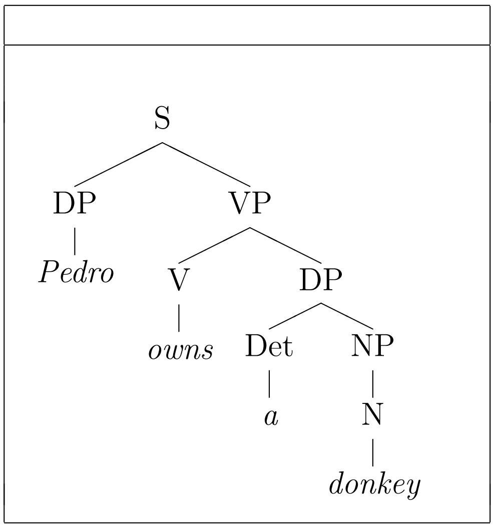

The first construction step to be performed on (2.2) decomposes the S-node 'semantically' into the contribution made by the subject DP and that made by the VP. The syntactic configuration consisting of the S-node and its two daughters is one of many in which one of the daughters plays the part of argument to the predicate that is contributed by its sister node. In all such cases the argument constituent contributes a *discourse referent* ('dref') to the representation of its sister – something that is implemented by inserting the dref into the position of the sister representation that was occupied by the argument phrase. The dref represents the entity or entities that the argument phrase refers to or quantifiers over. What this referent is, or what kind of quantification is involved, depends on the form of the argument phrase. When, as in the case before us, the argument phrase is a proper name, then, intuitively speaking, the dref represents the referent of the name (more exactly: the individual or entity to which the name is being used as a name of). For now we encode this information as a condition of the form 'Name(dref)'.[1]

In addition to putting the chosen dref into the position occupied by the argument DP we also place it in a separate component of the representation

[1]This is really a kind of ad hoc solution, though it is found in many presentations of DRT, as this was the way that proper names were dealt with in the original presentation of DRT and also in ((Kamp & Reyle 1993)) which served for many years as principal reference source. Later on we will treat proper names and other definite noun phrases as presuppositional: each definite DP comes with a presupposition to the effect that the interpreter should be able to 'identify' the referent. However, precisely what 'identifying' comes to is often difficult to spell out in detail, and varies from one type of definite DP to another. We will return to this matter eventually, in the part devoted to presupposition.

that is being constructed. This is the so-called *Universe* of the DRS. (A DRS always consists of two components, its Universe and its *Condition Set*. The latter contains all the (reducible and irreducible) *DRS-conditions*.) Note that our starting structure in (2.2) is already in DRS format, with an empty Universe (above the horizontal dividing line) and (below the line) a Condition Set whose only element is the syntactic tree that (2.2) displays. The construction steps of which we have just described the first transform this DRS into the final DRS, in which all Conditions have been fully reduced.
To apply the rule just described to the present case we choose some dref, x, say, and decompose the tree in (2.2) in the way the rule prescribes. The result is as in (2.3).

(2.3)

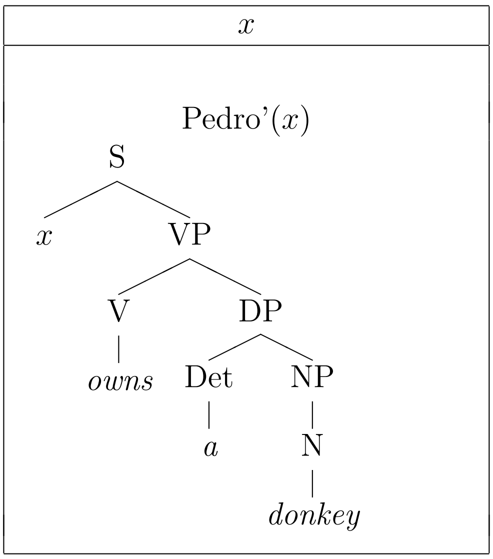

The next step (and only other step in this simple example) deals with the configuration of the VP node and its two daughters V and DP. Here we have once again a combination of an argument (the DP) and a predicate (the V) of which it is an argument; and the general treatment – choose a new dref to represent what the DP refers to or quantifies over, insert that dref into the argument position occupied by the argument phrase in the syntactic structure, separately represent the descriptive information of the argument phrase, and place the dref in a DRS Universe – is as it was for the first step.

However, the DP we are dealing with now is of a different type than the subject DP we dealt with in our previous construction step: it is an *indefinite description*, not a proper name (or, for that matter, a definite DP of any kind). The treatment of indefinite DPs in DRT (and in the independently developed *File Change Semantics* of Irene Heim (see (Heim 1982,1988))) was a novelty at the time when these accounts were proposed and it still sets these approaches apart from other ways of doing formal semantics, and in particular from the E-type and D-type approaches mentioned above. Indefinites are not treated as existential quantifiers (the assumption implicit in the way in which logicians had symbolized sentences with indefinites since Frege and Russell) but in a way that more closely reflects their traditional name of 'indefinite descriptions'. What this difference exactly comes to cannot be fully explained at this stage. The point will become clear when we turn to 'genuinely quantifying' DPs later on.

All that can be done at this moment is to point out in what ways the treatment of indefinites resembles that of definites. Common between the two is that – besides the operations mentioned in the one but last paragraph, which are the same for all argument DPs – in both cases the new dref is introduced into a DRS universe either at or above the level at which the construction rule is being applied. Exactly what that means cannot be appreciated at this point, at which all that we have so far seen is a single simple DRS with just one DRS-Universe, so that there is only one such universe where the dref can be put. (The matter will become clearer when we turn to quantifying DPs whose treatment is more complicated and leads to a differentiation of DRS levels.) Since, as said, there is in the case before us just one Universe that the new dref can be put into, that is where we put it. (So the dref for the indefinite ends up in the same Universe as the dref x we introduced for the proper name *Pedro*. We will soon see that that isn't always so.)

The second difference between an indefinite DP like *a donkey* and a proper name concerns the constraints that the DP places on the dref that is chosen to represent it: A proper name fixes the individual that is represented by its dref, where the content of an indefinite DP only imposes a restriction on what that individual can be. For instance, the indefinite *a donkey* imposes the constraint that what the dref represents must be a donkey - a constraint we implement in the form of the condition 'donkey'(y)' (assuming that y is the dref we choose in this case).[2]

[2]The formal similarity between the conditions 'donkey'(y)' and 'Pedro'(x)' is potentially misleading, as the conditions have a quite different meaning. 'donkey'(y)' is a case of normal predication, expressing that what y stands for must satisfy the predicate 'donkey''. The condition 'Pedro'(x) is really a stopgap for a process which identifies the individual that x stands for as the individual that the name *Pedro* is used to refer to on the given occasion.

The result of carrying out the second step, with y chosen as the dref chosen to represent *a donkey*, is as in (2.4),

(2.4)

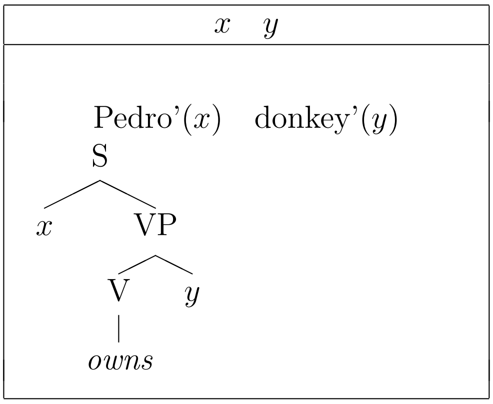

This is the complete DRS for the first sentence of (2.1). We can simplify this DRS by inserting the lexical semantic representation of the verb *owns*, which we take to be the 2-place predicate 'owns", and then throwing away the remaining syntactic skeleton, which at this point does nothing more than shore up the predication involving 'owns" and the arguments x and y. This notational simplification gets us (2.5).

(2.5)

x y
Pedro'(x) donkey'(y)
owns'(x,y)

To process the second sentence of (2.1) we start by inserting its syntactic structure as new reducible condition into the DRS (2.5) that we have just obtained for the first sentence. This gives (2.6).

(2.6)

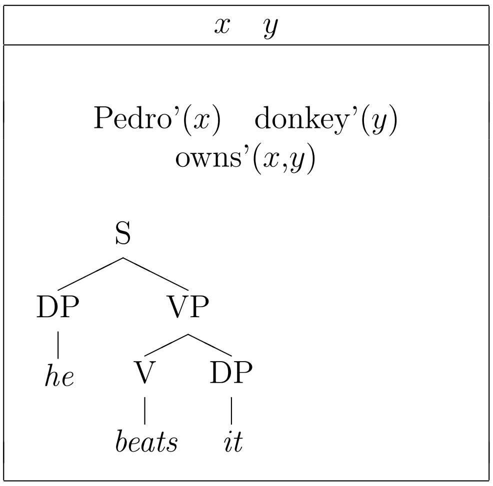

Once again the first rule to be applied deals with the combination S-DP-VP. This time the subject DP is a pronoun. We assume it as given that the pronoun is used anaphorically (rather than deictically) – an assumption that is almost inescapable when one deals with texts rather than spoken discourse – and in DRT this means that its antecedent must be found among the discourse referents in *accessible* DRS Universes. In the present case we have only one Universe and it contains two drefs, viz. x and y. Once again we assume that processing the DP involves the introduction of a new dref and we choose u for this purpose. Interpreting the pronoun as anaphoric to x or y can now be expressed via the identities '$u = x$' and '$u = y$', respectively, and that is how we will represent these respective options.

Which of the two options is it going to be? That is a question which DRT in its original form does not address, although the theory was consciously set up in such a way that it could be extended with accounts of 'anaphora resolution' which deal with this question explicitly. Intuitively, though, it is clear that the antecedent of *he* could only be x: Pedro is presumably a person; only persons can own other things, including donkeys, and the male pronoun *he* is used in English mostly to refer to persons. (There are some few exceptions to this, but they are not relevant here.) Individuals that are introduced as 'donkeys', on the other hand, are not the kind that can be referred back to by means of *he*. These considerations uniquely select x as anaphoric antecedent for u.

In all other respects the processing of the DP-VP combination is like what we saw when dealing with DP-VP reduction of the first sentence. So the result we get is that in (2.7)

(2.7)

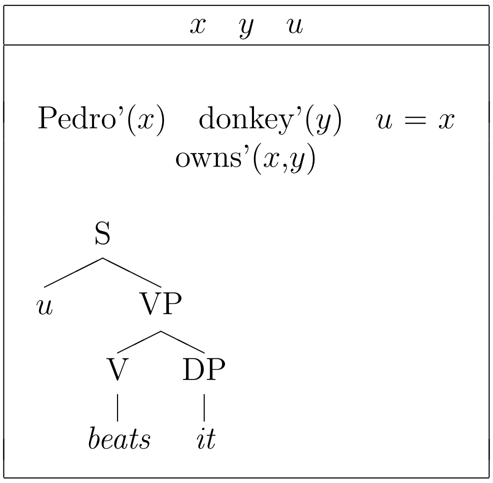

The Pronoun *it* that is direct object to *beats* in (2.7) is again treated as anaphoric pronoun and again the question comes up which of the discourse referents in the Universe of (2.5) should be selected as its antecedent. Considerations similar to those just brought to bear on the case of the subject *he* point to y. This time x is ruled out, and u for the same reasons as x, in virtue of the identity ‘ $u = x$’, which transfers those reasons from x to u. (An additional consideration is that once we have chosen x as the antecedent for the subject we cannot use it again as antecedent for the direct object, since that would have led to the interpretation that ‘x *beats* x’ and in English such a semantics can only be expressed with the help of a reflexive, as in ‘Pedro beats himself’, but not with a pronoun.) All other operations are as we now know them. Choosing v as the new dref, we get the result in (2.8.a), or, after inserting the semantic representation ‘beats’ for *beats* and throwing away the remaining syntax, in (2.8.b). (2.8.b) is the DRS for the two sentences of (2.1) together, in other words for the two sentence text as a whole.

(2.8) a.

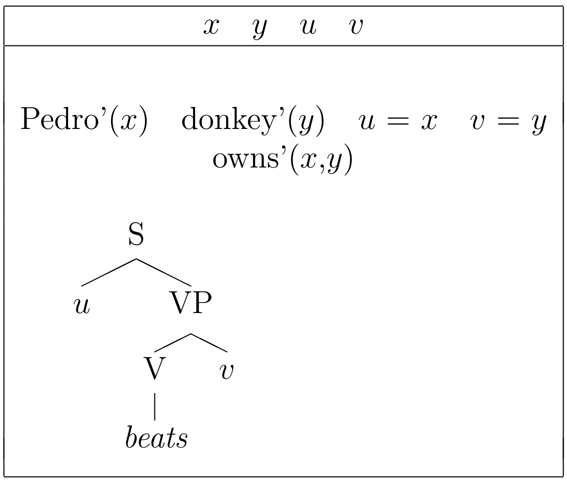

b.

$x \quad y \quad u \quad v$
Pedro'(x) donkey'(y) $u = x$ $v = y$ owns'(x,y) beats'(u,v)

The DRSs (2.5) and (2.8.b) are *simple* DRSs – DRSs whose conditions are all atomic predications. Simple DRSs can be thought of as partial models of the world, (2.5) as a model of the world as described by the first sentence of (2.1) and (2.8.b) as model of the world as described by the two sentences of (2.1) combined.

A simple DRS is a correct model of the world in case the world contains individuals corresponding to the drefs in the DRS's Universe such that these individuals satisfy the DRS's conditions More formally:

Def. 1 A simple DRS is *true* (of the world) iff there exists a function f that maps the discourse referents of its universe to entities (of the world) which satisfy the conditions that the DRS specifies for the drefs to which they correspond under f.

Formally, we may think of DRSs as the 'formulas' of some logical formalism, like predicate logic or the λ-calculus. One task of DRT has been to develop such formalisms – so-called 'DRS-languages' – that are suited for the representation of natural language (and some related purposes, such as the representation of thought, or of 'knowledge' or information, as in 'knowledge representation'). And the specification of a DRS-language involves, like the full specification of predicate logic or the λ-calculus, not only the syntax of the language – which specifies which combinations of symbols are 'well-formed formulas' (i.e. well-formed DRSs) – but also a precisely defined semantics.

Here we are at a kind of crossroads. There are two ways in which the semantics of a formalism can be made formally explicit. Formally there isn't much of a difference between them, but conceptually there arguably is. One way is to assume that the world in which we are and about which we speak is a certain way. It consists of objects, which have certain properties and stand to each other in certain relations and some of those properties and relations are denoted by the words of the language for which we seek a formal semantics (its 'predicate words'). The language may also have names for some of the objects and it may have terms that pick out objects in some other way, say according to some rule or set of rules. The sentences of the language make definite claims about the world, claims that are either true or false. And they do so – that is the central intuition – in virtue of on the one hand (i) their syntactic form and on the other in virtue of what real world objects satisfy the properties and relations denoted by the predicate words they contain. What we want from our semantic theory is an account which tells us exactly how the truth values of the sentences of the language are determined by these two factors.

How such an account can be given is discussed in detail in H&K. The H&K account is paradigmatic of all accounts of truth or truth conditions in formal semantics: it states (i) what makes atomic predications true – this is more or less trivial, since such predications are direct records of certain objects satisfying certain properties and of certain objects standing in certain relations, and such records are either true, when they fit the facts of the world, or false, when they don't[3] – and (ii) how the semantics of any syntactically complex expression is determined by the semantics of its immediate constituents. It is the second of these two tasks that is the real challenge. For one thing, the syntactic constituents of sentences are usually not sentences, and they need to be assigned other semantic values than truth values to make the theory work. That is why we need an assignment of semantic types to the different syntactic categories, define what type of semantic value goes with each of these semantic types and then formulate the connections between the semantic value of an expression and the values of its constituents as relations between the kinds of semantic values that the expression and its constituents should have in virtue of their syntactic categories.

What matters at this point, however, are not the details of such a compositional definition of semantic values, but what is at the bottom: the way the world is, in which the properties and relations for which the language has predicate words have their actual extensions. What the definition shows is how the truth value of each sentence depends on such facts, given what syntactic structure it has. We need not actually know wheat those exten-

[3]The formulation here is a little sloppy. The atomic predications I am speaking of typically consist of a predicate word combined with argument terms. These argument terms can be either 'variables' or constants and the 'variables' can be realized in various ways. Recall that in H&K there are no variables in the syntax of the LFs for which the semantics is defined, but only traces and their indices.

sions are to see that the definition accomplishes this. But of course for most sentences we do need such information if we want to determine whether they really are true or false.

The second option is to abstract away from the actual extensions of the properties and relations that the predicate words of our language are supposed to denote. That is, we do not refer in our semantic value definition to one particular world, in which the denoted properties and relations have their particular extensions, but to a whole range, or class, of such worlds – or, more accurately, to a class of abstract structures that can play the part of worlds in a definition of truth and other semantic values for expressions of our language in that each of them provides extensions for its predicate words. The semantic value definition must now specify what it is for a sentence of the language to be true in any model of the specified model class, so semantic values now identify relations between expressions of the language and models. In particular, truth becomes a relation between sentences and models. In order to get from such a definition of 'truth in a model' to a notion of truth simpliciter, one of the models in the class must be identified as specifying the actual universe of objects and the actual extensions of the properties and relations denoted by the predicate words of the language. That is presumably a practical impossibility. But it may be possible to establish constraints on what kind of model this model of the actual world would have to be, by finding out about what some of the objects are that make up the actual Universe, or how large the actual Universe is and about some of the things that do or do not belong to the actual extensions of certain predicates. And it may be that certain sentences are either true in all the models that satisfy these constraints (and thus are candidates for the role of 'the model that corresponds to the actual world') or else false in all those models. Of these sentences we can then affirm that they are true in the actual world, and thus true simpliciter, and of others that they are false simpliciter.

This second option is that of *model-theoretic semantics* in the strict sense of this term. It was this option that Montague took in his seminal work on natural language semantics. The method comes from the branch of mathematical logic known as *Model Theory*, where it was developed first and foremost by Montague's mentor and Ph. D. supervisor Alfred Tarski. One of the attractive features of the model-theoretic method is that it provides us with a conceptually satisfying characterization of the notion of logical entailment, or *logical consequence* to use the term usually reserved for this relation: a sentence C is entailed by sentences A_1, .., A_n iff in every model in which A_1, ..,A_n is true, B is true as well. In other words, logical entailment is preservation of truth from premises to conclusion in all possible models. Montague was of the opinion that it was essential for a semantic account of any language that it provide a conceptually well-founded characterization of logical entailment for that language. In particular, he thought that our most reliable semantic intuitions about the natural languages we speak are intuitions about what is entailed by what. It is through the predictions that a model-theoretic account of the language we speak makes about entailments that the theory can be tested: Do those predictions match our intuitions?

On this view the H&K approach wouldn't be quite enough as it stands. It would have to be supplemented with an explicit specification of a class of models, in which one 'designated' model would correspond to the actual world. Not that it is difficult to do this once a truth definition has been given in H&K format. But there is nevertheless an important difference in methodological stance between the two options, which is all too often swept under the rug.

Both options described above are also available for someone who wants to give a formal semantics for a DRS-language. As it was, the original formulations of DRT adopted the model-theoretic option. I personally cannot see a good reason against this. So in these notes we adopt the model-theoretic perspective too.

For a DRS language all of whose DRSs are simple, this means that we simply generalize Def. 1 to arbitrary models M.[4]

Def. 2 A simple DRS K is *true in* a model M iff there exists a function f that maps the Universe of K into the Universe $U_{(M)}$ of M, so that for each condition $P(d_1, .., d_n)$ in the Condition Set of K the elements f(d_1),..,f(d_n) of $U_{(M)}$ satisfy P in M.

2.2 Complex DRSs

One sense in which simple DRSs are intuitively simple is that they can be visualized as small worlds in the sense explained above. But when we want

[4]I will not go into the explicit definition of the model class $\mathcal{M}$ that consists of all and only the models M we will consider. There is nothing particularly interesting or surprising about such a definition, but it involves some set-theoretic technicalities that it would be awkward to have to go into here and without any significant profit. But there is no real need for such a definition.

a DRS-language that can capture the expressive power of quantification in a natural language like English, then we cannot limit ourselves to just simple DRSs. We need DRSs which have *complex DRS-Conditions* as well as simple ones. This is so for instance for sentences involving universal quantification. As our first example of this let us consider the sentence (1.6.e), which we repeat:

(1.6.e) John offended every linguist.

When we analyzed this sentence along the lines of H&K, we assumed as point of departure for the semantics the LF structure in (1.15), which we also repeat:

(1.15)

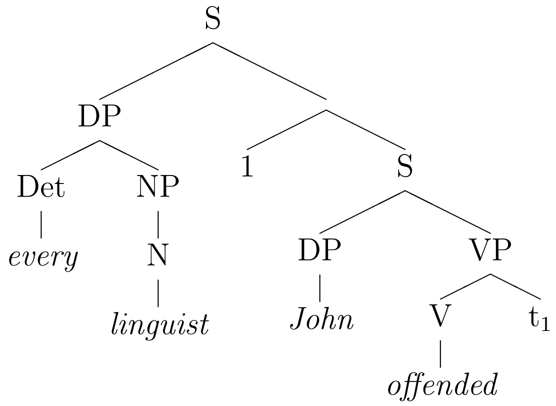

The treatment of this sentence in DRT does not require Quantifier Raising, and we will first show how its DRS can be obtained just from the surface structure, first given in (1.12). (2.9) gives the starting position for the DRS construction for (1.6.e), in which (1.12) is the only condition.

(2.9)
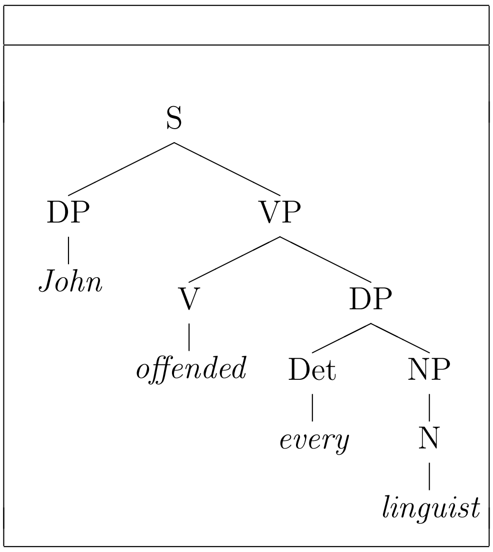

The first construction step is an exact replica of the first step in the DRS construction of (2.1) and we give the result without comment.

(2.10)
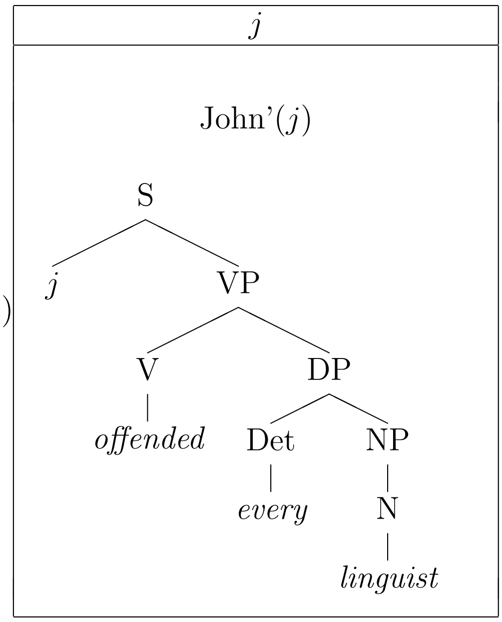

The next step must deal with the combination of verb and direct object DP and it is here that we see something fundamentally new. It is easy to verify – from the form of the discourse or from the truth conditions of the DCRS that we constructed for it – that all information involved in that sentence pair is analyzable in terms of existential quantification and conjunction over atomic predications. And that is the only information that can be represented in simple DRSs. Universal quantification is among the logical constructs that do not fit this general pattern – it cannot be reduced to combinations of existential quantification and conjunction. Therefore it must be given its own representational format. There just isn't anything else we can do.

The format that DRT has chosen for this purpose is that of a so-called *duplex condition.* A duplex condition is a structure that does justice to the fact that in natural language quantification typically takes the form of a 2-place operator, which combines two predicates - recall our first lexical entry for *every* in our discussion of MG. A duplex condition consists of two DRSs, one for the descriptive content of the quantifying DP and one for the predicate to which the DP is an argument. These two are connected by a third component, which is the actual quantifier. This third component contains (a) an indicator of the quantificational force of the duplex condition – in the present case, where the quantification is universal, this is the symbol $\forall$ for the universal quantifier[5] – and (b) the dref introduced to represent the DP. The combination of quantifier symbol and dref indicates that in the duplex condition the dref plays the role of bound variable. For reasons that will become clear later the dref is also introduced into the Universe of the restrictor DRS.[6] In the graphic display for duplex conditions we will be using the duplex condition resulting from the *every*-DP rule to verb and direct object of (2.10) looks as shown in (2.11).

[5]With other quantifying DPs, e.g. those beginning with the determiner *most*, it would be a quantifier expressing a different quantificational force (e.g. the one expressed by *most*).

[6]The fact that the dref occurs both as bound variable of the quantifier and also as member of a DRS Universe is likely to look suspicious to anyone with a proper education in formal logic: Doesn't this mean that the dref is bound twice over? The answer is: Not really. As the verification conditions for duplex conditions are formulated (see below) it is the binding by the quantifier that is directly relevant to the truth conditions. The quantificational binding overrules, so to speak, the semantic binding effects that are produced by the occurrence of drefs in DRS Universes in other representational configurations.

(2.11) 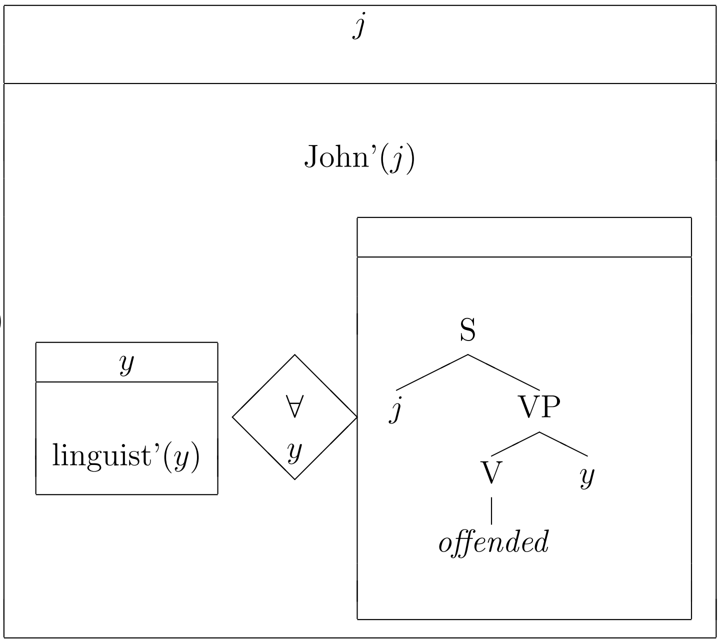

In general, the restrictor DRS or nuclear scope DRS of a duplex condition that has just been formed will contain conditions that are in need of further reduction. In which order those reductions are performed is then arbitrary. In (2.11) there are no reducible conditions on either side. All that can be done is to simplify the condition on the right, by throwing away the now superfluous syntax, with the result in (2.12).

(2.12) 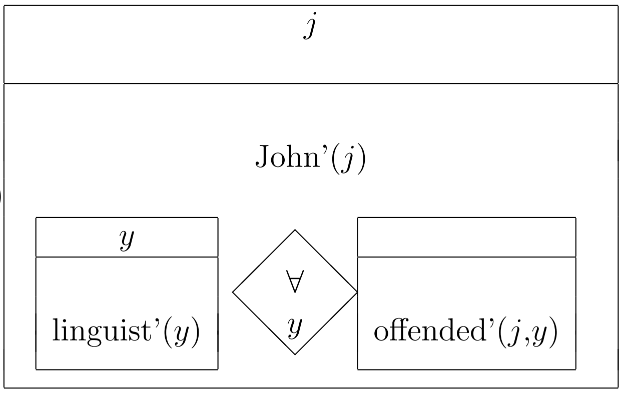

Although no QR is needed for the DRT-treatment of (1.6.e), it does not seem possible to do without QR altogether. The reason are inversely linked quantifiers as found e.g. in 'Some apple in every basket is rotten.' Since QR is needed for such cases, the simplest solution is to assume that QR applies to all quantifying DPs. In DRT, however, there is no need for the adjunction of indices below the raised quantifier phrases that is part of the implementation of QR in H&K. So we leave those adjunctions out. Instead we co-index the raised DP with the trace it leaves behind. For instance, for (1.6.e) this leads to the LF in (2.12).

(2.13)

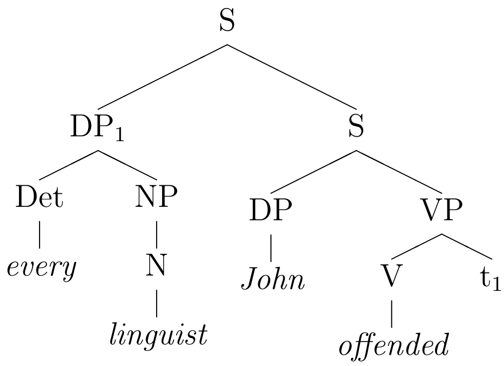

Let us see how the DRS construction for (1.6.e) goes, if we take this LF as the starting condition. That is, we now start with the DRS in (2.14).

(2.14)

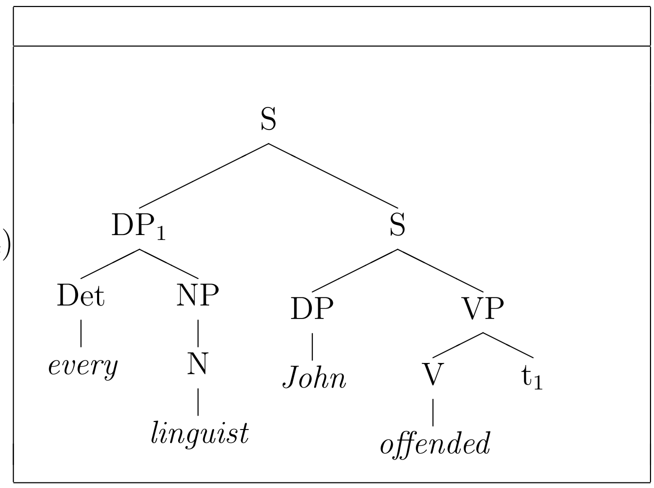

This time the first step is that which deals with the adjoined DP and thus the one that introduces the duplex condition. The only difference with the earlier application is that the dref chosen for the application – we assume that this is once again the dref y – now replaces the trace co-indexed with the DP that triggers the rule application. Otherwise everything is as in the application in the construction above:

(2.15)

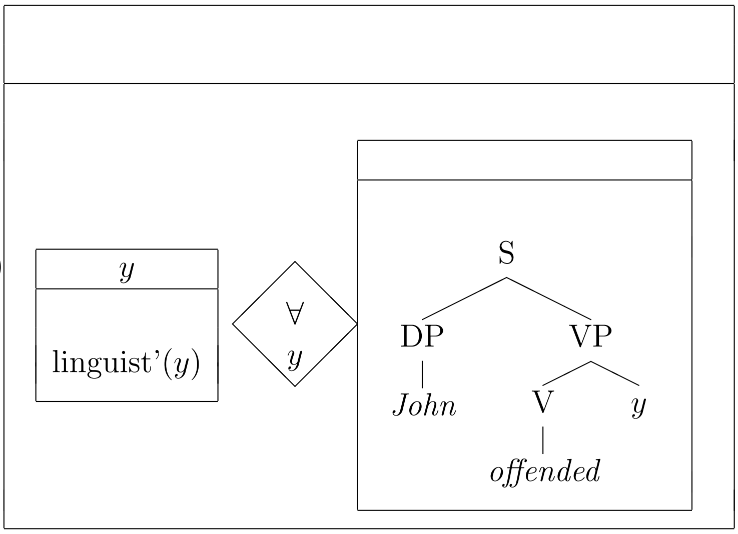

This time the condition in the nuclear scope DRS is reducible. The rule that needs to be applied is once again the one that deals with argument DPs that have the form of proper names. We have already applied that rule more than once. But the application before us raises a question that did not come up in our previous applications. It was said earlier on that the drefs that are chosen for the application of rules triggered by definite and indefinite DPs share the property that they must be placed in a DRS Universe that is either at the same level as the syntactic structure to which the rule applies or else at some higher level. At that point the statement made little concrete sense, since in the DRSs we were dealing with there was only one level, and one Universe. But with the advent of duplex conditions that situation has changed. In fact, there is a point connected with the new situation that we could already have made when going through our first DRS construction for (1.6.e). Consider the DRS (2.10), to which the duplex condition-introducing rule was applied, and the DRS (2.11) that resulted from that application. In (2.11) the dref y that was chosen for the application belongs to the Universe of the restrictor DRS of the duplex condition, and that Universe is at a level *below* the one of the condition to which the rule is applied. Intuitively and formally this means that seen from that latter level – that of the main DRS in (2.11) – y is no longer *accessible*: it does not play the role of the representative of some individual from the perspective of that level, but rather that of a quantificationally bound variable. We will presently see the implications of this in connection with the interpretation of anaphoric pronouns.

But right now our concern is with the placement of the dref j that will be used in the application of the proper name rule to the DP *John*. Given the position of the condition to which the rule is being applied - in the nuclear scope box of the duplex condition, all three DRS Universes in (2.15), the 'main' Universe at the top of the whole DRS displayed in (2.15), the Universe of the nuclear scope DRS and the Universe of the restrictor DRS count as 'at the same level or higher'. (That that is so also for the Universe of the restrictor DRS has not yet been argued, but we will soon come to that.) To which of these Universes should j be added? The answer is: to the main DRS Universe. Once again, we are not yet in a position to motivate this answer properly; that will have to wait for the presuppositional account of proper names and other definite noun phrases. For now we will just have to make do with the plain answer.[7]

With this last answer we know all we need to know in order to apply the proper name rule in the case before us. The result is given in (2.16.a) and the simplified version of that in (2.16.b)

(2.16)a.

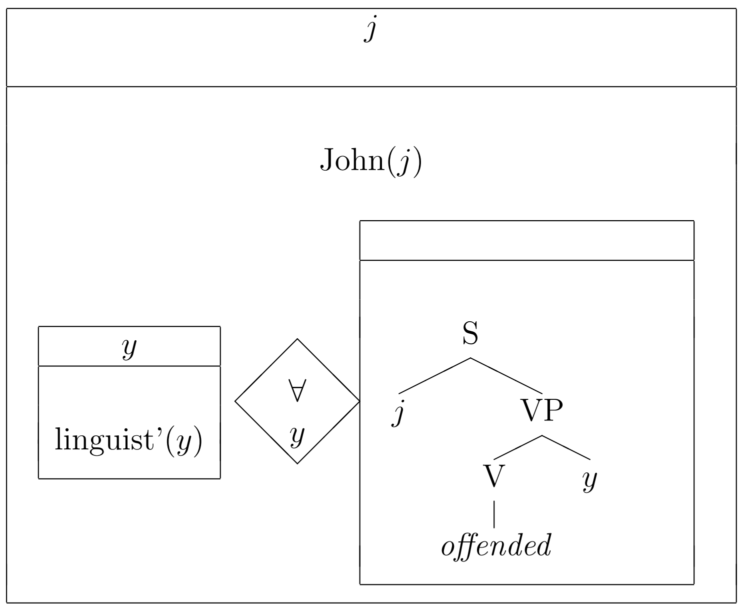

[7]Informally, the motivation can be stated along the lines indicated earlier: The most common use of a proper name, of which the occurrences in the examples considered here are all instances, serves to refer to some particular bearer of the name, and that is so irrespectively of where the name occurs in the sentence of which it is part. Because of this the name will always have maximal scope, irrespective of its syntactic position. The best way to capture this with the means available to us now is to make its representing dref an element of the outermost DRS Universe – in other words, of what we have been calling the 'main Universe' of the DRS.

b.

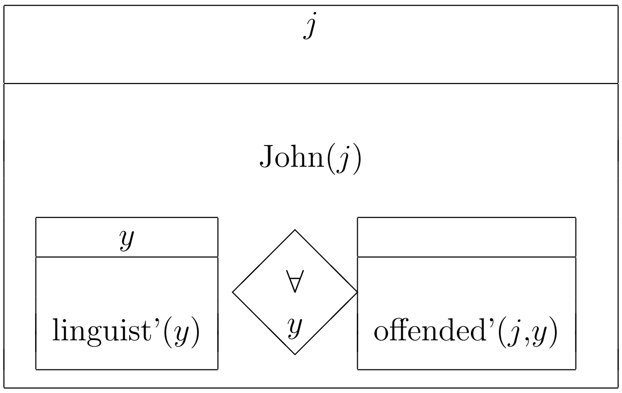

Note that this is the same DRS we obtained when we started the construction from the syntactic representation without raising (see (2.12)). And of course that is as it should be. With just one quantifier, raising cannot make a difference.

Before we continue our exploration of DRS construction for particular examples let us first return to the question of truth conditions for DRSs. The DRS in (2.12) and (2.16.b) is our first DRS with a complex DRS condition. What is it for such a DRS to be true in a model M? Let us refer to this DRS as 'K' and try to apply Def. 2. This definition tells us that K is true in M iff there is an embedding f of the Universe of K into the Universe of M which verifies all the conditions of K in M. The first question we need to answer is: what is in this case the Universe of K? For K displays several universes. The answer to this question is: the main Universe of K; from the perspective of K as a whole the drefs in the other two, *subordinate*, Universes do not function as representatives of particular individuals but rather – we already made this observation – as bound variables. This means that according to Def. 2 K is true in M iff there is an embedding f defined just for the dref j such that f(j) satisfies all the conditions in the Condition Set of K. Once again we now have a potential source of ambiguity. What in the case of K is its 'Condition Set? The answer parallels that we just gave to the question about the Universe of K: the Condition Set of K is the set displayed below the horizontal dividing line. In the case of K it consists of two conditions, the atomic condition 'John(j)' and the duplex condition below it.

So then: K is true in M iff there is an embedding f defined only on j with the property that it verifies these two conditions. What it is for f to verify the condition 'John(j)' in M we already know: f(j) must be the individual in M that the name *John* was used to refer to in the given utterance of the represented sentence. But what is it for f to verify the duplex condition of K?

Nothing we have said up to this point provides an answer to this question. So, in order to make Def. 2 applicable to DRSs like K we have to extend it with a clause that says what it is for an embedding to verify such duplex conditions.

What are we to say? Well, intuitively, given that we know what universal quantification means, it is pretty clear what we ought to say: for f to verify the duplex condition of K in M it should be the case that no matter how we extend it to an embedding function g which is also defined for the dref y and which verifies the conditions in the restrictor DRS, that g must also 'verify the nuclear scope DRS'. This last part of the formulation has been placed in scare quotes, because the notion of an embedding hasn't yet been defined for sub-DRSs of a bigger DRS. In the case we are considering, it is intuitively clear what we ought to say: each of the embeddings g just spoken of should verify the one condition of the nuclear scope DRS, viz. 'offended'(j,y)'. If this is how we explicate what it means for g to 'verify the nuclear scope DRS', then we have obtained the intuitively correct verification conditions for the duplex condition, and with that the correct truth conditions for K. For our application of Def. 2 to the present case now says that K is true in M iff the referent in M of the given use of the name *John* has the property that any way of assigning to y an individual of M that is a linguist also has the property that John offended him or her.

This works for the special case of the duplex condition of K. But it won't be applicable to arbitrary duplex conditions, not even to arbitrary conditions with the force of universal quantification. For in general, both the restrictor DRS and the nuclear scope DRS of a duplex condition can be of arbitrary complexity. Their Universes may contain any number of drefs and their Condition Sets any number of conditions (and those conditions may be complex conditions (e.g. duplex conditions) in turn). It takes some reflection to see how the verification conditions for universal duplex conditions should be stated so that they deal with this general case, and I just give the formulation here without going into further motivating detail:

(2.17) Let M be a model, and f an embedding function into M. (That is, f is a function which is defined for some set of drefs and which maps each of these to an individual from the set U_M.) Let C be a duplex condition of the form:

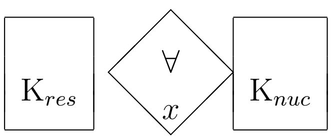

Then:

f *verifies* C *in* M iff every extension g of f to x together with the other drefs in $U_{K_{res}}$ which verifies in M all the conditions in the Condition Set of K_{res} can be extended to an embedding h which is defined in addition for the drefs in $U_{K_{nuc}}$ and which verifies in M all the conditions of the Condition Set of K_{nuc}.

Below we will see some examples which show the point of formulating the verification conditions of universal duplex conditions as in (2.17).

Note well that (2.17) only applies to universal duplex conditions. It is obvious that duplex conditions whose quantificational force is different, such as for example those whose quantifier is 'most', demand a different verification condition, which does justice to their quantificational force. But the generalization is less straightforward than one might have thought, because of the so-called *Proportion Problem*. For discussion see (Heim 1982,1988) and (Kamp & Reyle 1993).

A final remark before we return to matters of DRS construction. Comparison of the treatment of a sentence like (1.6.e) in H&K with either of the two DRT treatments presented of this sentence above may create the impression that DRT is much simpler when it comes to such sentences. But to a large extent that impression is misleading. A good part of what made the calculation of the truth value of (1.6.e) along the lines of H&K seem complicated had to do with the fact that what one is calculating there is the actual semantic value s of the sentence and its constituents. The construction of a DRS for a sentence does not accomplish that. The DRS is only a 'logical form' for the sentence, a structure that presents its semantic properties in a logically transparent form. That these 'logical forms' are 'logically transparent' is shown and made explicit by the model-theoretic truth definition for the 'logical form language' – in other words, the DRS language. And if we want to arrive at something like the truth value of (1.6.e) that is the final result of the calculation we went through when rehearsing the H&K treatment, then we will have to apply the truth definition for the DRS language to the logical form that DRT assigns to (1.6.e) – that is, to the DRS in (2.12)/(2.16.b) – and some suitable model M, which we take to reflect the real world. That is a good deal of additional work, which includes spelling out what it is for an embedding f to verify the duplex condition of this DRS.

This is one fundamental difference in perspective between the 'logical form' approach of DRT on the one hand and the semantic value approach that H&K share with the model-theoretic accounts of classical Montague Grammar. DRT does in two steps what these other accounts do in one. It factors the determination of the semantic values of sentences, discourses and other well-formed expressions into (i) a reorganization of the information that is contained in the sentence, discourse, etc. in a format that allows for an 'intentional' or 'referential' semantics in the smooth and simple way that we are familiar with from the artificial languages of formal logic, such as the predicate calculus or the typed λ-calculus.

This factorization has proved its practical usefulness. There is a wide variety of linguistic problems – the different uses of plurals as opposed to singulars, the intra- and inter-sentential effects of tense and aspect, the form and resolution of presuppositions, the semantics and pragmatics of information structure, the structure and content of propositional attitudes, the role and classification of discourse relations – where a DRT approach can be helpful. The reason why it can be helpful in dealing with these problems has to do with combination of two things. On the one hand DRSs are semantically and logically transparent in much the same way that the languages of formal logic are (like Predicate Logic, with which the representation format of DRT has much in common, as we noted, or the various versions of the typed lambda-calculus). This transparency is revealed by the model theory for the DRS languages to which these DRSs belong, though after a while of working with DRT this becomes sort of second nature to the user. On the other hand the principles of DRS construction show how the particular syntactic structures, lexical items and morphemes to which they apply make their contributions to the semantics of the sentences, texts and discourses to which they belong. They show the connectional between those sentences and so on, and the DRSs that they assign to them as semantic or logical forms. New syntactic constructions, words or morphemes, that hadn't been considered hitherto, may require new Construction Rules; but if these rules are identified in the right way they reveal the semantics of those new words etc.

Speaking more generally, the usefulness of DRT has much to do with its conception as a theory of *language interpretation* and not just as a theory of semantic values of object language expressions (and, more particularly, of the truth conditions for sentences). It is well to recall in this connection that DRT was intended from the start not just as a logical form approach to natural language semantics, but as one in which logical forms have a cognitive significance. It is because of their cognitive adequacy, the suggestion has been from the start, that DRSs can serve not only as identifiers of semantic content but also, and at the same time, as discourse contexts, which guide the interpretation of what comes next.

But apart from such practical advantages that a logical form approach like DRT may arguably have, there is also a more fundamental conceptual difference between it and the different 'semantic value' approaches we have referred to and seen a few applications of. Logical form theories are open to an interpretation that semantic value approaches are not open to (and, I think, it is fair to say, would not want to be open to), that according to which the logical forms that it assigns to sentences, texts etc are assumed to reflect, in some way, and at some suitable level of abstraction, the representational forms of the thoughts that interpreters extract from the sentences, texts etc that they interpret.

Thus, from the very start it has been part of the contentions of DRT that DRSs can perform the double duty of content representations and discourse contexts, because they capture, in some way and at some suitable level of abstraction the representational forms of an interpreter's thoughts. The difficulties and dangers that come with such a psychology-loaded conception of natural language semantics have often been commented on, and in particular they have been the focus of critical comments on DRT. Natural language semantics, it has been widely thought, should keep clear from any considerations of what is going on in the minds of those who use it. Rather, if there is anything we can say about what is going on in speakers' minds, then the results of natural language semantics should be among the things that help us explain how minds work when they deal with language.

An important question for one who advocates a cognitively non-neutral approach to the semantics of natural languages concerns what is concealed behind the qualifier above: that DRSs 'reflect, in some way, and at some suitable level of abstraction, the representational forms of an interpreter's thoughts'. Which properties of the logical forms that the theory proposes are supposed to be psychologically relevant and which are to be considered just artifacts of the chosen formalization? That is a hard question and it is a particularly hard one for an adherent to DRT today, because since DRT started there has been quite a bit of change in the forms that DRSs are supposed to have. In spite of all this I remain convinced that DRT is a valuable method not only for doing linguistics but also for investigating properties of the mind, especially those that manifest themselves in the us of language;

And I also believe that the study of how we express things in language can teach us things about this properties, even if the clues about mental structure that are hidden in linguistic structure can be hard to detect and decipher. In fact, my personal feeling is that the different transformations that DRT has gone in the course of the three and a half decades since its first formulations have brought us closer to an understanding of the representational properties of human linguistically expressible thoughts rather than dissipate DRT's original impulse in this direction and aren't just an admission that the approach got things wrong to start with and has been floundering ever since.

If this conviction is justified, however, then it points to what seems to me a real puzzle: If it is true that DRSs – or for that matter any other account of logically streamlined logical forms – capture important formal properties of how we represent content in thought, why do we what appear to be such quite different forms when we express those same contents in words? Why does our language not have that same logical transparency that such theories attribute to the forms in which we think? I believe that this is a very hard question. But it is not one that should be taken as a simple reductio of the conjunction that it queries. Why language should do things so differently than we might expect if we believe that our thoughts are well organized along the lines of the streamlined calculi of formal logic may have a variety of reasons. The most important one of these may well be that expressing thoughts was not what human language was developed for in the first place. It seems quite plausible that what made it possible for our species to develop into a species-with-language were communication-related advantages that had comparatively little to do with the expression of sophisticated content. By the time that language developed into a tool that was capable of doing that as well, its organization had already been fixed to such an extent that this new capacity – of expressing logically sophisticated thoughts in words – had to make do with a 'language engine' – a language-related cognitive architecture – that was already in place – not an optimal one perhaps for this particular kind of task, but better than nothing; and otherwise there was nothing.

If a story along these lines has any plausibility, the perhaps we have some right to be proud about how good we have become at navigating the streams of thought in this sub-optimal vehicle. But on the other hand we might have been much better at this if language and cognition didn't put so many obstacles in our way.

2.3 More Complex DRSs

We return to ground level and to the construction of DRSs, still by way of the top-down method. One of the uses of QR, we noted in our brief perusal of the H&K method, was that it can account for quantifier scope ambiguities in terms of how the scope bearing DPs are raised: Once an LF has been obtained in this way, that fixes the quantifier scope relations. But more than one such LF can be obtained from the given input string, and that is what makes the string ambiguous.

By way of example let us return to sentence (1.13), repeated here once more.

(1.13) Some philosopher offended every linguist.

The two possible ways of doing the Quantifier Raising for this sentence lead to the syntactic structures in (2.18.a,b). These are like the ones we considered in our earlier discussion of this sentence, but without the 'number adjunctions' that H&K use as triggers for the λ-abstractions that are part of their way of handling the semantics of quantifying DPs (compare (1.34)).

(2.18)a.

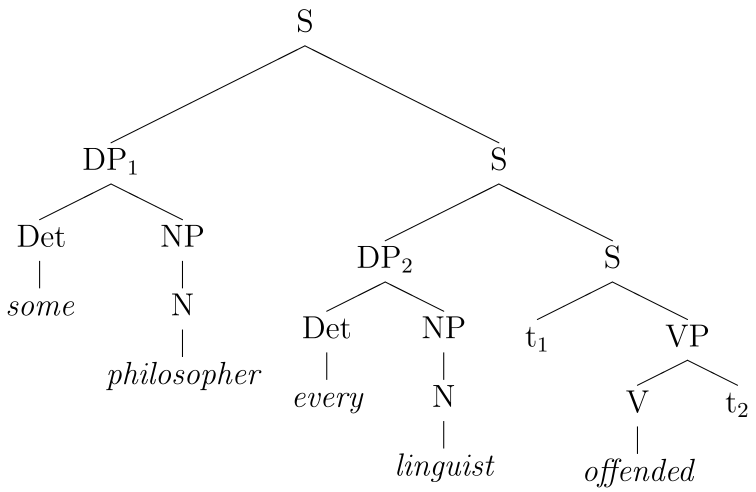

b.

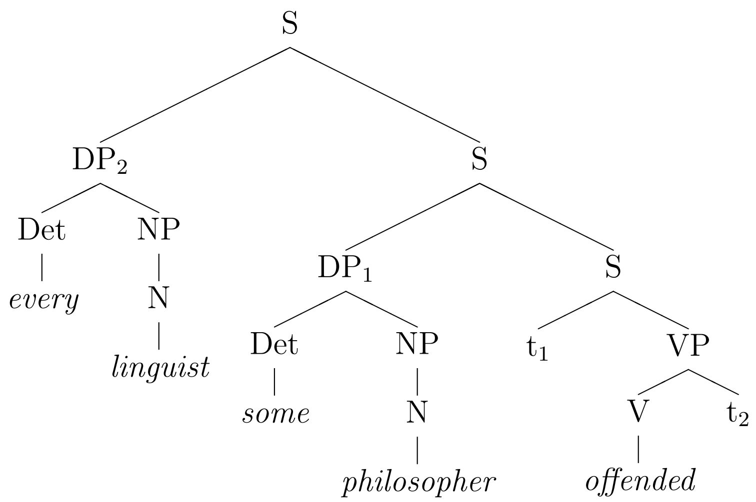

After the examples of DRS construction we have been through it is now more or less clear how to construct DRSs for these two LFs. But there still is one point that deserves explicit attention. It arises in connection with the interpretation of (2.18.b). The first step in the DRS construction for this LF introduces a universal duplex condition and the result is as in (2.19).

(2.19)

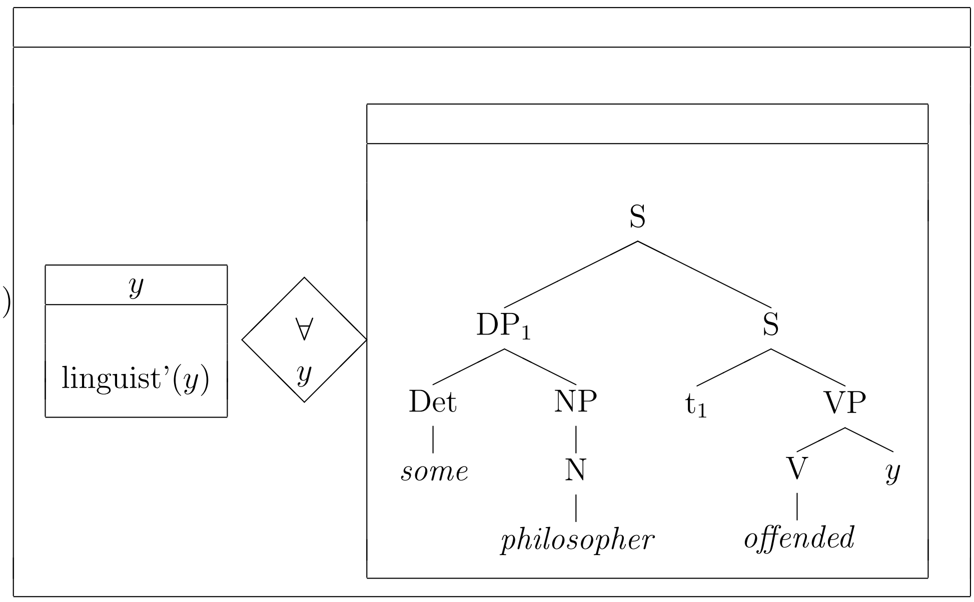

The next step must deal with the indefinite DP *some philosopher*. The one issue that hasn't been settled by what we have said about the treatment of indefinite argument DPs earlier is into which DRS Universe the dref chosen for the execution of the rule – let us use x – should be put. What definite and indefinite DPs have in common, we said earlier, is that their drefs must be place in a Universe at or above the level at which the rule is being applied. Subsequently it was stated (if not fully argued) that the drefs of standardly used proper names always go into the main Universe. With indefinites, we now add, the default is the opposite of this: the dref should be introduced into the Universe of the (sub-)DRS whose Condition Set contains the condition to which the rule is applied. Following this recipe we obtain the DRS in (2.20).

(2.20)

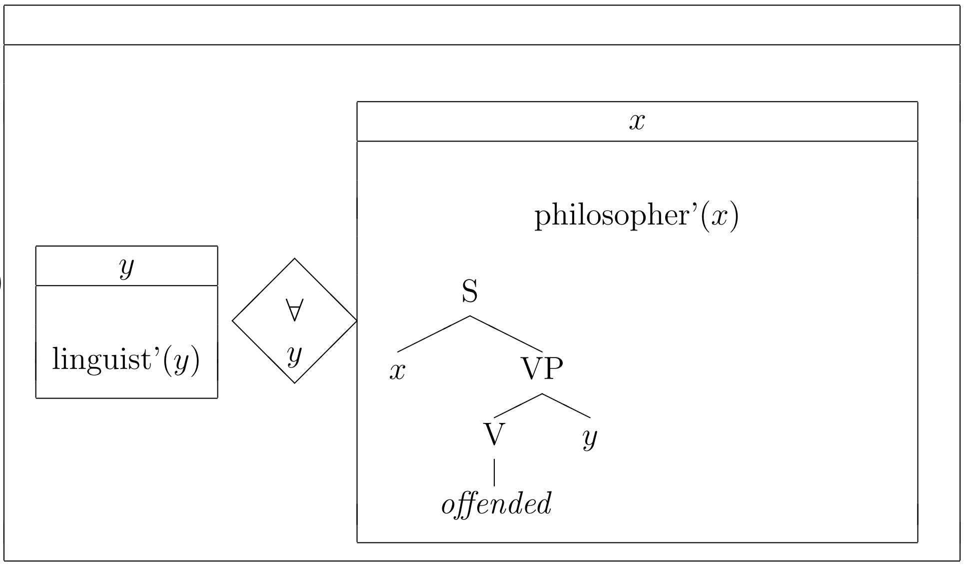

Exercise Using the same rules construct the DRS for the LF in (2.18.a).

The principle that the dref for an indefinite DP should be placed in the Universe of the DRS within which the rule for indefinite DPs is applied (more correctly: the DRS whose Condition Set contains the condition to which the rule is being applied) is not absolute. Indefinites can be used *specifically*, in the sense that the speaker has a particular individual in mind that she wants to talk about but uses an indefinite in order to avoid creating the impression that she thinks her audience can identify which individual that is. When an indefinite is understood as used in this way, then its dref should be entered into the Universe of the main DRS, to reflect the intuition that the utterance the speaker has made is about one particular individual, even if the interpreter doesn't quite know which individual that is.

If we allow this option for the DP *some philosopher* in (2.18.b), then we can construct a DRS that represents the ∃∀ reading even from this LF. (Exercise:

Show this!). In other words, there are two different ways in which to obtain this reading: as only possible reading according to (2.18.a) and as one of two possible interpretations of (2.18.b).

Here we reach the point where we should ask ourselves if this isn't too much of a good thing. In fact, there is a more fundamental question: Is it really defensible to subject indefinites to QR if at the same time we treat them as (largely) on a par with proper names and other definite noun phrases? This feels a bit like eating one's cake and having it. But perhaps there are really two cakes here – that indefinites are ambiguous between an interpretation as (existential) quantifier phrases and an interpretation as 'singular terms'. Your first reaction to this may be that this can only be the self-indulgence of a theoretician who is trying to make a virtue out his inability to make up his mind. But there are actually quite good reasons for such an assumption. (See in particular (Fodor & Sag 1982), (Kratzer 1998).)

In fact, the discussions over the correct interpretation of indefinites do not stop at the choice between quantifying DPs and terms. A third proposal that has had a good deal of currency is that indefinite should be analyzed as *choice functions*, that is functions from non-empty sets to members of those sets. More precisely, the interpretation of an indefinite of the form *a/some N* involves the introduction of a choice function variable f, assigns (S_N) to the argument position occupied by the indefinite and gets existentially bound somewhere in the resulting logical form, with narrow scope binding capturing non-specific and wider scope various forms of specific interpretation. (The two papers that introduced the choice function analysis of indefinites into formal semantics, (T.Reinhart 1997)and (Winter 1997), appeared at the same time; the proposals they make have much in common, but are not identical.) It would carry us too far to go into any of this in detail. In particular there is no question of giving detailed reasons supporting the decision to allow for the possibility of interpreting indefinites either as quantifying DPs or as terms. This is the assumption on which we will operate. Those who want a justification for it should consult the cited papers and the literature mentioned in there.

For (1.13) there are now two options: (i) *some philosopher* is interpreted as a quantifying phrase and (ii) *some philosopher* is interpreted as a 'singular term' (i.e. along the lines indefinites are treated within DRT). According to the first option *some philosopher* should be subject to QR, just like other quantifying DPs. Let us assume, as we have done so far, that QR of subject and direct object may occur in either order, giving rise to the two LFs in

(2.18). But if *some philosopher* can be interpreted as a genuine quantifier phrase, then presumably it too should (when interpreted in this way) give rise to a duplex condition in the DRS constructed from either of these LFs. For instance, (2.18.b) would yield the DRS in (2.21).

(2.21)

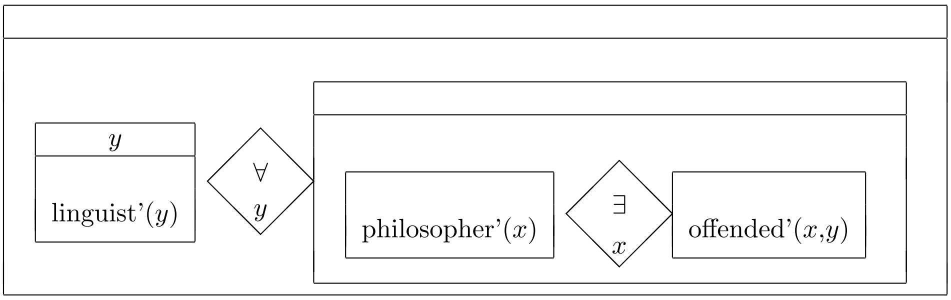

We now have a new type of duplex condition, so we also need a corresponding clause in our truth definition for DRSs to tell us what it is for an embedding f to verify such an ('existential') duplex condition in a model M.

<u>Exercise</u> State this clause.

<u>Exercise</u> Construct a DRS like the one in (2.21) for the LF in (2.18.a).

In addition we now also have the possibility of interpreting *some philosopher* as indefinite singular term. If we stick with our decision to apply QR only to genuine quantifying DPs, then on this interpretation *some philosopher* should not be raised, and we get just one LF for (1.13):

(2.22)

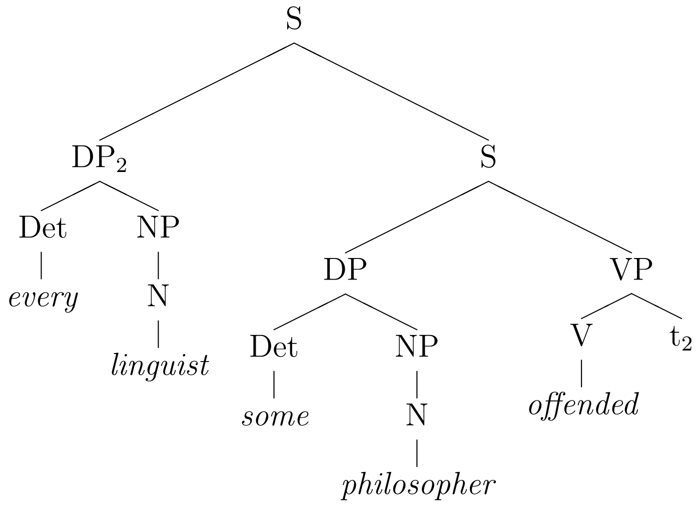

From (2.22) we can then construct two DRSs, the one in (2.20), which we get when we interpret *some philosopher* as a non-specific indefinite, and the one left as an exercise earlier, in which the dref x has been placed into the main Universe, reflecting the assumption that *some philosopher* is being used specifically.

2.4 Back to Donkeys

We have already dealt with the 'donkey discourse' given in (2.1), repeated here in good order. But application of the same idea to the donkey sentences in (1.44), also repeated, had to wait until we had dealt with some further issues these sentences raise. Now that we have dealt with those, it is time to return to the central issue that donkey sentences were meant to present.

(2.1) Pedro owns a donkey. He beats it.

(1.44) a. If Pedro owns a donkey he beats it.
(1.44) b. If a farmer owns a donkey he beats it.
(1.44) c. Every farmer who owns a donkey beats it.

Our first concern is with (1.44.c). The subject DP of this sentence contains a relative clause. So far our only encounter with relative clauses was in the context of inverse linking, where we assumed that the '$\exists\forall$ reading for a DP like *some apple in every basket* results from analyzing the PP *in every basket* as remnant of a reduced relative clause. There our concern was with reduced relatives, but the differences between reduced and non-reduced relatives are minor. The only difference between the reduced relative of (1.38) and the non-reduced relative of (1.44.c) is that in the latter the V node and the Comp node are overtly realized.[8] There is also a further difference between the syntactic representation of the RC of (1.44.c) in (2.23) and that of the reduced relative in (1.38) in that we are now no longer making use of the 'numeral nodes' that in the system of H&K serve to trigger lambda abstraction.

(2.23)

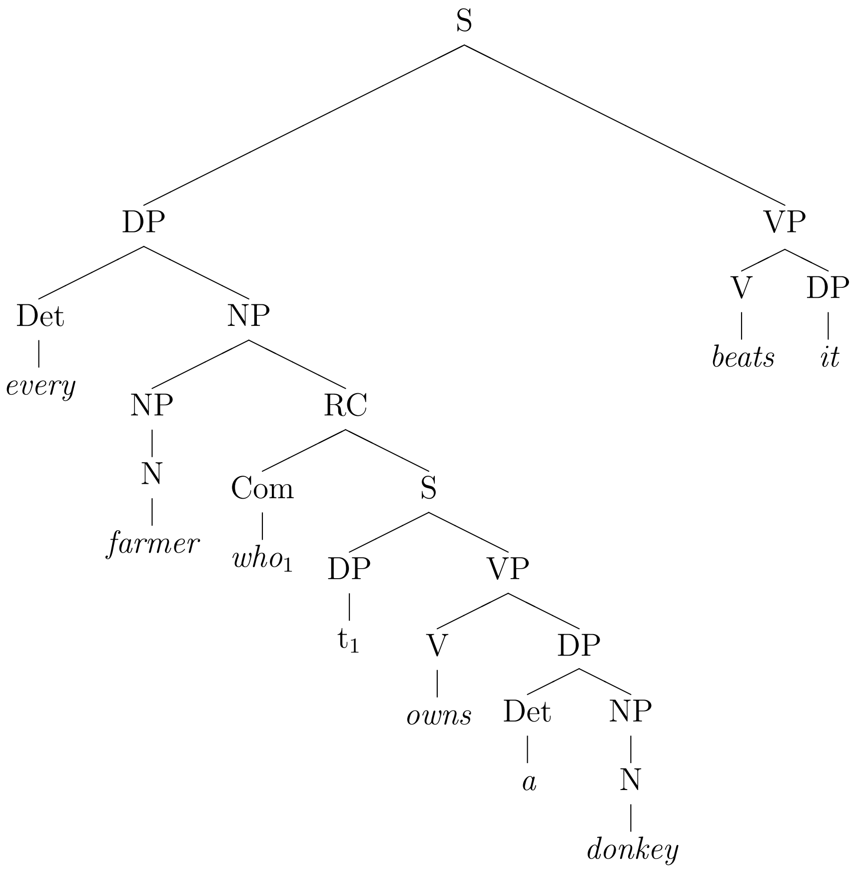

What exactly does (1.44.c) mean? Arguably the sentence is ambiguous between a reading according to which it says that for every farmer and any donkey that farmer owns the former beats the donkey. Arguably the sentence also has another reading in which *a donkey* has wide scope over *every farmer*, but that reading is exceedingly implausible, since it would imply

[8]One also finds a more limited form of reduction of relative clauses in English. In these only the relative pronoun is missing. This form of reduction is possible only when the argument position from which the pronoun is 'extracted' is not the grammatical subject position. A pair of examples are: *the woman John loved/gave the flowers (to)*.

that there was a single donkey that was owned by every farmer. So we will forget about that reading, and focus on the first one (which is the one that is assumed in all discussions of donkey sentences). So we are interested in constructing DRSs for (**??**.c) which capture this racing.[9]

How are we to construct such a DRS from the LF in (**??**)? That still depends on how we interpret the indefinite DP *a donkey*, as quantifying DP or as singular term. It is not hard to see, however, that a quantifying DP interpretation won't do in this case. For the existential quantifier it contributes to the DRSs will get narrow scope w.r.t. the universal quantifier contributed by *every farmer* (for details see below). But then the dref introduced for the indefinite won't be available as antecedent for the pronoun *it*. The DRS 'under construction' in (2.24) makes the predicament clearly visible:

(2.24)

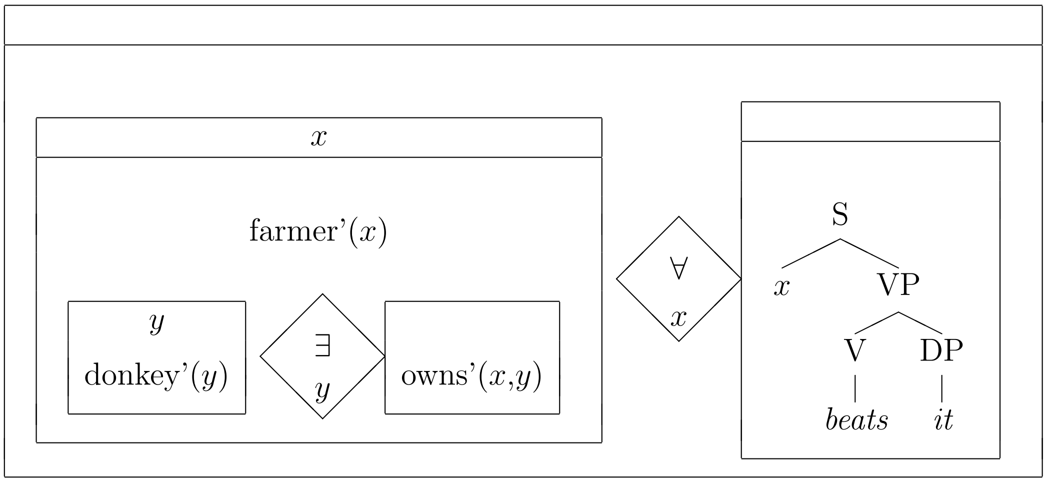

The problem with (2.24) is that the dref y, which we would like to use as anaphoric antecedent for the pronoun *it*, is not accessible from the position

[9] There is also a widely attested worry about what should be said about this sentence when it is used to describe situations involving farmers with more than one donkey. Is the sentence even felicitous when it is used to describe such a scenario? Some speakers say 'no' and others waver. (The following variant of (1.44.c): 'Every farmer who owns one or more donkeys beats them' does not provoke such reactions of doubt or refusal.) In DRT and a number of other approaches to the problems of donkey pronouns (though not all of them, see for instance (Cooper 1979), (Chierchia 1991), (Chierchia 1995)) this complication is ignored. The semantics we will give for this sentence treats it as applicable to situations in which farmers can have more than one donkey no less than in situations in they all have to most one; and in situations of the former type the treatment makes the sentience true only if every farmer beats each and very donkey he owns.

of the DRS condition that contains the pronoun (i.e. the one tree remaining in (2.24)). The principle appealed to here is that only those drefs are available as antecedents for the interpretation of a pronoun that is processed at a given level – that is: as part of a condition belonging to the Condition Set of some (sub-)DRS K – which are in DRS-Universes that are accessible from this Condition Set. One of these Universes is the Universe of K itself. But in complex DRSs, in which there are sub-DRSs as well as the main DRS, there may be additional Universes that are accessible from a given Condition Set. For an explicit definition of *dref accessibility* the reader is referred to (Kamp & Reyle 1993), Chs. 1, 2. However, the definition is pretty much what one should expect. What should perhaps be pointed out is that the Universe of the restrictor DRS K_{res} of a duplex condition is accessible from the Condition Set of the nuclear scope DRS K_{nuc}. But when a DRS occurs as proper sub-DRS of the retractor DRS of a duplex condition, then the drefs within the Universe of the first DRS are *not* accessible from the nuclear scope DRS. An example is the dref y in (2.24). y belongs to the restrictor DRS of a duplex condition that is a member of the Condition Set of the outer, universal duplex condition in (2.24). In this position y is inaccessible from the nuclei scope DRS of the universal duplex condition. So the pronoun *it* in the nuclear scope cannot be interpreted as anaphoric to y.

The situation would be different if we started from an LF for (1.44.c) in which the universal duplex condition was embedded within the existential duplex condition rather than the other way round. But this leads to the DRS that we don't want, which talks about a single donkey that might have been owned by several farmers. Our only hope, therefore, is to treat *a donkey* as a singular term.

But before we explore this alternative, we must first have a somewhat closer look at the construction that enables us to obtain (2.24) (even though, as we have seen, that construction cannot proceed any further). Our first task in detailing the steps of that construction is to identify the LF from which the DRS construction is to proceed. Since in constructing (2.24) we are assuming that both DPs, *every farmer* and *a donkey*, are quantifying DPs, both need to be subjected to QR. But we also assume, as in our discussion of inverse linking, that relative clauses are barriers to QR, i.e. that a DP within such a clause must be adjoined to the S below the RC node. With this restriction the LF for (1.44.c) is as in (2.25).

(2.25)

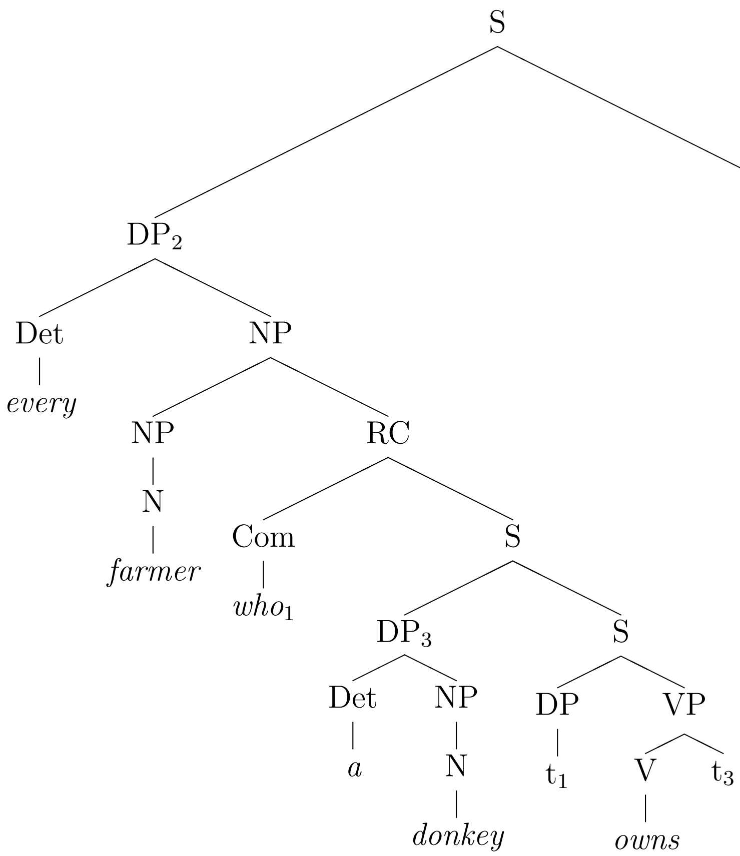

What we haven't yet dealt with the top down processing of combinations of NPs and relative clauses. The structures in (2.26) and (2.27) show how such combinations can be handled. (2.26) gives the result of applying the rule for universally quantified DPs. The only difference with earlier applications is that we now have to record the information that the dref introduced by the rule – x, in the application shown – must represent the individual or individuals to which the complex restrictor predicate applies. This predicate is given by the upper NP node of the subject DP. We encode the information by adding the chosen dref in parentheses behind this node. In this way we get (2.26).

(2.26)

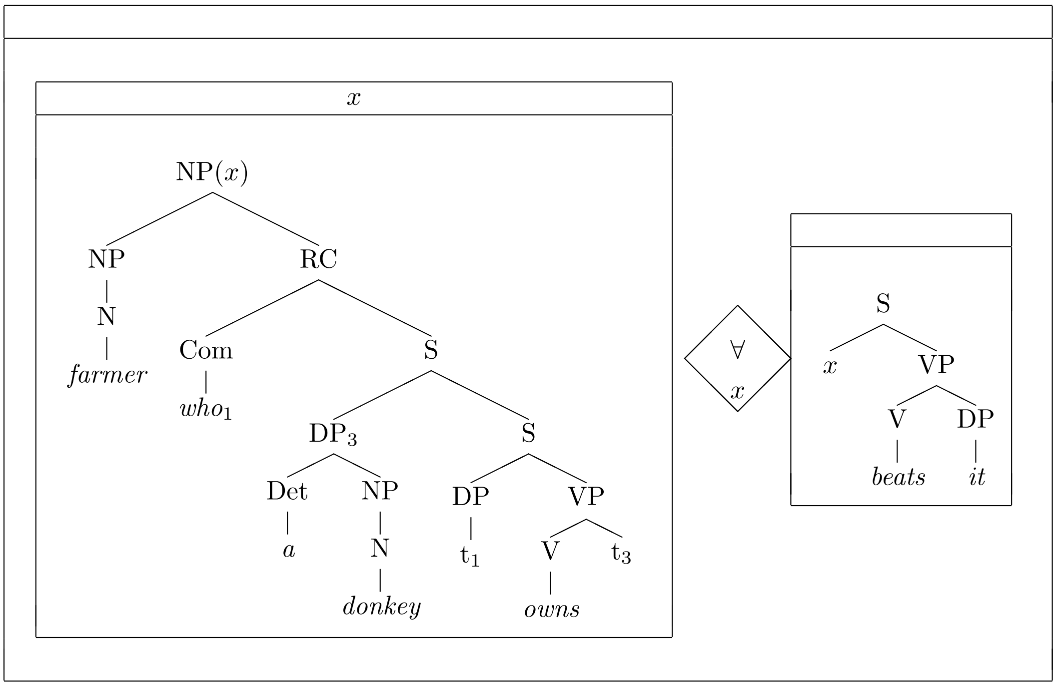

Reduction of the condition with root label 'NP(x)' takes the obvious form of treating the combination of lower NP and RC as a case of predicate conjunction. That is, the argument x is distributed over the lower NP and the RC. As before we reduce predicating the lower NP of x directly to the predication 'farmer'(x)'. Furthermore, the predication 'RC(x)' that results from this operation can, again quite obviously, be further reduced by inserting the argument x in place of the trace that is co-indexed with the relative pronoun *who*. The part of the syntactic structure above the S node of the RC can now be dropped, and as doing that makes the structure somewhat more perspicuous, that is what we do. Proceeding in this way we arrive at (2.27).

(2.27)

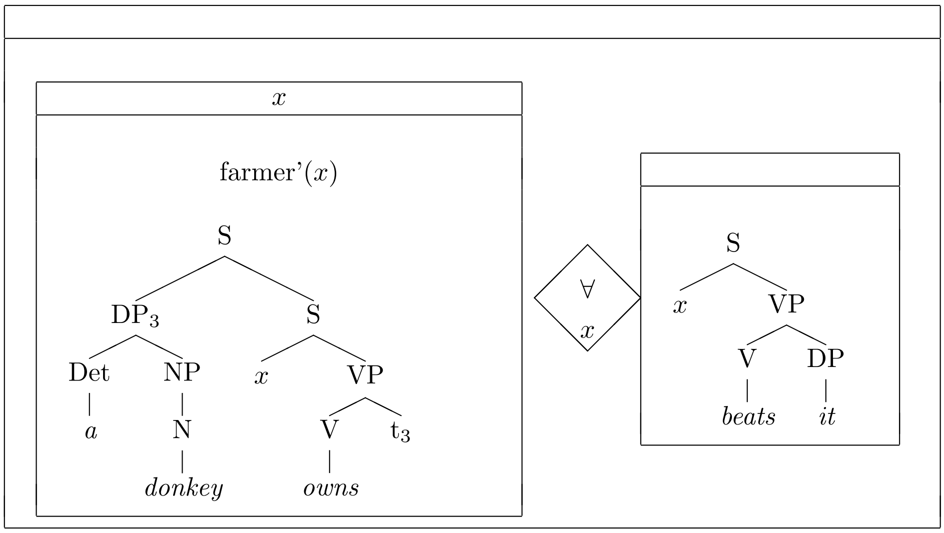

The remaining operation that has to be performed to reach (2.24) is the one that deals with the argument DP *a donkey.* Since in the construction of (2.24) this DP is treated as quantifying DP, its reduction leads to an existential duplex condition contained within the restrictor DRS of the universal duplex condition in (2.27). If y is chosen as dref for this operation, what we end up is indeed the structure given in (2.24).

Since as we saw (2.24) is a dead end, the only way to get a DRS for (1.44.c) that assigns it the truth conditions we want is to interpret *a donkey* as a singular term. In the light of what we have seen this is now more or less straightforward. A difference is that treating *a donkey* as singular term now carries with it that we do not QR it. That is, the LF from which we start is the one in (2.28).

(2.28)

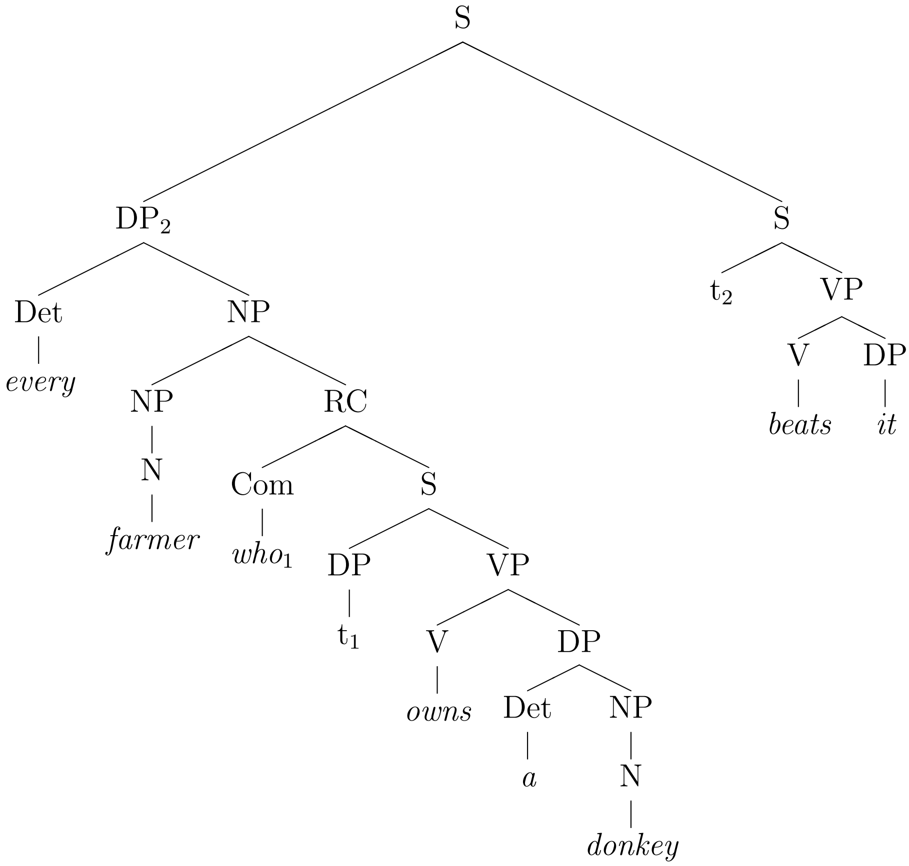

When starting from (2.28) application of the rule for quantifying DPs to the subject DP of (1.44.c) leads to the structure in (2.29).

(2.29)

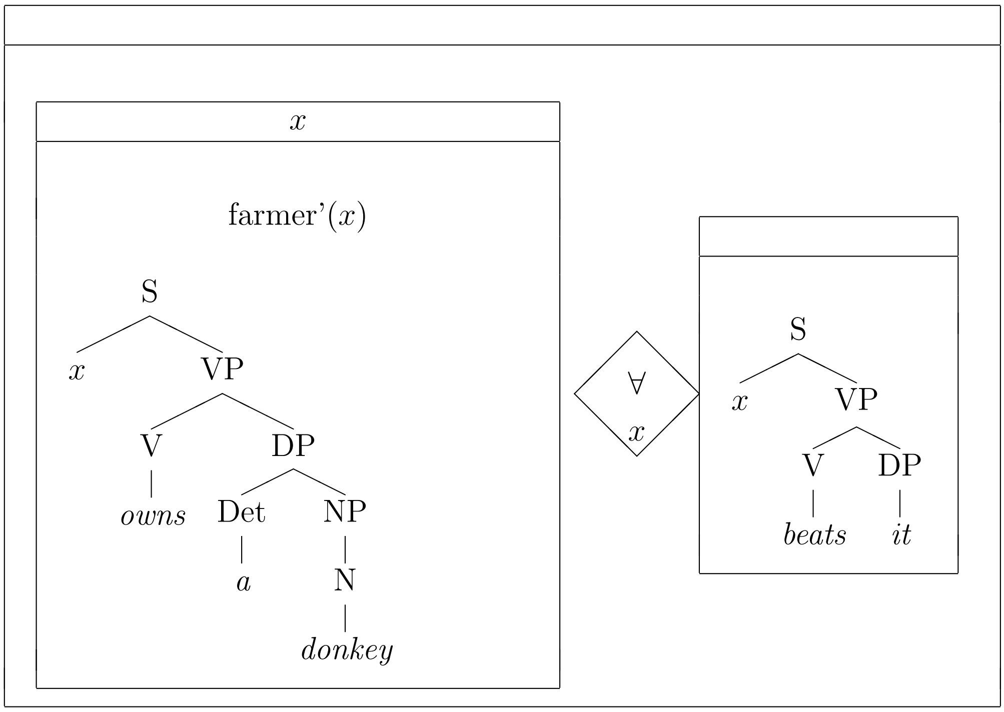

This time we treat *a donkey* as a singular term. As we have seen, this means that the dref y we choose must be inserted into the Universe of the DRS of the treated condition (or, in case the DP is treated as specific, into the main Universe, but this latter option leads to the counterintuitive interpretation we do not want). The result of this non-specific interpretation of *a donkey* is given in (2.30).

(2.30)

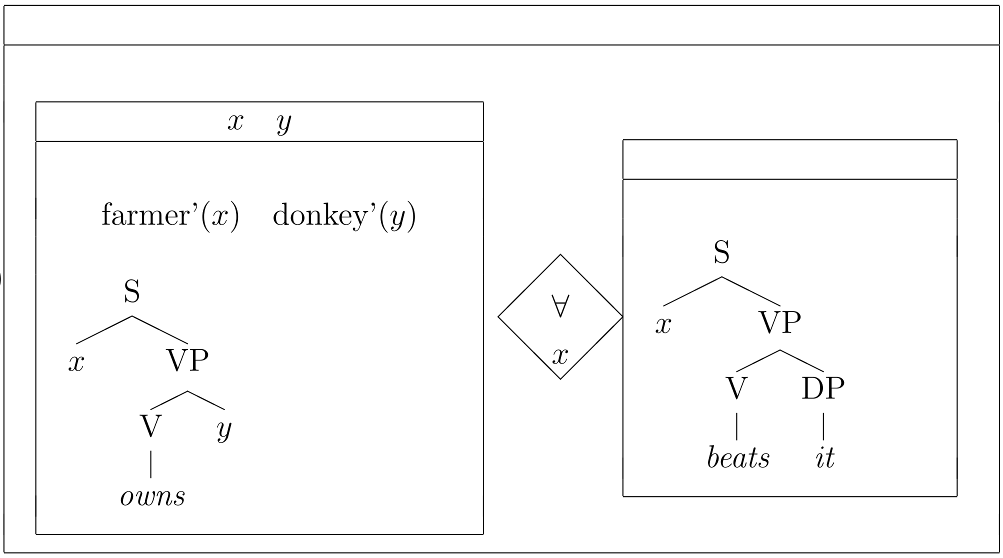

In (2.30) the dref y is accessible to the condition containing *it* and so the intended pronoun interpretation is possible. The result is shown in (2.31), where we have also eliminated the remaining bits of syntax inside the restrictor and nuclear scope DRS.

(2.31)

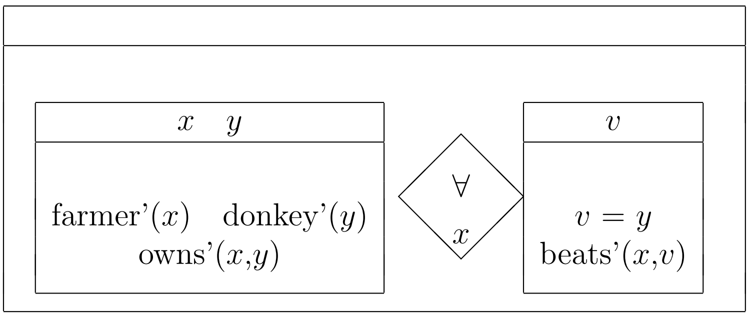

The main moral of this story is that in order to get the wanted interpretations of a donkey sentence like (1.44.c) we need to treat the indefinite as a term, and not as an existential quantifier. If one insists on treating the indefinite antecedents in such sentences as existential quantifiers nevertheless, then a quite different treatment ids needed for the pronoun, along the lines of E-type or D-type accounts.

What we have just observed for quantificational donkey sentences such as (1.44.c) also holds for the conditional variety exemplified by (1.44.a) and (1.44.b).

We first look at (1.44.a) and then briefly at (1.44.b). Both (1.44.a) and (1.44.b) have the form of a conditional consisting of an *if*-clause and a main clause. Furthermore, the Comp-position of the *if*-clause is realized as the particle *if*, which marks the clause as the antecedent of a conditional whose consequent is the clause to which the *if*-clause is adjoined. Thus the syntactic structure we assume for (1.44.a) is the one given in (2.32).

(2.32)

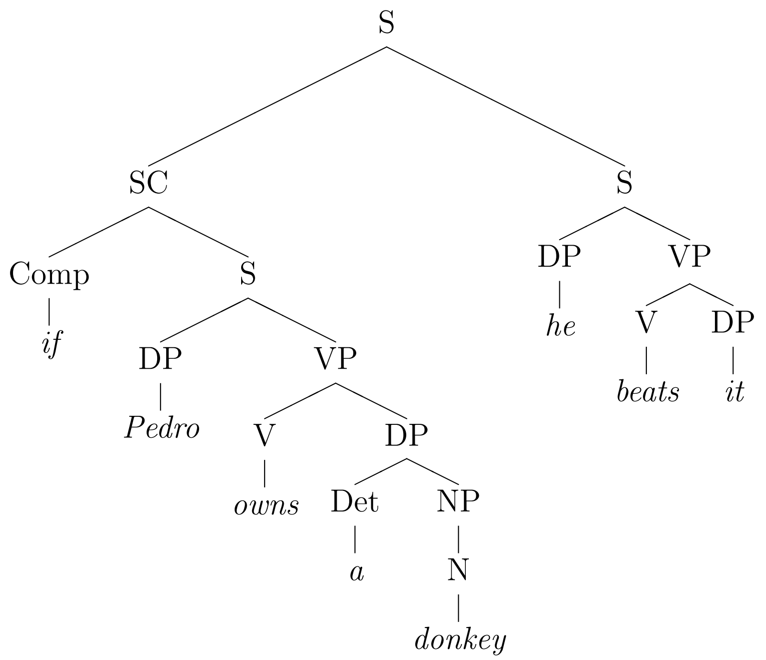

With this form of syntactic structure – of an *if*-clause adjoined to a main clause – comes a DRS construction rule (triggered by the combination of the S-adjunction of the Subordinate Clause and the specification of its Comp as *if*) which separates the *if*-clause from the main clause and puts the two in distinct DRSs K_{ant} and K_{con} which are connected by the conditional connective $\Rightarrow$. (This is a new type of complex DRS condition, for which we will still have to add a verification clause to the general verification definition.) The highest projection level of the SC has thereby done its duty and can be discarded. So what ends up in K_{ant} is just the S-structure of the *if*-clause. Application to (2.32) yields the result in (2.33).

(2.33)

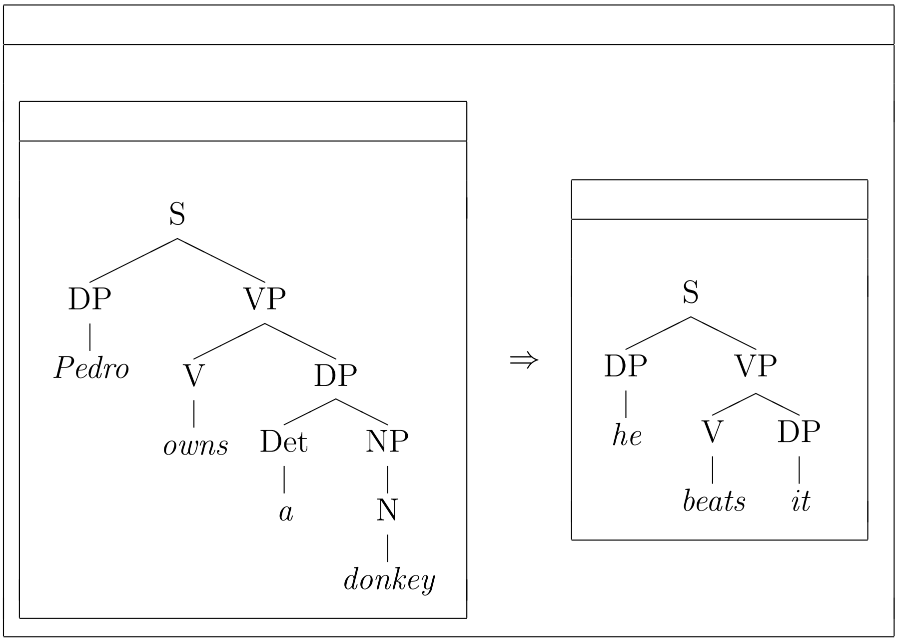

Two reduction steps are required for each of the conditions in K_{ant} and K_{con}. Those needed for the condition in K_{ant} involve (i) the construction rule for the Proper Name *Pedro* and (ii) an application of one of the rules applicable to indefinite DPs. The former rule leads to insertion of the dref chosen for the application – let this dref be p – and the condition specifying that this dref stands for the relevant bearer of the name *Pedro* into the main Universe and Condition Set. The choice of rule for dealing with the indefinite involves, as we have seen, one or two decisions: (i) whether to treat the indefinite as quantifying DP or as term, and, if the choice is made to treat it as a term, then (ii) whether to interpret as a specific or a non-specific indefinite. Strictly speaking the first choice has already been made in the case before us when (2.33) was adopted as LF from which to compute the semantics. For in (2.33) *a donkey* has not been QR-ed, which indicates that it is being treated as a term. (The reasons why a treatment as existential quantifier doesn't work here are the same as those we went through in connection with (1.44.c). In discussing (1.44.c) we also saw that the interpretation we wanted required that *a donkey* be interpreted as non-specific and these considerations apply here as well.)

Treating the name *Pedro* as usual and *a donkey* as non-specific leads to the structure in (2.34).

(2.34)

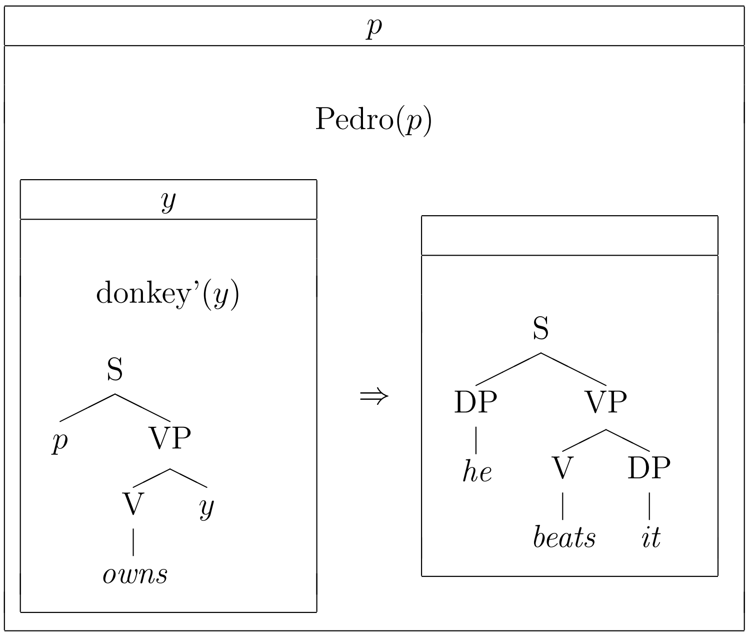

What remains are the reduction operations that still need to be performed on the condition in K_{con}. Both of these involve pronouns, and it is intuitively obvious how they should be interpreted: both pronouns should be interpreted anaphorically, *he* as anaphoric to *Pedro* and *it* as anaphoric to *a donkey*. (2.34) enables us to interpret the pronouns in this way because from the position of the Condition Set of K_{con} both the dref p and the dref y are accessible. Once again, this does not follow from anything we have so far said, but is part of the properties of conditional DRS conditions (those in which two DRSs are connected by $\Rightarrow$). For now let us just assume that this is so and complete the DRS construction accordingly. (2.35) gives the final result.

(2.35)

<table>
<tr><td colspan="3" align="center">p</td></tr>
<tr><td colspan="3" align="center">Pedro(p)</td></tr>
<tr>
<td><table><tr><td align="center">y</td></tr><tr><td align="center">donkey'(y)
owns'(p,y)</td></tr></table></td>
<td align="center">$\Rightarrow$</td>
<td><table><tr><td align="center">$u \quad v$</td></tr><tr><td align="center">$u = p \quad v = y$
beats'(u,v)</td></tr></table></td>
</tr>
</table>

<u>Exercise</u> Show that none of the other options for treating the indefinite *a donkey* in (1.44.a) leads to a DRS that captures the truth conditions we want (and that are represented (2.34)).

The truth conditions of natural language conditionals are as hot a topic of debate today as they ever were and that is as true for conditionals expressed with the help of *if*-clauses as it is for any other form that conditionals can take (in English or other languages). In choosing truth conditions for its conditional DRS conditions DRT took the lead from the classical tradition in formal logic in adopting what are in essence the truth conditions of the material conditional. But the verification conditions for conditional DRS conditions are nevertheless somewhat more involved, because care has to be exercised in relation to the potential members of the Universes of K_{ant} and K_{con}. In fact, this requirement makes the verification condition clause for conditional conditions look remarkably similar to that for universal duplex conditions. To wit:

(2.36) Let M be a model, and f an embedding function into M. Let C be a duplex condition of the form $K_{ant} \Rightarrow K_{con}$
Then:

f *verifies* C *in* M iff every extension g of f to the drefs in $U_{K_{ant}}$ which verifies in M all the conditions in the Condition Set of K_{ant} can be extended to an embedding h which is also defined for the drefs in $U_{K_{con}}$ and which verifies in M all the conditions of the Condition Set of K_{con}.

That this is in essence the material conditional can be seen by considering the special case in which both $U_{K_{ant}}$ and $U_{K_{con}}$ are empty. In this case the extensions g and h that are spoken of in the verification condition above coincide with f, so that the verification condition reduces to:

f *verifies* C *in* M iff (if f verifies in M all the conditions in the Condition Set of K_{ant}, then f verifies in M all the conditions of the Condition Set of K_{con});

and this is equivalent to:

f *verifies* C *in* M iff either f does not verify the conjunction of all the conditions from the Condition Set of K_{ant} or f does verify the conjunction of all the conditions from the Condition Set of K_{con} in M.

Let C_{ant} be the conjunction of the conditions in the Condition Set of K_{ant} and C_{con} be the conjunction of the conditions in the Condition Set of K_{con}. Then the last equivalence can be restated as:

C is verified by f in M iff the material conditional '$C_{ant} \to C_{con}$' is true in M (under the assignment that f may provide for 'unbound' drefs in this formula).

The truth conditions that (2.36) assigns to conditional DRS conditions go some way towards an intuitive justification for the accessibility rules of which we made use when completing the construction of DRS (2.35). But the following informal description of what these DRS conditions try to capture may be even more helpful. Intuitively, the purport of a conditional DRS condition $K_{ant} \Rightarrow K_{con}$ is this: The DRS K_{ant} describes a certain type of situation and the condition as a whole says that any situation that fits this description also fits the description that we get when we extend the description provided by K_{ant} with the descriptive material contained in K_{con}. In view of this understanding of what a conditional is trying to say, it is natural to expect that something like the following must be true: when interpreting what the consequent of a natural language conditional is trying to say the interpreter should be in a position, and should be entitled, to make use of the interpretation that he has already obtained for the antecedent when he interprets the consequent. But if that is how the interpretation of conditional works, then the drefs that have been introduced in the construction of the interpretation of the antecedent should be available when interpreting the consequent.[10]

[10]The 'classical' DRT that we are presenting in this brief survey adopts a classical model theory for its DRS-languages, of the sort described in the introduction to these Notes

This justifies the availability of y for the pronoun *it* in (1.44.a). But it doesn't justify the availability of p for the interpretation of *he*. That p, which is part of the main Universe, should also be accessible has to do with the fact that conditional DRS conditions of the form $\mathrm{K}_{ant} \Rightarrow \mathrm{K}_{con}$ are treated semantically as *indicative conditionals*: they make statements about the actual world and the actual world only. In model-theoretic terms this means that the conditional is evaluated in the same model as other material in the DRS that contains it. In particular, entities that the DRS asserts are actual by virtue of having drefs representing them in its main Universe will be available as anaphoric anchors. That is, the drefs in the main Universe should be available as antecedents for pronoun resolution – both of pronouns occurring in the consequent of the Natural Language conditional and and pronouns occurring in its antecedent.[11]

---

and assumed in what we have been saying about clauses of the verification definition as we went along. But other ways of formulating the semantics of DRS-languages are in principle possible as well, and situation-based accounts are among these. In such a situation-based semantics for DRT DRSs are evaluated with respect to situations. The verification definition is by necessity a partial one, with a given situation verifying some DRSs, falsifying certain others but as a rule leaving the question indeterminate for manny others; sun partiality is a generally acknowledged feature of the Situation Semantics approach. What speaks in favor of a situation-based semantics for DRT is the view according to which DRSs can be seen as describing situations – that they can be seen as complete descriptions of situations rather than as partial descriptions of possible worlds. (This view applies straightforwardly to simple DRSs – recall the discussion on p. 59 ff –but it can be extended to complex DRSs as well, using the existing proposal within Situation Semantics to handle conditionals, quantification and other logical operators.)

In a situation-based semantics for DRT it ought to be possible to interpret conditional DRS-conditions literally as statements about situations. For instance the conditional condition of (2.36) should now be interpretable as saying that any situation described by the antecedent DRS K_{ant} – thus, a situation consisting of a donkey that is owned by the (antecedently given) man Pedro – can be extended to a situation which contains the additional information that Pedro beats this donkey. In fact, it is this very idea that is found in Situation-theoretic accounts of the donkey problems, but without adopting the logical form-based dynamics of DRT, in which DRSs, DRT's logical forms, serve as discourse contexts that can be expanded with new incoming information. For the most explicit elaboration of the sitiuation-theoretic approach to donkey pronouns see the work of Elbourne, in particular his (Elbourne 2005).

[11]Note in this connection that this is not true for counterfactual conditionals. Suppose we had a type of conditional DRS condition that represented counterfactuals. The antecedent DRS of such a condition would be understood as the description of *non-actual* situations, with the counterfactual conditional as a whole saying that if there had been such situations then they would also have satisfied the description provided by the consequent. In the non-actual situations that such conditions speak of the entities represented by drefs in the main DRDs need not exist. Therefore it cannot be assumed without further argument that these will be available as antecedents for pronouns occurring in the

These considerations provide an informal motivation for the accessibility assumptions that we have been making use of in the examples above. But what we also want is a formal definition of accessibility. For that the reader is referred once more to (Kamp & Reyle 1993).

As is easily seen by applying the verification conditions in (2.36) to the $\Rightarrow$-condition in (2.35), the drefs in the Universe of the antecedent DRS of the condition act as if they were universally bound variables.

Exercise Show that the DRS in (2.35) has the same truth conditions as the one in (2.38) below, in which the $\Rightarrow$-condition has been replaced by a universal duplex condition in which the dref y is bound by the universal quantifier of that condition.

---

descriptions of such counterfactual situations. More specifically, this problem will arise when an individual that the DRS represents as part of actuality is not mentioned (directly or indirectly) in the antecedent of a counterfactual, and yet an attempt is made to resume it via a pronoun in the consequent. Here is an example:

(2.37) Pedro has bought a goat. If he had bought no goat, it would have been allowed inside the house.

This little discourse is bizarre because the worlds or situations that the antecedent of the counterfactual talks about the purchase that is said to have happened in the real world didn't occur, so there is no goat there that Pedro bought. So, in formally speaking, there is no goat for *it* to refer to in these worlds. (I am assuming that the use of *a goat* in the first sentence of (2.37) is non-specific. If the phrase were used specifically, it should be possible to construe the *it* of the second sentence as anaphoric to the particular goat that the specific indefinite was speaking of, so that the conditional as a whole would amount to the statement that if Pedro hadn't bought that goat, then he wouldn't have allowed it (i.e. that goat) inside the house (but, perhaps, given that the actually did buy the goat he does let her inside).)

As a first step in the direction of accounting for intuitions like the one I have just been trying to articulate, discourse referents in the Universe of the main DRS should not be automatically accessible for the interpretation of pronouns occurring in countercatual conditionals. (But of course this can't be more than a first step.)

(2.38)

$$\begin{array}{|c|}
\hline
p \\
\hline
\\
\text{Pedro}(p) \\
\begin{array}{|c|} \hline y \\ \\ \text{donkey'}(y) \\ \text{owns'}(p,y) \\ \hline \end{array}
\;\left\langle \begin{array}{c} \forall \\ y \end{array} \right\rangle\;
\begin{array}{|c|} \hline u \quad v \\ \\ u = p \quad v = y \\ \text{beats'}(u,v) \\ \hline \end{array} \\
\hline
\end{array}$$

After all this there isn't much that needs to be said about (1.44.b):

(1.44.b) If a farmer owns a donkey he beats it.

The only difference with (1.44.a) is that the subject of the *if*-clause is now an indefinite as well, just like its direct object. This means that the Universe of the DRS for the *if*-clause now has discourse referents for both the donkey and the farmer, and thus, according to the accessibility principles already adopted, that both of these are available as anaphoric antecedents for pronouns in the main clause. So the dref introduced by *a donkey* can be picked up, as in (1.44.a), by *it* and the one introduced by *a farmer* by *he*. (Again, how the choices are made is an aspect of anaphora that is not addressed in the versions of DRT that we are discussed and used in these Notes.) The resulting DRS is shown in (2.39).

(2.39)

$$\begin{array}{|c|}
\hline
\\
\hline
\\
\begin{array}{|c|} \hline x \quad y \\ \hline \\ \text{famer'}(x) \quad \text{donkey'}(y) \\ \text{owns'}(p,y) \\ \hline \end{array}
\;\Rightarrow\;
\begin{array}{|c|} \hline u \quad v \\ \hline \\ u = x \quad v = y \\ \text{beats'}(u,v) \\ \hline \end{array} \\
\hline
\end{array}$$

It is easily seen that the verification conditions for DRSs and DRS-conditions assign to (2.39) the truth conditions of a doubly universalized conditional, with one universal quantifier binding x and a second one binding y. If we look more closely at how these truth conditions emerge, it is tempting to think that what is involved here is a single universal quantifier which binds x and y all at once. (This is the universal quantifier over extensions g of the verifying embedding f that occurs in the verification conditions for $\Rightarrow$-conditions; in the present case the domains of the gs extend the domain of f with the two drefs x and y.) Seen in this light, the verification conditions for $\Rightarrow$-conditions capture the notion of *non-selective binding* that was first introduced in (Lewis 1975) in his analysis of quantifiers in natural language on the model of quantificational adverbs, like *always*, *mostly*, *rarely* and so on. There has been a fair amount of discussion in the wake of DRT and the (for these purposes equivalent) File Change Semantics of Heim over whether this is the empirically correct treatment. The general upshot of that discussion is that it is right for adverbial quantifiers but not for the corresponding nominal quantifiers. The reason why non-selective binding is not the correct analysis for nominal quantifiers cannot be detected in any straightforward way when we compare the quantificational adverb *always* with the nominal quantifiers *every* and *all*. But it emerges clearly in connection with a non-universal quantifier like 'most'. This quantifier can be expressed either nominally, with the help of the determiner *most*, or adverbially, by means of the adverb *mostly*. But the semantics of *most* and that of *mostly* are not fully equivalent. This is shown by well-worn examples involving the following sort of scenario: There are 25 farmers who own one donkey, 25 farmers who now two donkeys and one very rich farmer who owns 200 donkeys. Suppose that the first 50 farmers beat their donkeys, but that the last farmer (who has bigger fish to fry) beats none of the donkeys that he owns. Then the sentence in (2.40.a) is intuitively true: The sentence seems to make a statement about donkey-owning farmers, and to say that a majority of them are donkey beaters.

(2.40)a. Most farmers who own a donkey beat it.

b. Mostly, if a farmer owns a donkey, he beats it.

That is surely the case in the scenario just described, where there are fifty donkey beaters and only one non-beater. However, that is not what the non-selective binding account suggests. If 'most' binds pairs of variables representing farmers and donkeys they own, then that would suggest that for the sentence to be true there has to be a majority of such farmer-donkey pairs with the property that the farmer beats the donkey. But that condition is not satisfied in our scenario; for here there are 200 farmer-donkey pairs for which the farmer (the very rich one) does not beat the donkey and only 75 pairs for which the farmer does.

So, if we take it to be an essential part of the non-selective binding account that the kind of counting that is involved in the evaluation of quantifiers like 'most' is over the pairs (or, more generally, tuples) of individuals assigned to the non-selectively bound variables of the restrictor, then non-selective binding does not provide the right analysis for the nominal quantifier *most*. It seems intuitively right to extend the conclusion of this argument to other nominal quantifiers, including *every* and *all*, even if for these determiners the argument above (involving our scenario with the 51 farmers) doesn't work. (Exercise Why can't such arguments work for terms expressing universal quantification?)

In DRT terms, what is needed to deal with such nominal quantifiers are duplex conditions that involve direct binding of just one dref, but may involve 'secondary', or 'indirect' binding of the other drefs in the Universe of the restrictor DRS. (Thus in *Every farmer who owns a donkey beats it* the dref introduced by *every farmer* is directly bound and the dref introduced by *a donkey* indirectly.) When the distinction between directly and indirectly bound drefs is made explicit in the logical form representing the quantification, it then becomes possible to exploit this difference in a formulation of verification conditions that gets the truth conditions of (2.40.a) right. (There are some further complications however. For details see the discussion of 'most' in (Kamp & Reyle 1993).)

For (2.40.b) the facts appear to be different. Here the 'farmer-donkey' pair reading seems a possibility: You can, many have claimed, understand this sentence in such a way that it comes out as false in the above scenario. And that is because in this case it is possible to understand the quantifier as counting pairs. In other words, the adverbial quantifier *mostly* does seem capable of binding non-selectively, just as (Lewis 1975) had it. One way to represent this within out framework is to allow for duplex conditions in which the quantifier in the diamond in the center is not provided with any dref as argument. The interpretation of such a duplex condition is then – simplifying slightly – that what the quantifier counts are tuples of elements in the model that correspond to the drefs in the Universe of the restrictor DRS. Thus in (2.41), where the restrictor Universe has two drefs, the counting is over pairs of element from the model.

(2.41)

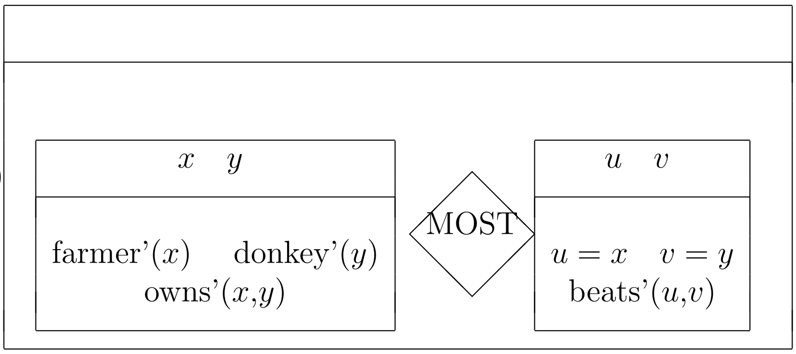

<u>Exercise</u> Give a formal statement of the verification conditions for duplex conditions of the kind illustrated in (2.41) (that is, duplex conditions in which the quantifier in the central diamond is not accompanied by a dref that it binds) which is in accordance with the informal explanation given above, and show that when this statement is applied to (2.41) the truth conditions are that the majority of the farmer-donkey paris such that the farmer owns the donkey have the property that the farmer beats the donkey.

The judgments of sentences involving adverbial quantification are subtle, however, and seem to depend on additional factors, in particular on aspects of *information structure*. Consider the following pair of sentences (free after Heim(?)).

(2.42)a. If a DRUMMER lives in an apartment complex, it is usually half empty.

b. If a drummer lives in an APARTMENT COMPLEX, he usually gets on with his neighbors.

(Capitals indicate that the constituent is given focal stress.)

Suppose the scenario is this. The dynamics of rental accommodation being what it is, drummers tend to have difficulties in finding places for rent, so that they often end up in places where few people really want to live and which, accordingly, have a good many vacancies. But then once a drummer moves into such a place, the non-drummers will gradually move away and be replaced by other drummers. So among the apartment complexes that have any drummers at all most will have a predominance of drummers. But let us assume that the filling up with drummers is a slow process, which doesn't keep pace with the exodus of non-drummers, so that apartment complexes with drummers also tend to be half-empty for quite a while, before they finally fill up again, with drummers filling up the vacancies. So most of the apartment complexes with drummers in them are half-empty, but if we count pairs of drummers and apartment complexes in which they live, then the majority of those do not involve apartment complexes that are half-empty, for the simple reason that there are so many drummers who live together in apartment complexes that have filled up with just them. Under these conditions (2.42.a) seems true, because intuitively it speaks of apartment complexes with drummers – this is the kind of apartment complex that is being counted by the quantifier *usually* in (2.42.a) – even though half-emptiness is not typical for the apartment complexes that occur in drummer-apartment complex pairs.

The explanation for our intuition that (2.42.a) talks about apartment complexes with one or more drummers living in them, rather than about drummer-apartment complex pairs or about drummers that live in an apartment complex, has to do with the focal stress on *drummer*. The effect of the focus on *drummer* is to make apartment complexes that have a drummer among the people living in them into the topic of the conditional as a whole. (The focus on 'drummer' has the effect of creating a contrast between apartment complexes *with* a drummer (the topic of the conditional) and apartment complexes without drummers, which is what the conditional is not about.) The over-all effect of this is that (2.42.a) comes across as a statement that quantifies over apartment complexes with drummers, much as if the wording had been: 'Most apartment compilexes with a drummer in them are half-empty.'

We see the inverse effect in (2.42.b). Here the stress on *apartment complex* has the effect that the sentence is understood as talking about drummers who are living in apartment complexes (as opposed to those living in other kinds of accommodation).

To show the truth conditional consequences of this by driving a wedge between quantification over such drummers as opposed to quantification over drummer-apartment complex pairs is a little harder. In fact, it is impossible so long as we assume that each drummer is living in just one apartment complex (or more generally, as living in just one place). For then there is a one-to-one correspondence between drummers living in apartment complexes and pairs consisting of a drummer and the apartment complex that he lives in. But we can change our perspective somewhat by taking into account that drummers (perhaps even more so than people in general) move from one accommodation to the next. It may be that normally drummers will, before they eventually reach the haven of an apartment complex with only fellow drummers, have lived in many apartment complexes from which they were expelled after complaints from the non-drummers. In this situation the number of drummer-apartment complex pairs may exceed (by a significant factor) the number of drummers who have ever lived in an apartment complex (i.e. those who have lived in apartment complexes at one time or another). If we assume in addition (and not unreasonably) that a drummer normally gets on well enough with other drummers living in the same complex, but that drummers and non-drummers do for the most part not get on well when they are neighbors, then (2.42.b) comes across as true. For on balance most drummers that live in apartment complexes live in apartment complexes occupied largely or wholly by drummers; so all or most of them have neighbors who are also drummers, and will get on with them. But if we understood the sentence as quantifying over drummer-apartment complex pairs, then the sentence majority would not come out true, because for the majority of such pairs, the drummer lives in a complex where her or his neighbors are likely to be non-drummers, and drummers and, as I surmised, non-drummers do for the most part not get on when they are neighbors.

It appears from this discussion (and not just because a the scenario is a little artificial and involves so many assumptions) that the semantics of adverbial quantifiers, and in particular of those like *mostly* and *usually*, is a complicated matter, and more complicated than the semantics of nominal quantification, where the syntactic structure of the quantifying DP always makes clear what is bound directly and what only indirectly,and thus what is quantified over and counted by the nominal quantifier. If we assume that non-selective binding is part of the semantics of adverbial quantification, then a further account is needed of how focal stress can undo the semantic implications (of quantification over tuples) that are usually taken to be part of non-selective binding.[12]

[12]Note that the proposal made above for the form of duplex conditions introduced by quantifying adverbs won't do for the cases just discussed, as in these the non-selective binding effect is overruled. One way to deal with the interpretations for (2.42.a) and (2.42.b) we have just discussed is to go back to our relier mode of representing quantifying and to introduce the dref that is being quantified over into the quantifier in the middle. (It is just that the considerations that lead to this insertion will now be more complicated than they are in the case of nominal quantifiers.) However, in the light of this correction, the suggestion made above, according to which non-selective binding is represented through an absence of drefs in the quantifier component of the duplex condition doesn't seem very natural any more. It would now seem more in line with the special provisions that have to be made for the specification of what the directly quantified variables are to always insert *all* the drefs that are part of the direct quantification into the diamond. In the case of non-selective binding this would mean that all the drefs occurring in the Universe of the restrictor DRS would also occur in the central diamond.

Developing a formal account of adverbial quantifiers with their sensitivity to focus effects is out of the question here, for one thing because it would have to include a formal account of the semantic properties of focus, topic and other information-theoretic concepts. So we have to leave the discussion of this kind of quantification at the informal and superficial level at which it has been conducted. But, summarizing and concluding, what the discussion has taught us is: (i) that adverbial quantification differs from nominal quantification in that the binding it involves is potentially non-selective; and (ii) that adverbial quantifiers are often hard to interpret because their potential non-selectiveness can be affected by information-theoretic effects like that of focus.

2.5 Negation

The standard repertoire of logical operators of the predicate calculus consists of &, $\vee$, $\neg$, $\rightarrow$, $\exists$, $\forall$ and (sometimes) $\leftrightarrow$. Within DRT two of these – & and $\exists$ – are captured structurally, & by putting conjuncts into the same Condition Set and $\exists$ by placing a dref in the local DRS Universe. The remaining ones each require for their semantic representation a special type of complex DRS condition. We have encountered two of those so far, viz. universal duplex conditions for $\forall$ and $\Rightarrow$-conditions for $\rightarrow$. Of the remaining three we will only have a quick look here at $\neg$. [13]

Negation is among the logical operators that cannot be defined in terms

[13] For $\vee$ see ((Kamp & Reyle 1993)). The biconditional is curiously problematic in DRT. Normally the biconditional is defined as the name 'biconditional' suggests: $A \leftrightarrow B$ is analyzed or defined as $(A \rightarrow B)$ & $(B \rightarrow A)$. In DRT forming such a conjunction doesn't automatically give the desired result, because the left hand DRS of a $\Rightarrow$-condition may contain one or more drefs which act as universally quantified variables with scope over the conditional. Forming the converse $\Rightarrow$-condition of the given condition in order to get the effect of the biconditional won't do in such cases, as the result would be that these drefs are now in the Universe of the consequent box, which makes them into existentially quantified variables with scope restricted to the consequent of the conditional. In order to get the right form for the second conjunct of the conjunction of conditionals it is in general necessary to exchange the Universes of the antecedent DRS and the consequent DRS. I do not know of a useful and perspicuous notation to indicate this exchange. (One could of course define a DRS condition of the form $\mathrm{K}_1 \Leftrightarrow \mathrm{K}_2$ as shorthand for a pair of conditional DRS conditions $\mathrm{K}_1 \Rightarrow \mathrm{K}_2$ and $\mathrm{K}'_2 \Rightarrow \mathrm{K}'_1$, where K'_1 and K'_2 are the results of subjecting K_1 and K_2 to the necessary dref exchange. I haven't so far found the effort worth making, but in DRT applications in which biconditionals are prominent there could well be a point in doing so.)

of conjunction and existential quantification. So it cannot be represented structurally within DRT and therefore needs its own type of complex DRS-condition. Since negation is a 1-place operation, this DRS-condition will be built out of one DRS (and not two, in the manner of ⇒- and duplex-conditions). For a simple example consider the sentence in (2.43).

(2.43)Pedro doesn't own a donkey.

Let us assume that the negation in (2.43) is adjoined to VP and that it is morphologically realized as *doesn't*. this less to the assumption that the syntactic structure of (2.43) is as in (2.44).

(2.44)

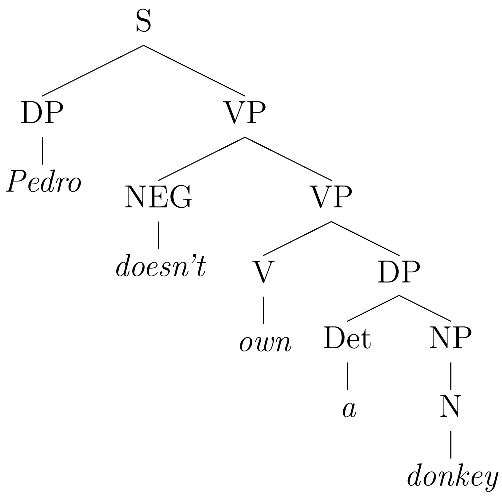

The first step of the top down construction method when applied to (2.44) is familiar and leads to (2.45).

(2.45)

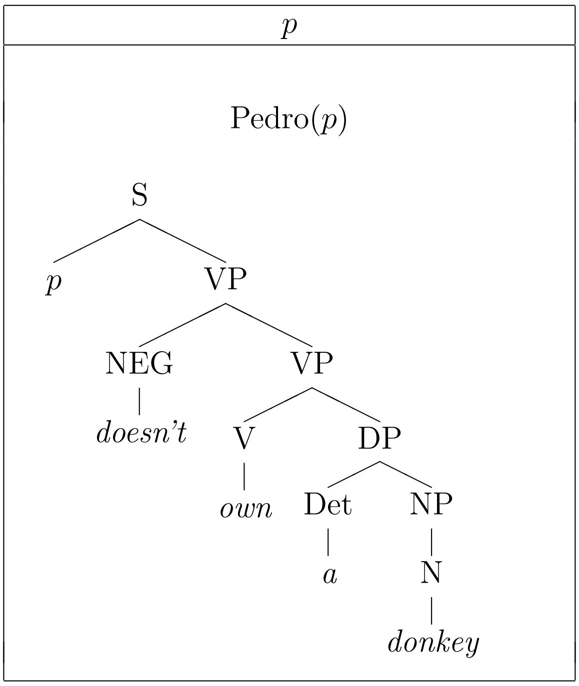

The next step is the one that matters. The effect of the NEG-node is to introduce a DRS-condition which consists of the negation sign ¬ applied to a DRS which serves to represent the content of the part of the syntactic tree to which the negation has been syntactically adjoined. In the present case this leads to (2.46).

(2.46)

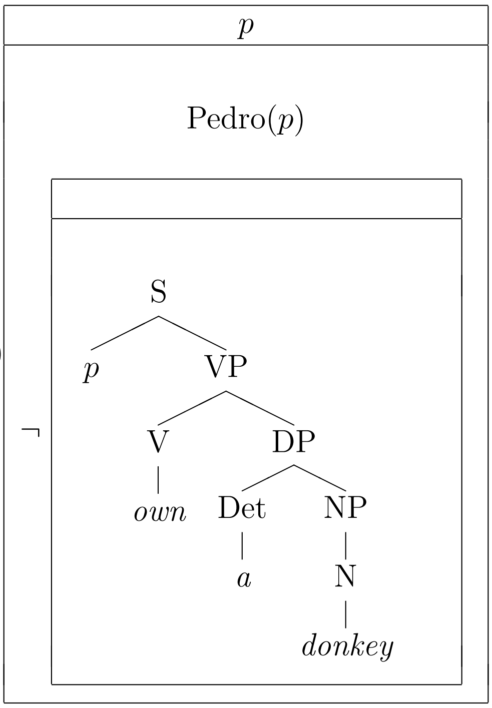

At this point there is only one more operation to be performed, that which deals with the indefinite direct object phrase. Assuming a narrow scope reading for the indefinite as non-specific singular term, which places the new dref into the Universe of the local DRS, we get, after eliminating the last remnants of syntactic structure the DRS in (2.47).

$$(2.47)\quad \begin{array}{|c|}\hline p \\ \hline \\ \text{Pedro}(p) \\ \\ \neg\ \begin{array}{|c|}\hline y \\ \hline \\ \text{donkey'}(y) \\ \\ \text{own'}(p{,}y) \\ \hline \end{array} \\ \\ \hline \end{array}$$

What should be the verification conditions of ¬-conditions? Intuitively the answer is clear: an embedding function f verifies such a condition in a model

M iff it is not possible to extend f in a way that verifies the DRS that is in the immediate scope of ¬. Formally:

(2.48) f *verifies* ¬ K *in* M iff there is no extension g of f to the Universe of K which verifies the conjunction of all the conditions from the Condition Set of K_{con} in M.

<u>Exercise</u> Use the verification clause in (2.48) to spell out the verification conditions for (2.47) and verify in this way that (2.47) captures the intuitive truth conditions of (2.43).

It is part of the general definition of accessibility in DRT that the drefs in the Universe of a sub-DRS that is the DRS of a ¬-condition are not accessible from the level of the DRS which contains this condition (let alone from any DRS which contains this latter DRS as a sub-DRS). This is confirmed by the embedding conditions for ¬-conditions, which treats the drefs in the Universe of the DRS prefaced by ¬ as bound variables. It is also confirmed by examples like that in (2.49).

(2.49) ?? Fred doesn't own a car. It is red

There are other examples, however, that seem to contradict the predictions made by this part of the definition of accessibility. (2.50) has two of them.

(2.50) a. It is not true that there is no rabbi at this wedding. He is standing behind that palm tree over there. (Beaver)

b. If Fred does not have a car then he does not need to park it.

One possible suggestion for dealing with (2.50.a) is that when a DRS containing a dref in a certain inaccessible position can be transformed (by some simple and straightforward logical transformation) into a logically equivalent one in which that dref occupies an accessible position, then this operation may be carried out prior to using the (now accessible) dref as anaphoric antecedent.

As regards (2.50.b) there has been a suggestion that the negated content in the consequent can be understood as extending the negated content of the antecedent of the conditional. (Thus, if the structure of the DRS of this sentence is as in (2.51), then interpretation of the consequent may take the form of duplicating the negated condition in the left hand side box and then processing the non-negated main clause of the conditional as part of the DRS of the duplicated ¬-condition.)

(2.51)

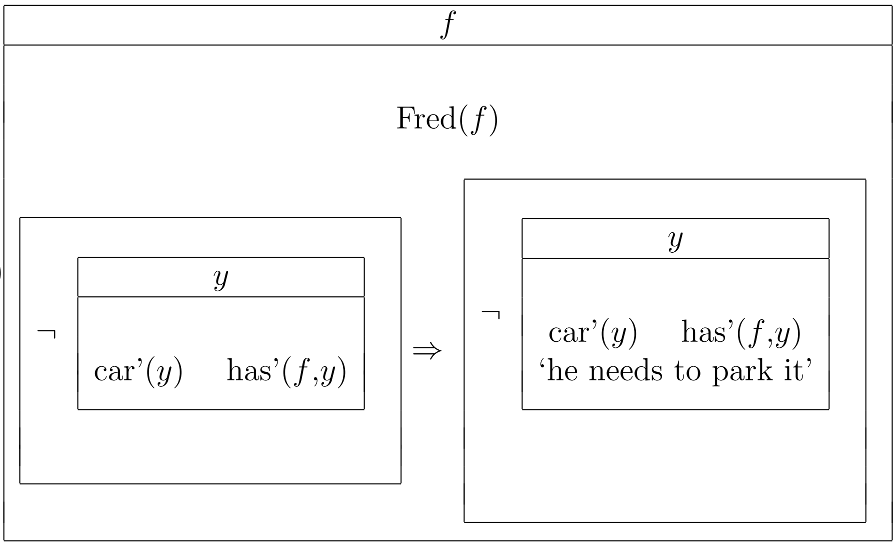

Exercise

Either (i) give reasons why you think these solutions are not viable or (ii) fill in the necessary details.

This brings us to the conclusion of this partial review of the original top-down approach to DRS construction. The choice of negation as the operator with which to conclude this review has two reasons. First, negation will play a significant role in later parts of these notes. The second reason is independent of the purpose and continuation of these notes. Within DRS languages negation has a special status in that ¬-conditions are the only complex DRS-conditions that need to be added to the vocabulary of basic DRSs in order to obtain a DRS-language with the expressive power of full first order logic. The intuitive reason for this is that in classical logic &, ∃ and ¬ form a 'functionally complete base' for the traditional set of operators mentioned above – in the sense that all remaining operators can be expressed with the help of these three – and that in DRT & and ∃ are represented structurally, so that no special DRS-conditions are needed on their behalf. (For details see ((Kamp & Reyle 1993)).)

An implication of this observation is that the contributions made by universal duplex conditions, ⇒-conditions (and the ∨-conditions that are used in DRT to represent disjunctions, but which we have left out of this review) can be mimicked by constructs involving just ¬-conditions. For Logical Form purposes – i.e. for the representation of universal quantification, conditionals and disjunctions as they occur in English and other natural languages – such reductions are awkward, in that they make the relation between the natural language sentences and their Logical Forms much less perspicuous. But when DRS-languages are used for other purposes than providing DRSs as Logical Forms for bits of natural language – e.g. when they are used as general Knowledge Representation Languages – then the reduction to the stripped down version in which ¬-conditions are the only complex DRS-conditions may have its uses, e.g. in formal proofs that make use of induction over DRS complexity.

2.6 Plural Pronouns and other Plural DPs

This section gives a selective review of a treatment of plurals in DRT that is presented in considerable detail in *From Discourse to Logic* (see (Kamp & Reyle 1993), Ch. 4). The purpose of this review is two-fold: (i) to point out that plural pronouns are, on the face of it, less restrictive than singular pronouns as regards antecedent selection; and (ii) to introduce the *mereological* view of the relation between the referents of singular and plural phrases, of which we will make use later on in our discussion of singular and plural definite and indefinite descriptions.

While this section is not about the top-down method as such, the few examples of DRS construction involving plurals that I will give will make use of this method, since that is all we have available at this point. So in that sense the section can be seen as a sort of appendix to the previous one. But on the other hand its main purpose is to do more of the groundwork for what is to come; and this groundwork is of a quite different sort from what we have just gone through.

2.6.1 Partee's Ball Examples

One of the central principles behind the DRT-based treatment of anaphoric singular pronouns that was presented in the last section was that their interpretation always requires an 'antecedent' dref that is already part of the representation at the point when an interpretation for the pronoun has to be found (and moreover this dref must be in a position that is 'accessible' from the position occupied by the pronoun). Perhaps the most dramatic illustrations of this constraint are minimal pairs like the following, originally due to Partee:

(2.52)a. One of the ten balls is not in the bag. It is under the sofa.

b. Nine of the ten balls are in the bag. It is under the sofa.

c. All but one of the ten balls are in the bag. It is under the sofa.

In (2.52.a) the pronoun *it* in the second sentence can be interpreted as anaphoric to the missing ball mentioned in the first sentence. In (2.52.b) this is not possible, although the first sentence seems to be conveying the same information as the first sentence of (2.52.a) of some set of ten balls, nine are in the bag and one is not. (If *it* has any interpretation in this sentence, it is as anaphoric to *the bag*.) The explanation of this difference suggested by the version of DRT we have just reviewed is quite straightforward: The subject phrase of the first sentence of (2.52.a) introduces two drefs, one for the missing ball and one for the set of ten balls of which this ball is a member. The dref for the one missing ball can then serve as antecedent for *it* in the second sentence. In contrast, the subject DP of the first sentence of (2.52.b) introduces besides a dref for the set of ten balls a dref for some nine-membered subset of that set, but not for the one ball that makes the difference between them. Thus no such dref representing a single ball is available when *it* needs an interpretation, and the best the interoperation can come up with is to link the pronoun to the dref representing the bag. (2.52.c), which is in some sense even more similar to (2.52.b) than (2.52.a) is, confirms the pattern: Here, once again, processing of the subject DP of the first sentence leads to the introduction of two drefs, one for the set of ten balls and one for the missing ball; and the latter is then available for the interpretation of the pronoun.

To make this explanation formally explicit we need a few more principles about DRS construction than we have been considering so far. Some of these will have to be stopgap measures at this point; they will be replaced by more fully developed treatments later on, and for now all I can do is plead with you to find them plausible.

Our first stopgap measure concerns the definite descriptions that occur in the first sentences of (2.52.a) and (2.52.b). They each contain two definite descriptions, *the ten balls* and *the bag*. Eventually these will be treated as triggers of 'reference identification' presuppositions, just as the other two types of definite DPs we have already encountered, the pronouns and the proper names (though, as noted, each of these DP-types generates its own type of presupposition, with its own form and own resolution constraints). Our temporary stopgap assumption for definite descriptions is to simply let them introduce a dref, which is characterized as the maximal satisfier of the

DP's descriptive content,with a proviso for additional tacit constraints, which are not expressed overtly in the DP itself, but are supposed to be recoverable from the context in which it is used. We will assume that all these tacit constraints are packed into a single predicate C, but, for now, say nothing more about what C might be, or how it may be recovered. For the plural definite description *the ten balls* this amounts to the following: the DP introduces a dref X – we use a capital letter in such cases to indicate that what the dref represents is a set of two or more individuals and not a single individual; this is the semantic contribution made by the plural morphology of the DP – together with conditions that express: (i) that X contains all individuals that are balls and satisfy the additional constraint C, and (ii) that X has ten elements, or that the 'cardinality of X is 10', in more official set-theoretical jargon. There is more than one way in which maximality can be expressed. The one that proves most conducive to our purposes is that which says that the referent consists of satisfiers of the descriptive content of the description (more often than not in conjunction with some contextually given predicate C) and that there is no proper superset that also satisfies this condition. One way to express this – here for the case of *the (ten) balls* – in our DRS formalism makes use of the set-theoretical relations 'element' and 'proper superset', for which we use the standard symbols $\in$ and $\supset$. To express the second condition we use the familiar device consisting of two vertical strokes to denote the cardinality of the set denoted by the term they flank. But we will see in the next two paragraphs that the use of this notation is not essential.

The maximality condition, thus conceived, is given in (2.53.a) and the cardinality condition in (2.53.b).

(2.53)a.

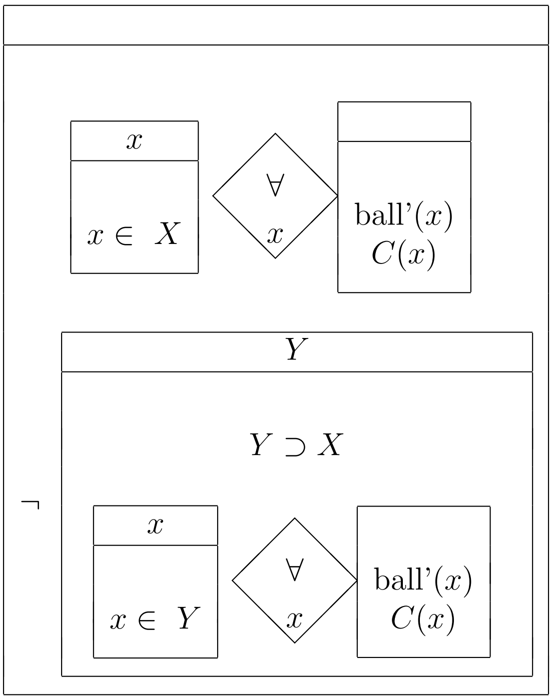

b. $|X| = 10$

N.B. The condition in (2.53.b) may give the impression that the DRS-language to which it belongs is one that includes number theory. But that is not the case. Strictly speaking the DRS-language to which '$|X| = 10$' and similar cardinality conditions belong does not contain constant terms, such as '10', for natural numbers. Rather, the condition '$|X| = 10$' is to be understood as an abbreviation of the formula which expresses that X has 10 members in logical notation – that is: in the way in which the formula '$(\exists u)(\exists v)(u \in X \mathbin{\&} v \in X| \mathbin{\&} (\forall w)(w \in X \rightarrow w = u \vee w = v))$' expresses that X has exactly two members. More precisely, '$|X| = 10$' is to be understood as short for a set of simple and complex DRS-conditions that are equivalent to the formula of predicate logic that expresses that X has ten members. Note well: that the full specification of a DRS construction algorithm that can deal with DPs of this general form (i.e. *the* + cardinal term + NP) presupposes that the construction algorithm includes a special module that deals explicitly with the conversion of number terms used in the chosen fragment of English. This module must be able to convert phrases containing number denoting expressions - such as 'two', '2', 'ten', '10', 'twelve', '12', 'twenty one', '21', 'one hundred', '100', '10^2' and so on into sets of DRS conditions in which no such references occur. In particular, it must be able to convert DPs in which a number denoting expression occurs as the modifier of a noun, as in *the ten balls*, *the last hundred days*' etc. into sets of DRS conditions which jointly say of the plural dref representing the DP that it consists of the number of elements in the set denoted by the expression. Writing out such a module in full detail is a painful and not particularly rewarding exercise, in part because of the idiosyncrasies that are found in the naming systems for natural numbers that we find in most natural languages. (English is a comparatively mild case of idiosyncrasy in this domain.)

However, as the formula above for the case of '2' indicates, the DRS-language needed to represent sentences like those in (2.52) must have an element that was not part of the DRS-languages considered in the DRT review presented in the last section. This is the relation constant $\in$. The relation expressed by this constant – that of an individual being a member of a set (or, as will be proposed below, an atomic part of a mereological complex) – is essential for the expression of cardinalities in terms of classical predicate logic, and it is al;so needed to deal with the semantics of *partitive constructions* of the sort we find in a DP like *one of the ten balls*, in which the word *of* is used to express that the individual denoted by *one* is a part of the set denoted by *ten balls*.

2.6.2 Partitives and Cardinals

This brings us to the syntax and semantics of the partitive DPs *one of the ten balls* and *nine of the ten balls*. From an intuitive semantic point of view both of these are indefinites: they introduce a new element into the discourse. In the first case this is a individual – some ball – and in the second it is a set – some set of nine balls. But apart from this difference between individual and set the two phrases look like they are very much alike, and it should be our 'null hypothesis' that they have the same structure.

In deciding what this structure is it is somewhat easier to start with the second phrase. The word *nine* seem to have, like all 'plural cardinals' (viz the words *two*, *three*,.., *eleven*, ..) two uses: (i) as proper names (of the number in question; this is a use typically found in mathematics, e.g. in *five plus seven equals twelve*); and (ii) as prenominal modifiers, as in *two pounds of sugar*, or *ten balls*. When used in this second way cardinals behave much like adjectives. But they nevertheless have a special status. This is clear when we compare, say, the phrase *ten balls* with one in which *ten* has been replaced by a regular adjective, such as *red* or *big*. *red* in *red balls* says something about each of the balls which belong to the set of balls that is denoted by the phrase as a whole. *ten* doesn't do that. It doesn't say anything about the individual balls in the set but makes a statement about the size of the set. That is particularly clear when prenominal cardinals and adjectives are combined in a single phrase, as in *ten red balls*; and it is also significant that in such phrases the order of cardinal and adjective cannot be (easily) reversed: *red ten balls*, if perhaps not outright ungrammatical, is very marked.

We will therefore treat prenominally used cardinals as belonging to their own syntactic category, 'CardP' (for "Cardinality Phrase'). And we will assume, consistently with the fact that prenominal cardinals typically precede prenominal adjectives, that CardPs enter the structure of the noun phrase at a higher level than prenominal adjectives. This level is one at which it has been decided whether the phrase as a whole is to be used to describe a single individual or a set of two or more individuals. In English that decision is typically made manifest through the difference between singular and plural morphology. (We will see below that that assessment isn't quite accurate in general, but for now the generalization will do and be useful.) One way in which we can make this idea formally precise is to assume that CardPs are the specifiers of a new projection level – that of 'NumP' (for 'Number Phrase') – and that the head of that projection (labeled 'Num', for 'Number') carries the feature that decides between plural and singular. (This is a binary feature with the two values 'sing' (for 'singular') and 'plur' (for 'plural') When the specifier position is filled with a 'plural cardinal' (i.e. a cardinal for any of the numbers from 2 upwards), then its input, given by the sister node Num', must of course be a phrase that denotes a set of two or more members. So the feature imposed by Num must be plur.[14][15]

[14]One way to secure this is by treating the plural cardinals as carrying a selection restriction to the effect that their inputs must be 'plural' in this sense. We will return to this later, when we have said enough about presupposition. I am assuming that selection restrictions are a species of presupposition.

[15]Prenominal cardinal phrases can have considerable complexity. Here are some examples: *at least one*, *more than two*, *one or two*, *exactly three*, *less than or equal to seventeen*, *between five and ten*, *either at most five or at least eight*, *at least five thousand and if the opposition is right more than ten thousand*. I do not know if there is any definitive statement of what the exact range of English cardinality phrases is. (A systematic corpus-based search would help. Has it been done?), but these few examples seem to show clearly enough that there is non-trivial complexity here and that a characterization of the totality of such phrases requires a recursive definition of the sort familiar from generative syntax. This is not the task before us right now. But the sheer extent and the potential complexity of number phrases seems an additional reason for classifying them as a syntactic category in its own right, to which the grammar assigns a suitable slot in the syntactic structure of the noun phrase.

What has been said so far implies that the syntactic structure of the DP *the ten balls* should be as in (2.54).

(2.54)

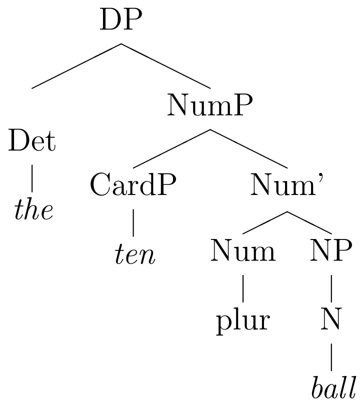

But what about *nine* in *nine of the ten balls*? Superficially the phrase as a whole looks like the result of adjoining the PP *of the ten balls* to an NP. And this is how we will analyze it. This means that the NP to which *of the ten balls* is adjoined is given, somehow, by *nine*. But what syntactic role does *nine* play in this NP. The answer to this question becomes clearer when we look at another use of plural cardinals, illustrated by (2.55).

(2.55)Mary bought eight books and Fred bought nine.

In this sentence *nine* is clearly functioning as an argument phrase, just as *eight books*, and thus as a DP. It is also clear what this DP refers to: it refers to nine books; 'book' has to be recovered from the context (which is easy enough in this case) and after this recovery *nine* can be seen as playing the same part in the reconstructed phrase as *ten* does in *ten balls*. The one difference between the occurrence of *nine* in *nine of the ten balls* and *nine* in (2.55) is that in the former there is no need for contextual recovery of the nominal head (the constituent of category N). The empty N gets its content from the partitive PP that is adjoined to it at the level of NP. (But this has to do with the semantics rather than the syntax of the phrase.)

This gets us most, but not all of the way. *nine of the ten balls*, as it occurs in the first sentence of (2.52.a), is a DP, not a NumP. But the missing step is straightforward. Plural indefinites can have an empty determiner, as we see in simple sentences like 'Fred bought books.' or 'Books were strewn across the table.'. It is reasonable to assume that this is also the case for the subject phrase *nine of the ten balls*. With this assumption we get as syntactic structure for *nine of the ten balls* the tree shown in (2.56).

(2.56)

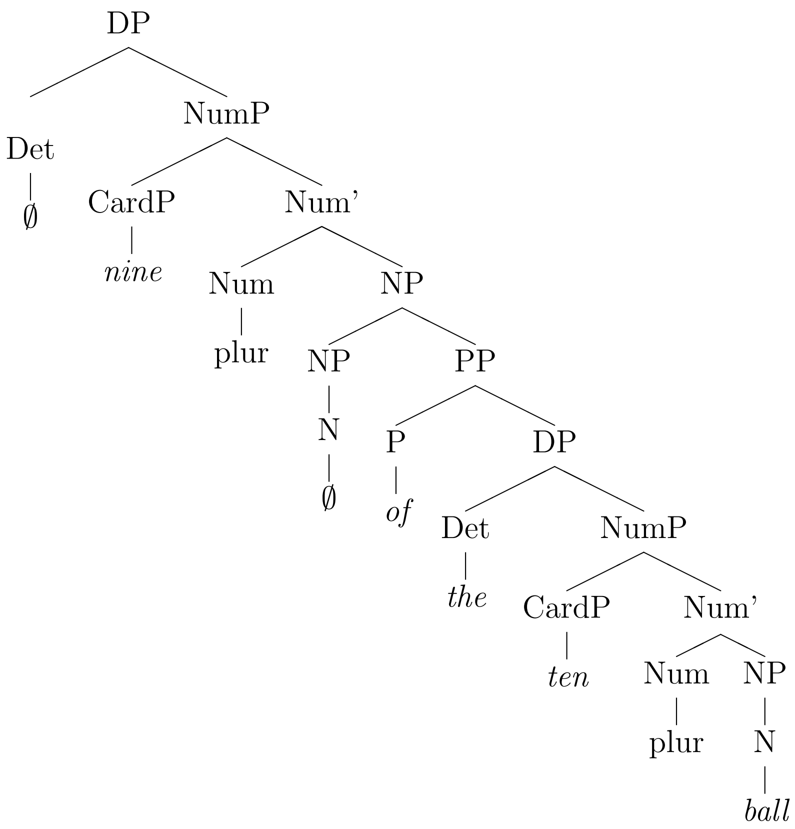

Now that we have decided on a syntactic structure for *nine of the ten balls* it is easy to extend this analysis to the singular DP *one of the ten balls*. The reason why we deferred looking at this DP, although it is the first one we must deal with when constructing DRSs for (2.52.a) and (2.52.b) in that order, is that the word *one* has a wider range of uses than the plural cardinals. Not only can it be used as the proper name of the number 1 and as a prenominal cardinal (as, say, in 'I bought only one book.'); it can also be used as an indefinite pronoun, as in 'If one realizes one has offended somebody, one should apologize.', or as pro-Noun, as in 'Mary bought a book and I bought one too.'.[16] But even though these alternative uses may suggest alternative analyses of the DP *one of the ten balls*, we take the parallels between *one of*

[16] I assume that there is an actual ambiguity in *one* as it occurs in a sentence like 'Mary bought some books before I bought one.' One one analysis of this sentence *one* is an indefinite pro-Noun, the N constituent of an indefinite DP with empty determiner. On a second analysis *one* is a prenominal cardinal, which modifies a phonologically empty Num' constituent. In speech these two interpretations will be distinguished in that the second structure is realized by putting a stress on *one*, whereas the first structure is verbalized by destressing *one*. I take it that in the sentence 'Mary bought a book and I bought one too.' *one* functions as an indefinite pro-Noun and that in 'Mary bought two books and I bought one.' it functions as a prenominal cardinal. But more should be said about what

the ten balls and *nine of the ten balls* to be decisive: In either case the first word of the string acts as a prenominal cardinal. One point worth noting in this connection is that *one of the ten balls* is, like *nine of the ten balls*, missing an overt determiner. For plural indefinites this is standard, but for singular indefinites it is not. It appears that in front of *one a* gets suppressed, perhaps for purely phonological reasons.[17]

This settles the syntax of *one of the ten balls*. In analogy with the syntactic structure of *nine of the ten balls* in (2.56), its structure of should be as in (2.57).

(2.57)

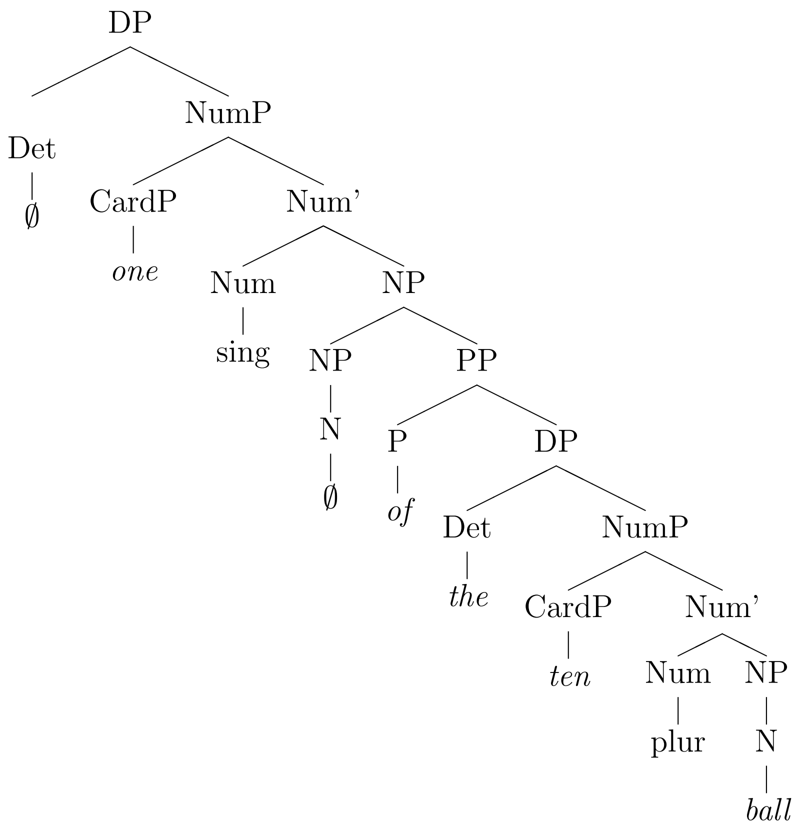

the disambiguating factors are that corroborate these claims. This is not the place for that, nor am I the right person to do it.

[17]Note that this is true also when *one* has another status than that of a prenominal cardinal, e.g. in ‘Mary bought several books after Fred bought (*a/$\surd\emptyset$) one.’

2.6.3 Singular and Plural Forms of Verbs

Now that the distinction between plural and singular is being made explicit for noun phrases, the question must be raised whether we shouldn't adjust our assumptions about the syntax of verb phrases as well. After all, English verb phrases also come in two forms, singular and plural, and this distinction isn't purely syntactic, it isn't just a requirement of number agreement of the finite verb with the subject; there are also semantic implications. Many plural verb phrases allow both for a collective and for a distributive interpretation. For instance, *hired a new secretary*, as in *They hired a new secretary.*, can be understood as describing a collective hiring of a single new secretary by the set of people denoted by the subject term, but it can also be understood distributively, as reporting on an occasion when each of the people in this set hired his own secretary, so that, presumably there were as many secretaries hired as there were people hiring.

I do not think, however, that these semantic considerations justify additional syntactic structure, with nodes at which the distinction between singular and plural is made or where the choice is made between collective and distributive readings of plural verb phrases. One reason is that the distinctions between plural and singular interpretations of verb phrases and that between collective and distributive readings of 'plural phrases' also arises in connection with arguments of verbs other than their grammatical subjects, where morphological agreement plays no role. Furthermore, in passive constructions it is the argument that functions as direct object in active constructions which now takes over the syntactic role of subject and thus becomes the noun phrase with which the verb must morphologically agree. In other words, the semantic differences just spoken of do not pair up with morphological number agreement in any straightforward manner.

This is not to say that the morphological phenomena would not require some systematic account somewhere in the syntax. But that is a matter that I propose we set aside here. And with that I propose that we do not build any provisions for the semantics of verb phrases into their syntax, but leave these distinction to be made when the syntactic structure is converted into its semantic representation. This means that at some point in the course of constructing the DRS decisions will have to be made about how it is to be interpreted – and, accordingly, about what its semantic representation should be – which are *not* dictated by the syntax. There is a kind of tradition within generative linguistics, embraced as much by Montagovians as by those who work with some form of Chomskyan generative syntax, that all such choices should be syntactically predetermined – that all ambiguities should be resolved at the level of syntax, with the possible exception of lexical ambiguity. For someone who acknowledges a level of semantics, whether in the sense of logical forms that are semantically and logically motivated (such DRSs) or in the form of semantic value assignments, it is hard to see why this should be the way the grammar of a language like English should be, or why it should be such a good thing to set the grammar up this way. What we will do here goes against it – only mildly, but it does.[18]

To summarize this discussion of plural and singular verbs: We assume the same syntax for verb phrases that we have been assuming so far. The number morphology on the verb will be ignored. We will just make sure that in the sentences we will be considering number agreement is always observed.

One last preliminary decision. The verb phrases of the sentences in (2.52) are *copula constructions*, consisting of a copular verb – the verb *to be* – and a *copula complement*. Copula complements can take three syntactic forms. They can be (i) Adjective phrases, such as *clever*, *quite clever*, *related to someone who has climbed Everest*, *feared by all and father to most*: (ii) Prepositional Phrases: *in the bag*, *on top of the roof*, *between Florence and Siena*: or (iii) noun phrases, such as *a famous rock singer*, *the owner of the BMW garage downtown*, *an asshole*. And these phrase types can also be combined into more complex copula complements, e.g. *clever and always on top of things*, *clever but a stickler for detail*, *either smart but frivolous or else just plain stupid*, *a professor, pleased as pie with himself and never on time*. Intuitively it is clear that copula complements always play the role of predicates that VPs consisting of a copula and a copula complement attribute to the subject. exactly what the role of the copula is hard to articulate precisely at this stage; this matter is better left to a point when we have learned enough about tense. So at this point we will focus on the *combination* of copula and copula complement and treat these combinations semantically simply as predications of the referential argument of the subject term. Syntactically

[18]In the spirit of this conception of ambiguity treatment, the assumption is often made that there must be some node in the syntactic projection of a plural verb that can house a 'distributivity operator'. When this position is filled by the (usually invisible) distributivity operator, then that forces the distributive interpretation, if the operator is not present in the syntax, then this gives rise to some other interpretation of the verb phrase, which is taken to have some kind of default status. Of course you can always set things up this way. But there ought to be good reasons for locating such decisions in the syntax. If there aren't, and the policy to make some aspect of the syntax responsible for the distinction rests on some general methodological principle, then that should be made explicit; or else it should be made plain that this is just some ultimately arbitrary design choice.

we will nevertheless analyze such combinations as syntactic compounds that are built from copular verb and copula complement as distinct constituents.

2.6.4 Back to our Balls

Before wee can present a syntactic structure for the first sentence of (2.52.a) there is one further issue we must settle. This is how we are to deal with the negation of this sentence. (The matter is orthogonal to what makes the discourses in (2.52) important for the present discussion, but a decision about it has to be reached nevertheless.) Earlier, in section 1.10, in which negation was discussed for the first time we treated negation as an adjunct to VP and the finite verb (a form of *do* to which the negation is directly attached – its 'support verb' as terminology has it – was implicitly assumed to be part of the overt realization of the negation. That is an oversimplification of the issue, to which we will return when we get to Tense and Aspect. In the first sentence of (2.52.a) we are facing a similar problem, but there is one difference. Whereas in the earlier sentences the support verb *do* was distinct from the main verb, in the sentence we must deal with now negation is supported by the verb form *is*, a finite form of the only verb there is in the sentence, viz. the copula *be*. To keep the syntax of our sentence as close as possible as the one we used for our earlier cases of sentence negation, we assume that the verb form *is* has been moved from its position as main verb into the VP adjunct position where it can do its duty as negation support.

With these and the earlier decisions we have made on behalf of the syntax of the first sentence of (2.52.a), the syntactic tree take the form shown (2.58). (For formatting reasons the structure had to be broken into two parts, with the VP displayed separately.

(2.58)a.

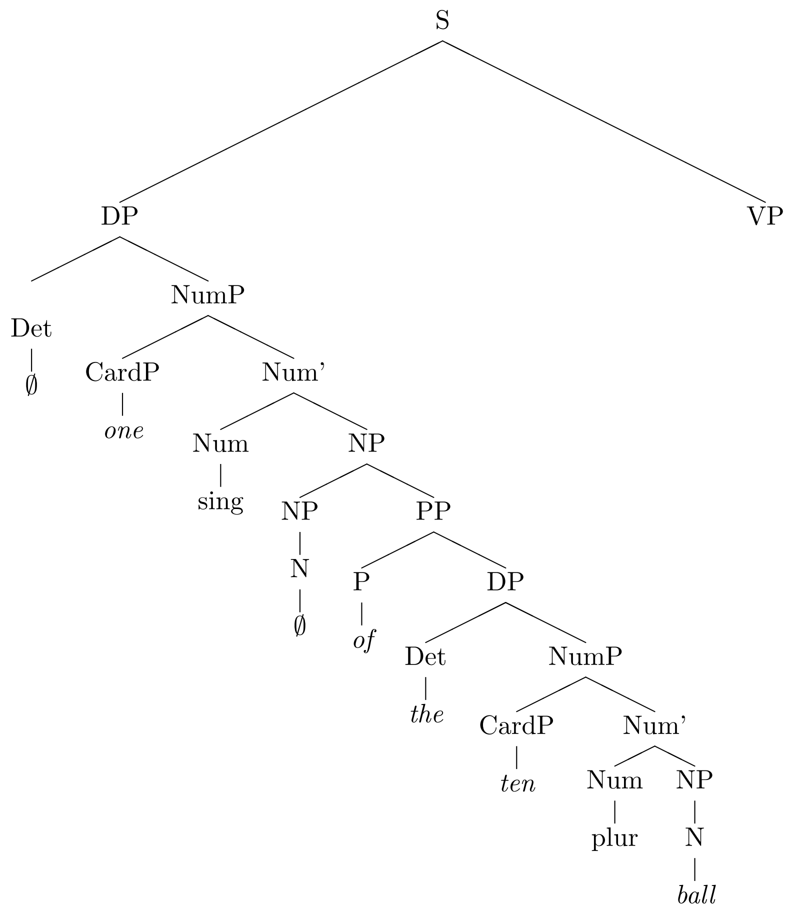
S
DP
VP
Det
∅
NumP
CardP
one
Num'
Num
sing
NP
NP
N
∅
PP
P
of
DP
Det
the
NumP
CardP
ten
Num'
Num
plur
NP
N
ball

b.

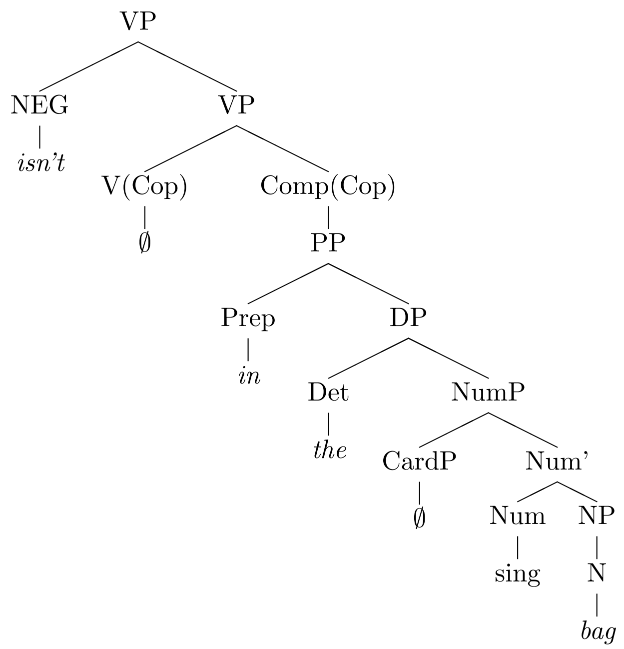
VP
NEG
isn't
VP
V(Cop)
∅
Comp(Cop)
PP
Prep
in
DP
Det
the
NumP
CardP
∅
Num'
Num
sing
NP
N
bag

Given what we have seen in earlier constructions and what has been said just now about the syntax and semantics of the first sentence of (2.52.a), the top-down construction of a DRS for the structure in (2.58) is fairly straightforward. Our first step decomposes the S node into a condition involving the sister to the Det constituent of the subject DP and one involving the VP. The sister constituent of the Det node is now one of the category 'NumP, but that makes no difference at this point. Since both conditions are complex and both are to be predications of the referential argument of the subject, we proceed as we did before when we encountered this kind of situation: we attach the dref for the DP in parentheses to the top nodes of the two subtrees.

But what dref should be chosen at this point? We now have a problem because we do not know yet whether the phrase we are interpreting is a singular or a plural. We could solve this by another policy of looking down into the structure of the phrase. This time, however, we will avoid the need for this by choosing a dref that is neutral between the kind we should choose if the phrase is a singular – recall that in that case we want to choose a lower case letter, just as we have been doing so far – and the kind we should choose in case the phrase turns out to be a set-introducing plural, which by our convention ought to be a capital letter. The neutral symbol we introduce at this point is a lower case Greek letter. In the present instance our (arbitrary) choice is the letter ξ. So it is ξ that we attach to the NumP node.

Finally, since the subject DP of the first sentence of (2.52.a) is an indefinite, we introduce ξ into the Universe of the DRS in which the operation is carried out, i.e. in the Universe of the main DRS (which is the only one we have). This gives us (2.59).

(2.59)

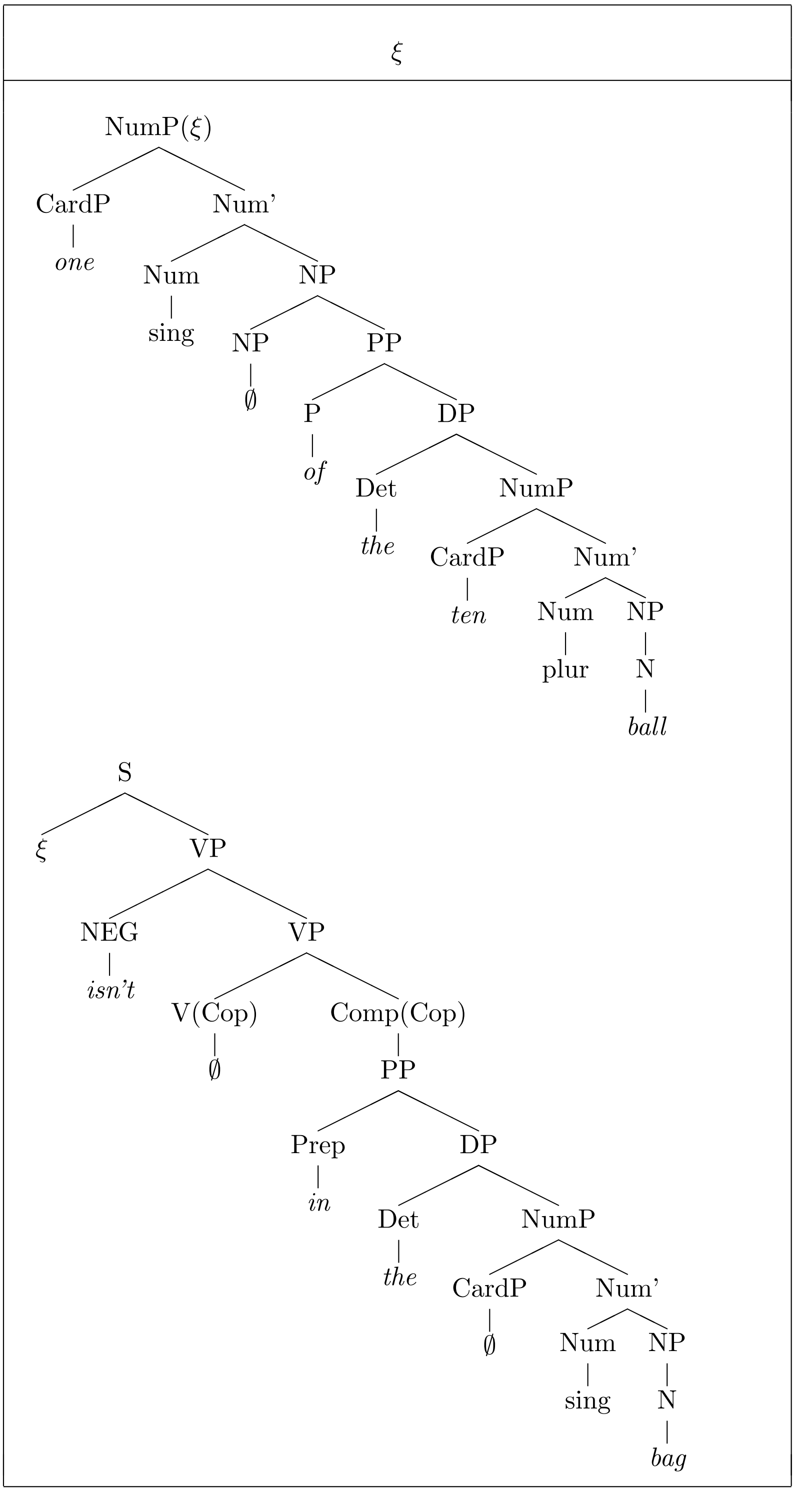
ξ
NumP(ξ)
CardP
one
Num'
Num
sing
NP
NP
∅
PP
P
of
DP
Det
the
NumP
CardP
ten
Num'
Num
plur
NP
N
ball
S
ξ
VP
NEG
isn't
VP
V(Cop)
∅
Comp(Cop)
PP
Prep
in
DP
Det
the
NumP
CardP
∅
Num'
Num
sing
NP
N
bag

We first deal with the interpretation of the NumP structure. The first reduction step concerns the combination of CardP and Num'. CardP is the cardinal *one*, which impose on ξ the condition that its cardinality equals 1. This is expressed by the condition '$|\xi| = 1$'. In order to save space we now remove the NumP- and the CardP-nodes while transferring '(ξ)', which expresses the predication relation, from the node labeled NumP to the one labeled Num'.[19] The result is (2.76).

[19]We could also keep the CardP node, the transfer of '(ξ)' from NumP to Num' being a sufficient indication that the NumP node has been dealt with. But for space and perspicuity reasons discarding seems preferable.)

(2.60)

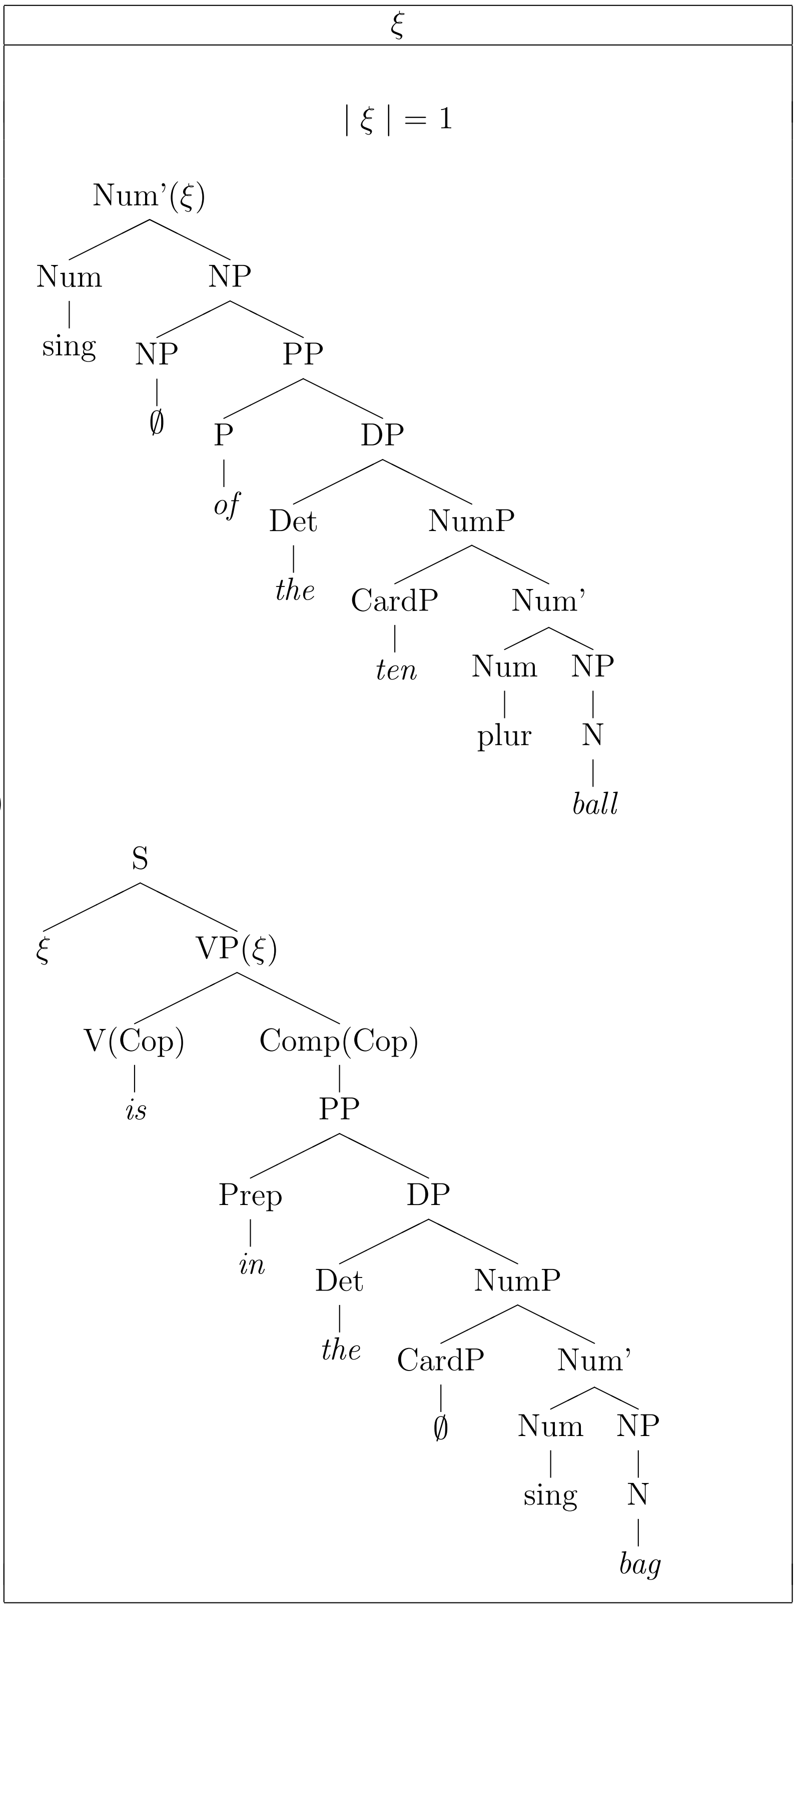
ξ
| ξ | = 1
Num'(ξ)
Num
sing
NP
NP
∅
PP
P
of
DP
Det
the
NumP
CardP
ten
Num'
Num
plur
NP
N
ball
S
ξ
VP(ξ)
V(Cop)
is
Comp(Cop)
PP
Prep
in
DP
Det
the
NumP
CardP
∅
Num'
Num
sing
NP
N
bag

The next step is to account for the contribution by the feature value sing of Num. This feature tells us that the complement of Num must be interpreted as a predicate of individuals. We implement this by replacing the 'non-committal' dref ξ by a dref designed to represent an individual. Let us use x for this purpose.

As far as the condition is concerned that prompts this step, the Num'-structure that remains after the immediately preceding step, this operation seems perfectly in order. We can proceed essentially as in the last step, discarding the top node of the Num' structure in (2.76) together with its left daughter. But now there is something awkward. We have to replace ξ by x not only in the condition on which the present step is being performed, but also elsewhere (in the Universe of the DRS in the condition '$|\xi| = 1$' and as argument of the VP). This may look like a violation of compositionality and a quite unwarranted one at that. But no worry!. This problem will resolve itself automatically later on when we switch from a set-theoretical to a mereological ontology.

Once again we discard the nodes that have now done their work (the Num- and the Num'-node). The result of this last step is displayed in (2.61).

(2.61)

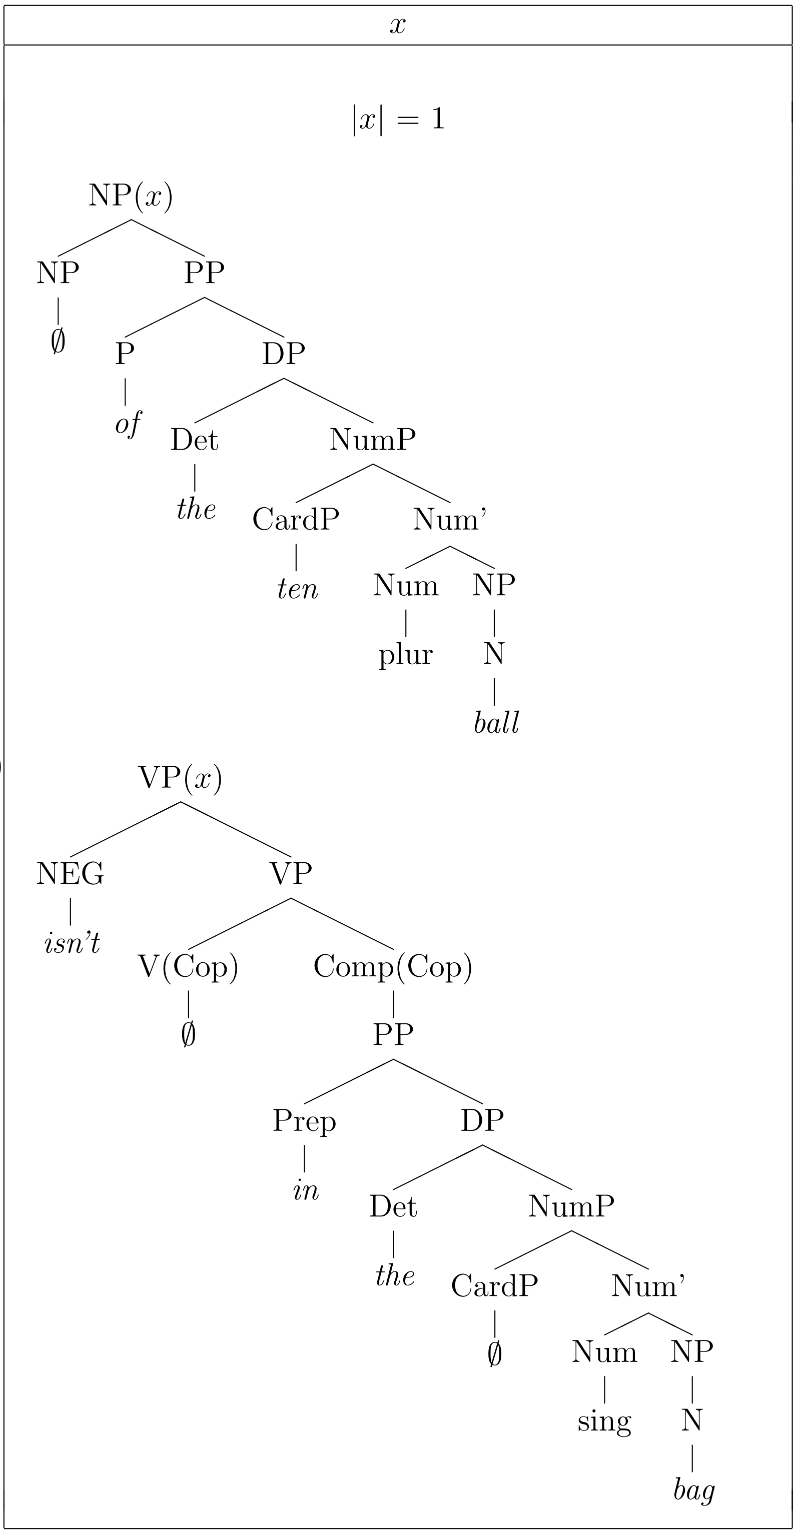
x
|x| = 1
NP(x)
NP
∅
PP
P
of
DP
Det
the
NumP
CardP
ten
Num'
Num
plur
NP
N
ball
VP(x)
NEG
isn't
VP
V(Cop)
∅
Comp(Cop)
PP
Prep
in
DP
Det
the
NumP
CardP
∅
Num'
Num
sing
NP
N
bag

Note another awkwardness that results from these operations. The cardinality condition now has the form '$|x| = 1$'. But is that condition coherent? What it wants to say is something that has now become a triviality, something like: 'if you take an individual x, that is just one'. When we think in terms of Set Theory and its standard notation, this is not the right form in which to express this triviality. Rather, what we should have is that the singleton set consisting of just x has cardinality 1: '$|\{x\}| = 1$. We could stipulate that the cardinality condition is now modified into this form. But no such stipulation is needed. This too is a matter that will fall into place when we move to a mereological ontology. So I will leave the apparent incoherence for what it is until we reach that point.

Next we have to deal with the combination of the empty N and its PP-adjunct. This combination can be treated as a case of ordinary 'predicate modification', but where the modified predicate is void, so that of the conjunction that normally results in such cases only the second conjunct, given by the adjunct, makes a real contribution. (We can think of the contribution of N as that of the trivial predication, which is satisfied by everything. Such a trivial predication could be added explicitly to the resulting representation but there seems no point in doing so. If we don't, then what we get is the structure given in (2.62).

(2.62)

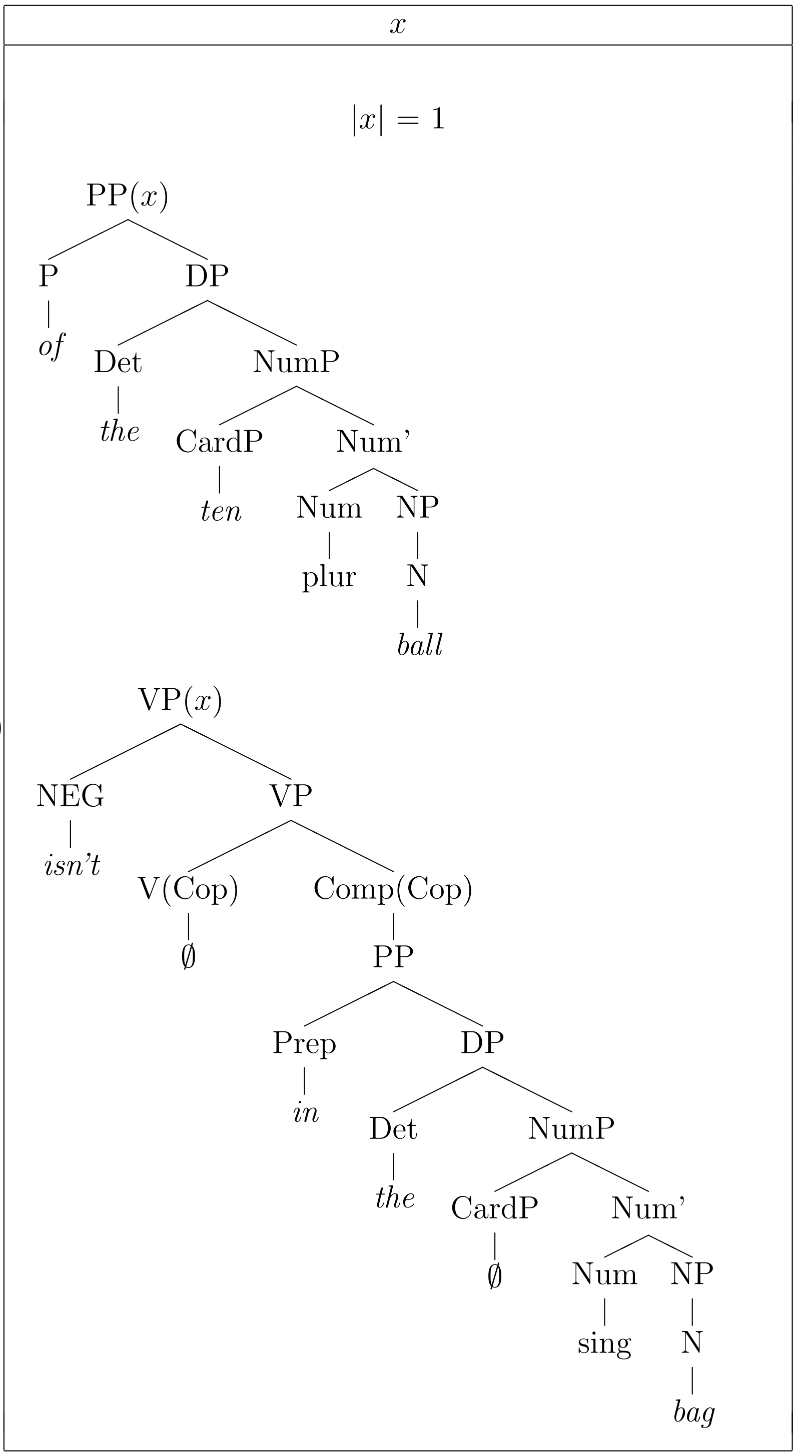

We must now deal with the contribution made by the 'partitive preposition' *of*. One problem that interpretations of occurrences of most prepositions have to cope with has to do with the fact that they are highly polysemous. Usually the intended meaning of a preposition occurrence can be inferred from the context in which it occurs. But it is by no means clear how we actually do this. And in many cases where it is intuitively clear how we do it, it may still be very difficult to state in formal terms exactly what the mechanisms are that are involved. We set this problem aside and take it as given what it is that *of* does when it occurs in a partitive construction like those found in the first sentences of (2.52.a,b). But what is this contribution? In the case before us right now it is simply that the individual, represented by the dref which is displayed as argument of the PP, is an element of the set denoted by the DP that is governed by *of*.[20] We will express the relation that *of* expresses in (2.62) by using the set-theoretic symbol '$\in$'. This too is a decision that will be reconsidered when we move to a mereological ontology.

To express the condition which says that the PP argument x stands in the $\in$-relation to the set denoted by the embedded DP we need to be able to refer to that set. That is, we must, at this point, introduce a dref representing this set. Let us choose η for this purpose. Since η is the set described or denoted by the embedded DP, we make that explicit by attaching '(η)' to its DP-node, and we also add η to the Universe of the DRS. So what we get is:

[20] It appears to be a general constraint on the grammaticality of partitive constructions that the embedded DP be a plural. At a minimum the DFP must denote a collection of two or more individuals. There are DPs that can do this while being morphologically singular. An example is the DP *the crowd*. Perhaps phrases like *one of the crowd* or *nine of the crowd* are not that bad. But may speakers seem to prefer *from* to *of* in such contexts, preferring e.g. *one from the crowd* to *one of the crowd*. I will not make efforts to account for this constraint but silly assume that the inputs to DRS construction satisfy it.

(2.63)

$x \quad \eta$

$|x| = 1 \qquad x \in \eta$

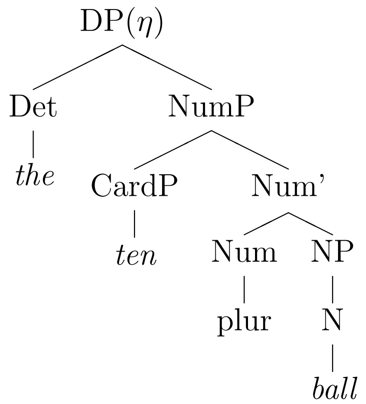

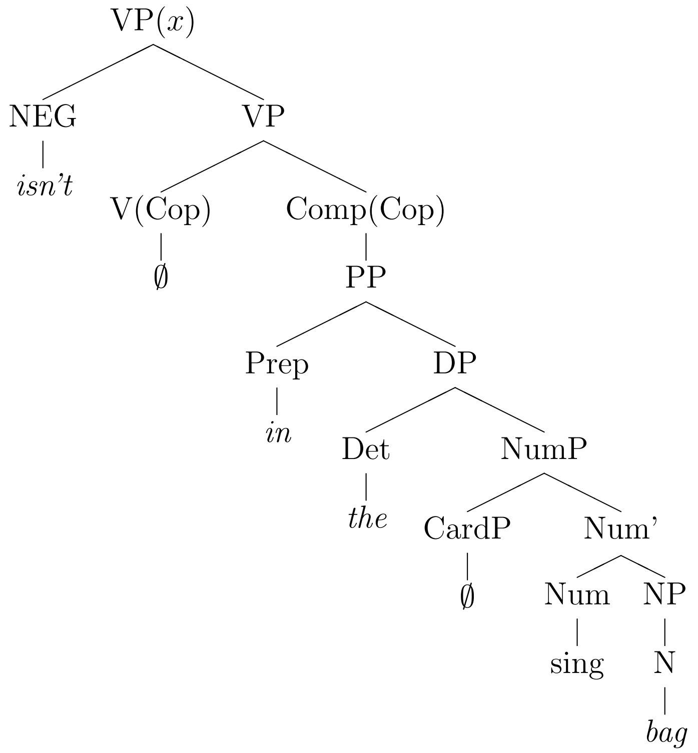

The only part of the subject DP that remains to be interpreted is the embedded DP of the PP. The first constraint that is imposed on the referential argument η of the DP is the maximality conveyed by the definite article *the*. To represent this constraint we follow in essence the recipe indicated in

(2.53.a). But in doing so we stumble on yet another problem. If we apply the recipe directly to the predicate identified by the NumP-tree, then we do not get what we want. For blind application would lead to the conjunction of the condition that η satisfies this predicate, viz. of being a set of 10 balls (and satisfying the additional contextual constraint C), and a condition that there is no proper superset η' of η which also satisfies this condition. But as it stands that cannot be what we want. For if η has 10 members, then obviously there can be no proper superset of η which also has 10 members. That is true, but it is true independently of whether η is a maximal satisfier of the predicate or not. It is clear that this is not what want, and it is also clear what we do want: in stating the second conjunct of the conditions that express maximality the CardP node has to be ignored.

There is no particular difficulty in stating the rule for processing the maximality expressed by *the* in such a way that things come out as we want them to: we just stipulate that the predicate which η' should satisfy is not given by the NumP-tree, but by the subtree headed by Num'. But this really does feel like a hack; and it is a hack that will not disappear when we pass to a mereological ontology. It appears, rather, that the maximality constraint of which we have been assuming that it is associated with the determiners of definite descriptions enters into the construction of definite descriptions much farther down. In fact, there is strong evidence that the definite article *the* itself enters the structure of definite descriptions at an earlier stage and that the traditional assumption, which we have been following here, that *the* occupies the Determiner position is either wrong, or that it ends up in a higher position via some kind of movement operation while nevertheless making its semantic impact in its base position. We will return to this issue later. For detailed discussion see (Coppock & Beaver 2014)).

For now we will, *nolens volens*, adopt the hack. What this comes to when made explicit in detail is this: we pass the referential argument η from the DP- to NumP-node and at the same time create a new condition saying that there is no proper superset η' of η that satisfies the predicate expressed by the subtree whose root is Num'. Moreover we add to both DRSs a condition to the effect that η satisfies an additional 'contextual' predicate C, without specifying what that predicate is. With these provisos what we get is the structure in (2.64):

(2.64)

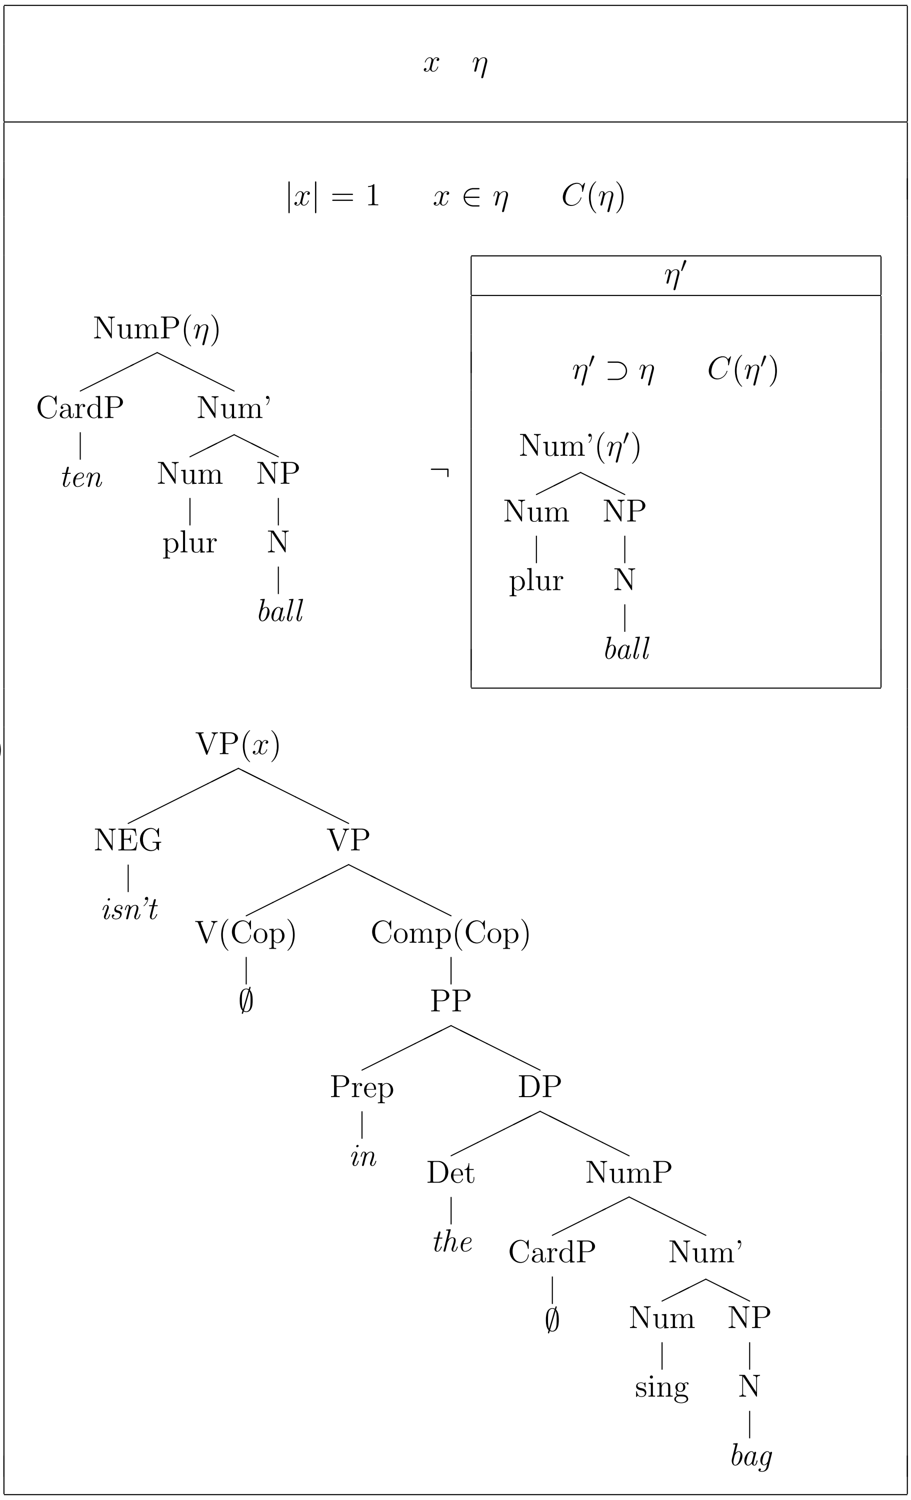

Our next step must deal with the NumP-node of the NumP-tree in the main Condition Set of (2.64). Since the operations operations involved in this step are familiar, we proceed directly to the result.

(2.65)

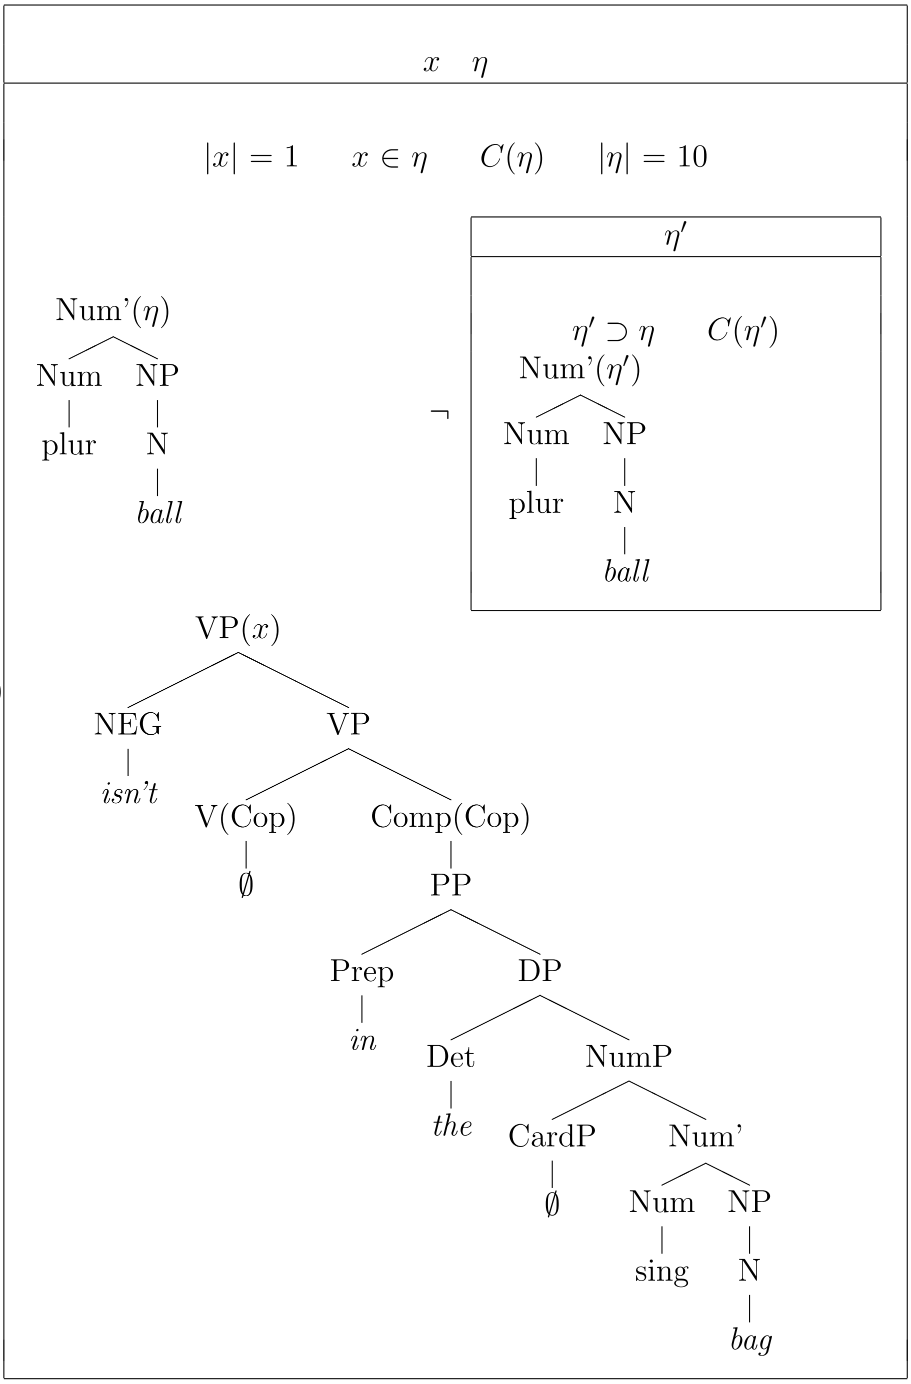

The next two steps that have to do with the subject phrase concern the contribution made by the feature 'plur'. 'plur' carries the information that the referential argument is a set of two or more elements rather than a single individual. We implement this by replacing η and η' by capital letters, choosing Y cane Y'. Again we discard the Num- and Num'-nodes, which at this point have done their work. The result is as in (2.66).

(2.66)

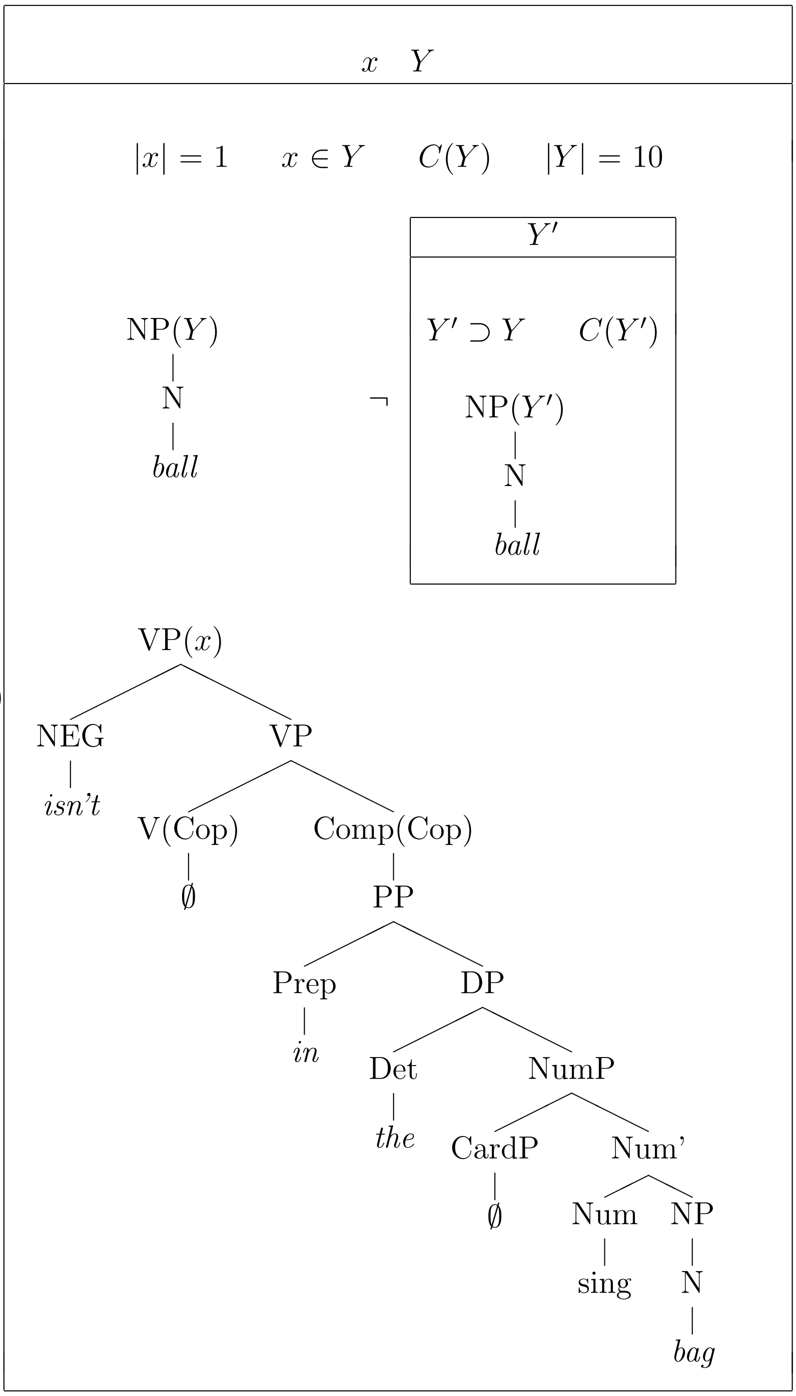

We can now also throw away what remains of the syntax of the subject phrase (i.e. the N- and NP-nodes). There is a question, however, what exactly should be understood by the conditions 'ball'(Y)' and '$C(Y)$' (and likewise for 'ball'(Y')' and '$C(Y)$'). Intuitively it is clear what the condition 'ball'(Y)' should say: it should mean that every member of Y is a ball. Likewise, in may applications '$C(Y)$' will amount to the claim that each member of Y satisfies the condition C. However, this is not always so and for this reason we leave the conditions '$C(Y)$' and '$C(Y')$' untouched.

There are various ways in which the principle that 'ball'(Y)' means that every member of Y satisfies the predicate 'ball' ' can be made explicit. The most common proposal in the literature is to admit an operation which turns ordinary nouns N – nouns whose extensions are understood as sets of individuals – into predicates *N whose extensions consist not only of the individuals in the extension of N but also of sets of such individuals. Thus if a and b belong to the extension of the predicate 'ball' ', then a, b and the set $\{a,b$ $\}$ will all be members of the extension of '*ball' '. For now we let this be, taking it as given that if X represents a set of individuals, then N(X) just means that N is true of every member of the set. With this convention the operation of the throwing away of the remaining syntax made up of the N- and NP-nodes becomes entirely straightforward.

Simultaneously with these operations we also perform the first reduction step on the VP-tree. The first step of this reduction detaches the VP adjunct and places its sister – now with '(x)' attached to its top node to make it into a condition that predicates the content of this structure of x – as sole condition of a new negated DRS. See (2.67).

(2.67)

$$x \quad Y$$

$$|x| = 1 \quad x \in Y \quad |Y| = 10 \quad \text{ball'}(Y) \quad C(Y)$$

$$\neg \; \boxed{\begin{array}{c} Y' \\ \hline Y' \supset Y \quad \text{ball'}(Y') \quad C(Y') \end{array}}$$

$\neg$

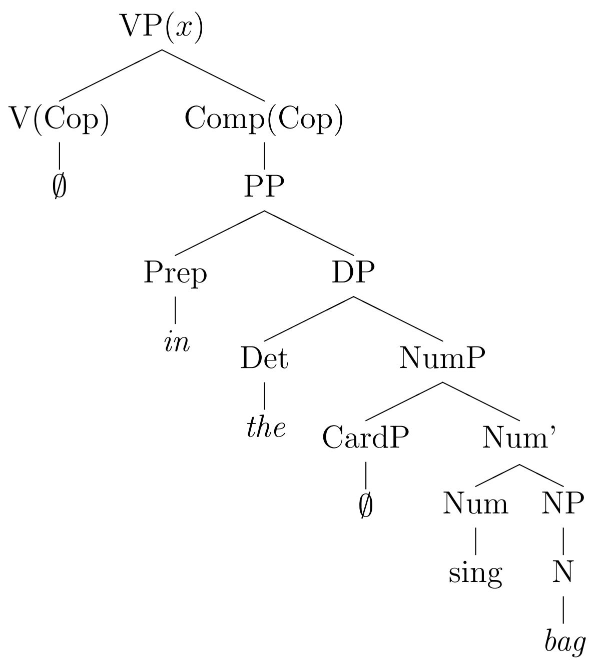

To reduce the VP-labeled condition in the negated DRS at the bottom of (2.67) we understand the V(Com) node as simply passing the argument x on to its complement. This gives us (2.68).

(2.68)

$x \quad Y$

$|x| = 1 \quad x \in Y \quad |Y| = 10 \quad \text{ball'}(Y) \quad C(Y)$

$\neg$ Y'

$Y' \supset Y \quad \text{ball'}(Y') \quad C(Y')$

$\neg$

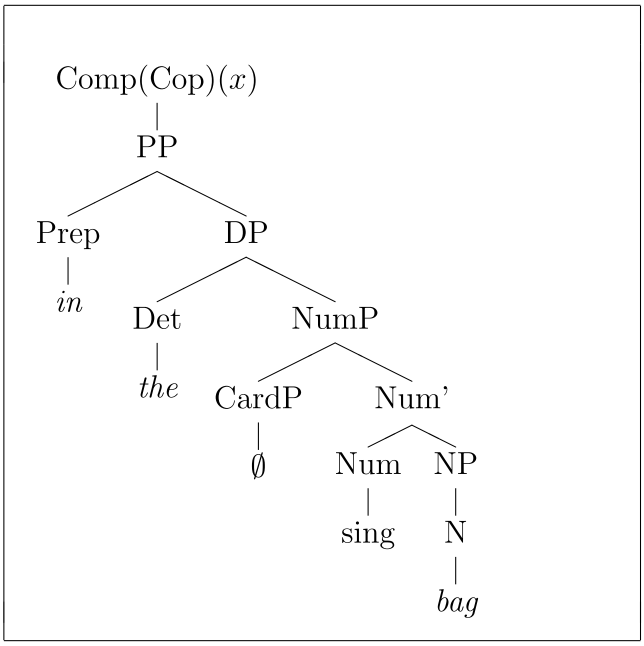

The remaining reducible condition says that x satisfies the predicate expressed by the copula complement. Since nothing happens between this node and the PP-node it dominates, this just amounts to x satisfying the predicate dominated by the PP node. The PP we are dealing with here is a 'canonical' one in which the preposition expresses a genuine relation between the referential argument x and the argument supplied by the DP that is governed by the preposition. To represent this relation we have to introduce a dref for the second referent, and at the point where we are we have no access to what kind of argument this will be – single individual or multi-membered set. So we choose again a neutral dref, ζ, in order to state the relation expressed by *in* as a relation between x and it and attach '(ζ)' to the DP-node of the governed DP. There is also a question as to which DRS Universe ζ should be inserted into. Here we face a similar problem that we encountered when dal ing with proper names and indefinite descriptions. Earlier we argued that the discourse referents for proper names shouls always be inserted into the main Universe, and that the accompanying DRS condition, which says that the dref stands for the bearer of the name should be inserted into the main Condition Set. We also noted that the matter could be dealt with properly only when we have a way of handling presuppositions, and postponed a proper discussion until then. With definite descriptions the matter is more complex. They presuppositions can sometime be justified locally, in which case their representing drefs should end up in some subordinate Universe. But there are also many cases where their presupposition is justified at the global level, in which case their dref and the Condition(s) specifying what they stand for must end up in the main Universe and Condition Set, respectively. The case before us is evidently of this latter sort: The phrase *the bag* can only be understood as referring to some particular existing bag, of which the sentence says that a certain ball is not in it. Again a proper treatment of all this must wait until we have a proper presupposition treatment. Here we simply insert, without further argument, ζ and the condition given by the syntactic structure for the DP *the bag* in their highest possible positions.

(2.69)

$$x \quad Y \quad \zeta$$

$$|x| = 1 \quad x \in Y \quad |Y| = 10 \quad \text{ball'}(Y) \quad C(Y)$$

$$\neg \; \boxed{Y' \;\Big|\; Y' \supset Y \quad \text{ball'}(Y') \quad C(Y')}$$

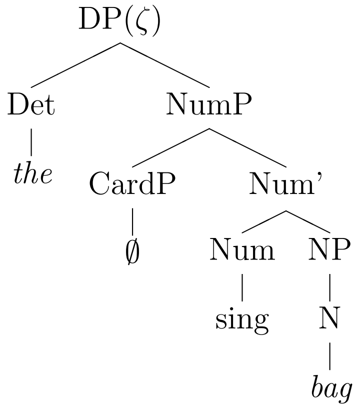

$$\neg \; \boxed{\text{in'}(x, \zeta)}$$

At this point all the operations needed to deal with the remaining reducible condition are familiar. So we move directly to the final DRS, leaving the remaining reductions as an exercise.[21]

[21]The occurrence of $\supset$ in the DRS below again seems wrong when its second argument is an individual (x) rather than a set. This is one of the problems that will be resolved when we adopt metrology.

$$
(2.70)\quad \begin{array}{|c|}\hline x \quad Y \quad z \\ \hline |x| = 1 \quad x \in Y \quad |Y| = 10 \quad \text{ball'}(Y) \quad C(Y) \\ \neg \begin{array}{|c|}\hline Y' \\ \hline Y' \supset Y \quad \text{ball'}(Y') \quad C(Y') \\ \hline \end{array} \\ \neg \begin{array}{|c|}\hline \text{in'}(x, z) \\ \hline \end{array} \\ \text{bag'}(z) \quad C'(z) \\ \neg \begin{array}{|c|}\hline Z' \\ \hline Z' \supset z \quad \text{bag'}(Z') \quad C'(Z') \\ \hline \end{array} \\ \hline \end{array}
$$

This has been a rather long haul and the outcome may seem rather modest given what we set out to achieve. With (2.70) we have obtained a DRS for the first sentence of (2.52.a) that can serve as discourse context for the second sentence, and in particular make it possible to get the intended interpretation for the pronoun *it*. Evidently this is possible since the main Universe of (2.70) contains a dref for the missing ball.

(Exercise: Use the principles introduced in the course of the construction above for the first sentence of (2.52.a) and our earlier principle for pronoun interpretation to obtain an interpretation of the second sentence in which *it* does refer to the missing ball.)

It should also be clear that when we construct a DRS for the first sentence of (2.52.b) on the basis of the same principles that have been used in constructing the DRS for the first sentence of (2.52.a), then no dref for the missing ball will be forthcoming. (Rather, we will now have drefs for (i) the set of ten balls and (ii) for some nine-membered subset of that set.) So our pronoun principle predicts that no interpretation of *it* as referring to the missing ball is possible in this case.

2.6.5 Once more: Non-distributive Interpretations of Plural Verbs

There is one aspect of the DRS construction for the first sentence of (2.52.b) for which there is no precedent in what we have shown so far. This is the plural form of the verb in the verb phrase *are in the bag.* Here, as for other plural verb phrases, the general question arises how the semantic relation between the verb phrase and its plural subject compares with that between the verb phrase and singular subjects. In fact, this question is a special case of a more general one: How is the semantic relation between a verb and a plural argument phrase occupying any of its argument positions connected to the relation between it and a singular argument occupying that same position? In general, there are at least two ways in which the relation between a plural DP and its verbal predicate can be analyzed. The first is the *distributive* interpretation, according to which every member of the denotation of the plural argument DP satisfies the predicate expressed by the verb phrase. For most verb phrases distributive interpretations are possible and sometimes they are the only ones.

But in addition to a distributive interpretation of the predicate-argument relation in which they stand to their subjects, many verb phrases also allow for one or more non-distributive interpretations. As an example consider the sentence in (2.71).

(2.71) Five lawyers hired a new secretary.

On its distributive interpretation (2.71) says that there was some set of five lawyers, each of whom hired a new secretary. (So presumably five new secretaries were hired.) But there is also a non-distributive – also called 'collective' – interpretation, according to which the five lawyers acted as a group, hiring a single secretary between them.

Non-distributive verb phrase interpretation is a topic in its own right, about which much has not yet been said. It is mentioned here only in passing, the main point being that non-distributive interpretations exist. For our current example – the first sentence of (2.52.b) – the issue is not directly important, since *are in the bag* has (as far as I can tell) only a distributive interpretation. And until further notice we will only consider distributive interpretations of verb-argument combinations.

Even when we focus exclusively on the distributive interpretations of such combinations, there remains the question how these interpretations are represented in logical form and what precisely triggers such representations if and when they arise. For now I will simply assume that when the argument phrase introduces as its representative a dref that represents a set of two or more members, then the distributive interpretation of the combination of the DP and its predicate takes the form of universal quantification over the set represented by this dref – a duplex condition in which the restrictor says that the dref bound by the quantifier belongs to the set represented by the DP and the nuclear scope DRS contains the predication expressed by the verb or VP.[22]

(2.72) gives the (uncompleted) DRS for the first sentence of (2.52.b) that we get when we apply this rule to its S node.

[22]When there are interpretation options for plural predicate-argument combinations besides the distributive one, these will, on the present conception of DRS construction, give rise to yet other construction rules It is sometimes suggested that the choice between the distributive option and others should be made explicit within the syntactic structure from which the semantic representation is derived. For instance, the distributive interpretation will, according to such proposals, be triggered by a distributivity operator DIST that occurs at the right place in the syntactic structure that serves as input to the DRS construction. But of course, shifting such decisions to the syntax doesn't give us by itself any answer to the real question: When should the syntactic structure be assumed to contain such an operator and when should it not?)

(2.72)

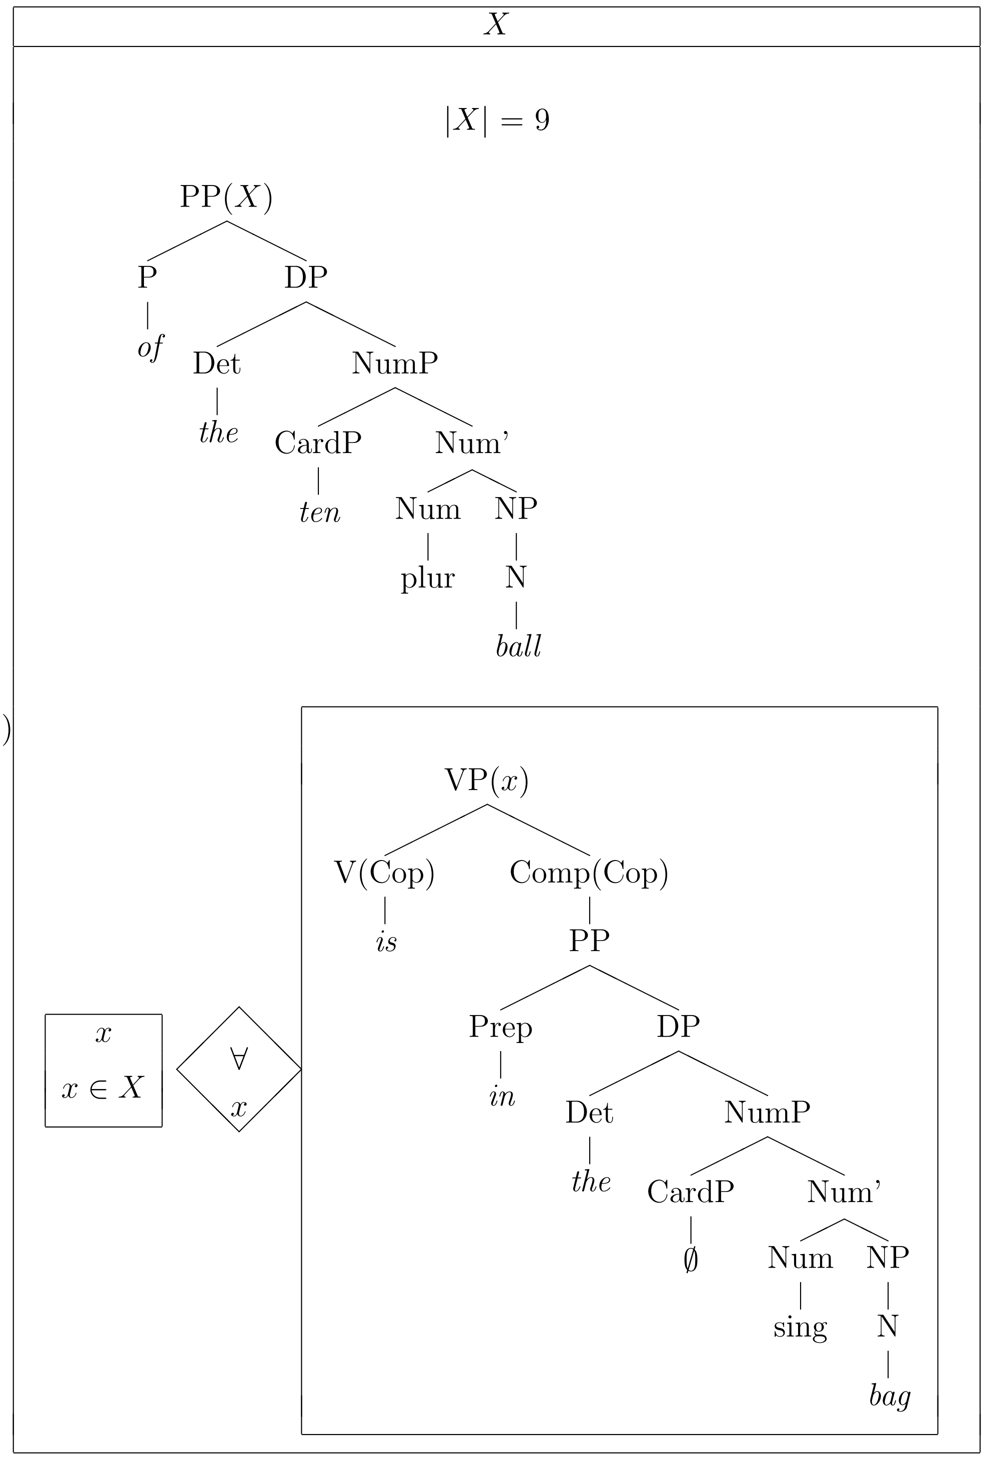

Exercise Carry out the remainder of the DRS construction for (2.72) and then construct the extension of the DRS you obtain in this way with the beginnings of a construction for the second sentence of (2.52.b) to the point where this construction cannot be carried any further, thereby showing that *it* cannot be interpreted in the discourse context that your DRS for the first sentence of (2.52.b) provides.

Exercise As we noticed, in (2.52.c) the interpretation of *it* as referring to the missing ball is possible again, and for intuitively obvious reasons. Formulate principles for the interpretation of the first sentence of (2.52.c) which leads to a DRS that does provide the wanted anaphoric antecedent for *it*.

2.7 The Methodological Implications of Partee's Ball examples

The central point that Partee's example in (2.52) is meant to establish is that propositional content as traditionally understood is not sufficient to account for certain aspects of pronominal anaphora: The first sentence of (2.52.a) and the first sentence of (2.52.b) are meant to express the same proposition, but nevertheless the pronoun *it* in the second sentence can be construed as referring to the missing ball in the one case but not in the other. To clinch the point, however, the interpretations of the two first sentences must assign them truth conditions that are unequivocally identical. And that is not so for the DRS in (2.70) that we constructed for the first sentence of in (2.52.a); nor is it so for the DRS that we get for the first sentence of in (2.52.b) if we proceed along the same lines that we followed in constructing the DRS in (2.70). For as things stand, the latter DRS only says that there is a missing ball, but not that it is the only one. Similarly, the DRS that we get for the first sentence of in (2.52.b) when we follow the same construction principles only says that nine of the ten balls are in the bag, but without excluding the possibility that the tenth ball is there as well. In order to clinch the case, therefore, we must modify the sentences so that this doesn't happen. The two sentences must yield DRSs with exactly the same truth conditions. One way to do this is to replace the examples in (2.52.a,b) by the following variants:

(2.73) a. Exactly one of the ten balls is not in the bag. It is under the sofa.

b. Exactly nine of the ten balls are in the bag. It is under the sofa.

But now we need processing principles for phrases like *exactly one* and *exactly nine*, which must guarantee that the resulting representations are equivalent.

Exercise Formulate principles for the interpretation of *exactly one* and *exactly nine* which in conjunction with the principles formulated above do produce truth-conditionally equivalent sentences for the first sentences of (2.73.a,b).

The phenomenon illustrated for *it* in (2.52) and (2.73) has a direct analogue for the plural pronoun *they* in examples like those in (2.74).

(2.74)a. Two of the ten balls are not in the bag. They are under the sofa.

b. Eight of the ten balls are in the bag. They are under the sofa.

Arguably such example pairs are perhaps even more remarkable than the one involving singular pronouns, for in general plural pronouns appear to be more flexible than the singular ones and to be usable also in cases where their antecedent is not available in the form of an accessible dref in the discourse context as is, but has to be constructed from other material that the discourse context does contain. One example is (2.75.a). (2.75.b,c), almost randomly chosen, give some further illustration of this aspect of plural pronouns.

(2.75)a. Freddie took one ball out of the bag. Andy took out another one. They are under the sofa.

b. John took Mary to Acapulco. They hated the place.

c. John took Mary to Acapulco. There they met Fred and Suzie. The next morning they set off on their sailing trip.

The *they* in (2.75.a) is naturally interpreted as referring to the combination of the ball taken out of the bag by Freddie and the one taken out of the bag by Andy. Somehow it is permissible to form a 'set' , or 'sum', out of the two mentioned balls and use that as the antecedent for the pronoun. (2.75.b) and (2.75.c) illustrate the same point. In (2.75.b) *they* refers to the combination of John and Mary. That is also the case for the first *they* of (2.75.c); but the second *they* is ambiguous between (at least?) three interpretations, one in which it refers to John and Mary, one in which it refers Fred and Suzie and one in which it refers to the four of them.

2.8 Mereology vs. Set Theory

To make this intuition – that the pronouns in (2.75) refer to sets or sums formed out of elements that were introduced explicitly into the discourse context – explicit within our current formal framework we have to make a decision that I have tried to postpone so far by talking about 'sets or sums'. By the distinction between 'set' and 'sum' hangs a longer tail – quite a long one in fact. Here I'll make the story as short as I can, tailoring it to our particular concerns. The terms 'set' and 'sum' belong – and are paradigmatic for – two different theoretical approaches to a very basic structure, which is central to the foundations of mathematics but which also proves to be fundamental for the semantics of natural languages; and there it arises in particular in connection with the distinct semantics of singular and plural. The two approaches are known by the names of *Set Theory* and *Mereology*. Within mathematics Set Theory seems to have won the day. In fact, it did quite a long time ago. Set Theory provides one of the frameworks for developing all of pure mathematics – all objects that are talked about in different branches of mathematics (numbers, functions, vectors, vector spaces and so on) can be construed as sets of one kind or another. (There are alternative ways of providing general foundations for mathematics, such as Category Theory, which does not use the concept of set (and set-membership) as basic, but I am unaware that Mereology survived for long as a serious contender – if it ever was one.)

The sets of Set Theory provide what is in essence a layered totality: You have sets at the lowest level, then sets that can be formed from these sets, then sets that can be formed from all that has become available so far, then sets that can be formed of what this last operation has made available and so on. The picture is straightforward when we start with some collection D_e of entities that are not sets. Let a, b, c,.. be members of this collection. Then the first layer, or 'stratum', of sets will include sets such as $\{a,b\}$, $\{a,c\}$, $\{a,b,c\}$, but also the *singleton sets* $\{a\}$, $\{b\}$, $\{c\}$, .. . All these sets are distinct from the elements from which they are formed. That seems plain for sets consisting of more than one element, but it is also true for the singleton sets: $\{a\}$ is not the same as a etc. The difference between singletons and their members is essential to how Set Theory works.[23]

[23]The 'layerdness' of standard Set Theory is not strict but *cumulative*. For instance, at the second level of set formation, we do not only find sets like $\{ \{a,b\}, \{a,c\}\}$ $\{\{a,b\}, \{b,c\}, \{a,b,c\}\}$ or $\{\{a,b\}\}$ (the singleton set whose only member is the set consisting of a and b) and $\{\{\{a,b\}, \{a\}\}$, but also sets like $\{a, \{a,b\}\}$ and $\{a, b, \{a,b\}\}$, which consist of a mixture of sets and members of D_e.

The fundamental notion of Set Theory, apart from the notion of 'set' itself, is that of *set-membership* the relation that holds between a set an each of its members (or *elements*, another word for the same notion). This relation is traditionally denoted as '$\in$', with the members of the set mentioned in first position and the set itself in second. (Thus we write' $a \in \{a, b\}$ to express that a is a member of the set $\{a,b\}$.)[24]

The fundamental difference between Set Theory and Mereology is that the central relation of Mereology is the *part-whole* relation, which we will denote here as $\preceq$. $\preceq$ is not like $\in$, which holds, and only holds, between entities of different strata. ('$a \in \{a, b\}$' can be true only because $\{a,b\}$ belongs to a higher stratum than a, '$\{a,b\} \in \{\{a,b\}, \{b,c\}, \{a,b,c\}\}$' can be true only because$\{\{a,b\}, \{b,c\}, \{a,b,c\}\}$ belongs to a higher stratum than $\{a,b\}$, and so on.) The entities related by $\preceq$ all belong to a single 'stratum'; or better, there are no strata in Mereology. The strata found in the standard way of doing Set Theory are, according to Mereology, an artifact of the way Set Theory is conceived, and should be dispensed with if possible. And the central claim that some advocates of Mereology have made is that it has found a way of doing that: by axiomatizing the fundamental properties of $\preceq$ it is possible to provide the same general foundation for mathematics and semantics that Set Theory provides through its axiomatization of the properties of $\in$.

The single stratum of Mereology contains everything that populates the different strata of Set Theory – its individuals (in case there are any) the sets that can be formed from these and all that can be generated within Set Theory from those. it is quite common in Mereology to refer to all these things – all that is included in its single stratum – as 'individuals'. This termonology can be confusing to someone socialized and operating within the world of Set Theory. To the distinction that is made in set Theory between individuals and sets corresponds in Mereology that between *atomic* and *non-atomic* individuals.

This is not the place to go into the details of these formalizations, let alone

[24]It is seven possible to set up Set Theory – and this is the version usually considered by Set Theorists, who study the realm of sets for its own sake – in such a way that $\in$ is the only primitive notion. On the conception underlying this set-up everything is a set, so there is no need to introduce 'set' as a predicate that distinguishes sets from non-sets. And the 'cumulative hierarchy' of sets spoken of above now starts with the 'empty set', a set with no members whatsoever, which can be defined without reference to any other sets.

to evaluate, on the basis of such formalizations, the respective merits of Set Theory and Mereology as a framework for the foundations of mathematics. But in relation to the special interests and needs of a theory of singular and plural in natural languages like English Mereology has much to say for itself. So from now on we adopt Mereology, without further ado, as our ontology.

To my knowledge the attractions of a 'mereological ontology' as a basis for the semantics of singular and plural terms in natural languages were first clearly recognized by Godehard Link, who brought Mereology into Formal Semantics. (See in particular (Link 1986)). Link had a number of related reasons for this proposal. One was an intuition that it is unnatural to have to assume that singular and plural referring trams refer to entities of different logical types, with singular terms referring to individuals and plural terms referring to sets of two or more individuals. (According to the usual set-theoretic definition of domains in models for the λ-calculus, and systems based on it such as the Higher Order Intensional Logic used in Montague Grammar, the individuals belong to the base domain D_e and the sets to the higher functional domain $D_{<e,t>}$.) This intuition seems particularly plausible in connection with phrases like *one or two women* – what sort of thing does this phrase refer to, an element of D_e or an element of $D_{<e,t>}$? And how can we represent the contribution of this term to the semantics of sentences containing it (such as, say, 'One or two women came.')? If we assume, as we have been doing, that this semantic representation takes the form of a DRS, then this representation will have to contain a dref representing the subject sand that dref will have to stand for what is either of type e or of type $<e,t>$. So the dref has to act like a variable that can range over a domain that consists of elements of different logical types, and that is very much against the spirit of the λ-calculus and systems based on it, which have very strongly shaped the thought about semantics and ontology since Montague started using them. (We will shortly see another example which requires a semantic representation with a variable that would have to range over a combination of D_e and $D_{<e,t>}$ on a set-theoretic conception.)

A further, more specific reason that Link had for preferring a mereological ontology had to do with the definite article *the*. Mereology allows for a simple uniform treatment of the contributions that *the* makes to singular and plural DPs. In both cases, Mereology enables us to say, *the* expresses the condition that the DP refers to the *maximal satisfier* of its descriptive content. (Recall our treatment above off the DP *the ten balls*, to which we will return presently.) The difference between singular and plural definite descriptions is now that the feature 'singular' imposes the constraint that what is described is an atomic individual, whereas the feature 'plural' imposes the constraint that what is described is a non-atomic individual.

The mereological ontology we adopt involves the following formal assumptions:

(i) The 'domain of individuals' D is partially ordered by the relation $\preceq$. ($\preceq$ is a weak partial order of D, meaning that it is (a) reflexive, (b) antisymmetric, and (c) transitive.)

(ii) $< D, \preceq >$ is a *complete upper semi-lattice* in that for every non-empty subset X of D there is an element ΣX of D with the following two properties: (a) $(\forall x)(x \in X \rightarrow x \preceq \Sigma X)$; (ii) if d is any element of D such that $(\forall x)(x \in X \rightarrow x \preceq d)$, then $\Sigma X \preceq d$. (So, intuitively, ΣX is the smallest element of D (in the sense of $\preceq$) that contains all elements of X as parts.)

When the set X consists of two members – say, $X = \{a, b\}$ – then we write '$a \oplus b$' instead of ΣX $(= \Sigma\{a, b\})$.

N.B. this abbreviation specification entails that $\oplus$ is (a) *idempotent*, (b) *commutative* and (c) *associative*. For instance: $a \oplus a = a$ (idempotency) ; $a \oplus b = b \oplus a$ (commutativity), and $(a \oplus b) \oplus c = a \oplus (b \oplus c)$ (associativity).

(Exercise: Prove these identities.)

The *atomic* elements of a part-whole structure $< D, \preceq >$ are those elements of D which do not properly contain any other element of D: For any d of D, *atomic*(d) iff there is no d' in D distinct from d such that $d' \preceq d$. The *non-atomic* elements of D are all the others.
For the semantic framework within which we are working, adopting a mereological ontology comes to this. The models with respect to which DRSs are evaluated for truth and falsity now each have a mereological structure for their 'domain of individuals'. An embedding function f that verifies a DRS K in such a model must map the drefs in K's Universe that are specified as drefs for atomic individuals onto atomic individuals in the model and drefs that are specified as drefs for non-atomic individuals onto non-atomic individuals. drefs that are not specified one way or the other may be mapped either on atomic or on non-atomic individuals of the model.
Before we say more about the examples in (2.75), let us, now that we have switched to a mereological ontology for our models, return to our earlier treatment of the phrase *one of the ten balls*. We repeat the most directly

relevant representation from the sequence of representations that were displayed earlier to show the DRS construction for the sentence *One of the ten balls is not in the bag.*, viz. (2.76).

(2.76)

ξ

$|\ \xi\ | = 1$

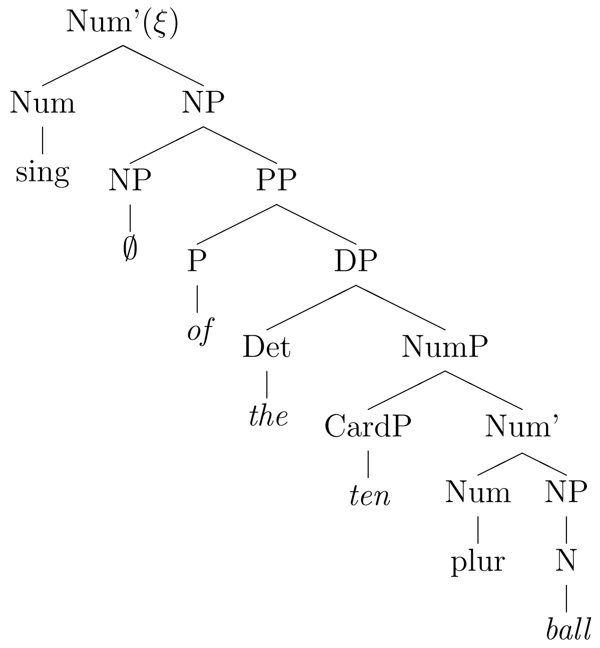
Num'(ξ)
Num
sing
NP
NP
∅
PP
P
of
DP
Det
the
NumP
CardP
ten
Num'
Num
plur
NP
N
ball

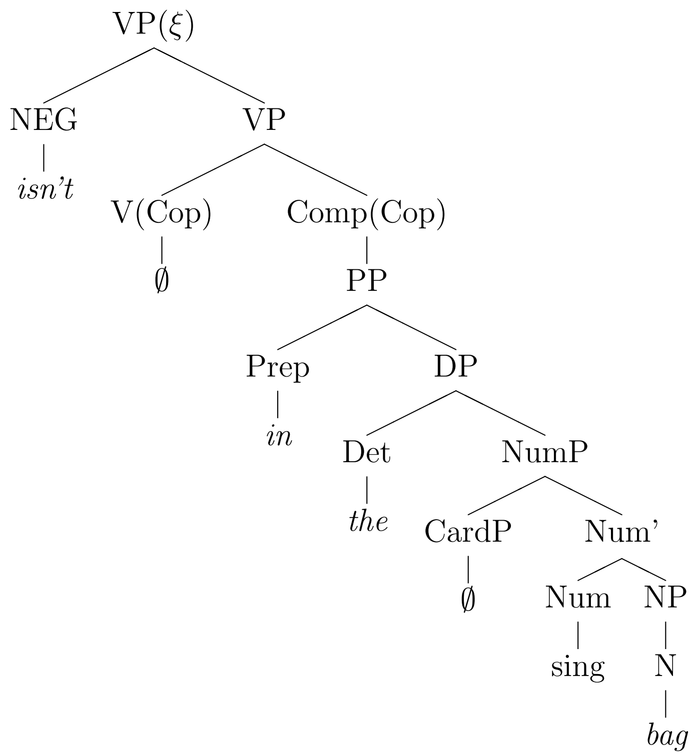
VP(ξ)
NEG
isn't
VP
V(Cop)
∅
Comp(Cop)
PP
Prep
in
DP
Det
the
NumP
CardP
∅
Num'
Num
sing
NP
N
bag

Recall one of the difficulties we ran into when constructing the semantic representation of *one of the ten balls* that occupies the upper part of (2.67). Our first step in the construction of this representation involved introducing for the DP *one of the ten balls* as a whole the 'neutral' dref ξ, since at that plaint the information provided by the feature 'sing' was not yet available. But then, when that feature did become available, we changed ξ into x, to indicate that we were dealing with a single individual; but the awkward side to that was that we had to change *all* occurrences of ξ to x, which at that point were already spread over different DRS-conditions. With our new ontology this is no longer necessary. We can introduce two new predicates, *atomic* and *non-atomic*, into our DRS language, corresponding to the 'atomic-non-atomic' distinction in our models, and reinterpret the contribution made by 'sing' to be that of adding the condition 'atomic(ξ)' to the Condition Set to which the condition containing the 'sing' feature belongs. Once this condition is added there is no need to change ξ into x, since the condition constrains the possible valuations of ξ to atomic individuals anyway.[25] For the embedded DP *the ten balls* the story is similar. Instead of replacing η by Y we add the condition 'non-atomic(η)'.

2.9 Synthesizing Antecedents for Plural Pronouns

At long last we return to the examples in (2.75). First, the DRS construction for (2.75.a), repeated here.

(2.75.a) Freddie took one ball out of the bag. Andy took out another one. They are under the sofa.

We have seen applications of all the rules we need to construct the DRS for the first sentence, provided we treat the verb as a 3-place predicate 'x took y out of z'. The resulting DRS is shown in (2.77).

[25] We can of course still replace all occurrences of ξ by x while at the same time omitting the condition individuals has 'atomic(ξ)' as a way of simplifying the notation, but that is a different matter. What matters is that the earlier problem no longer arises for the official notation, which involves adding 'atomic(ξ)'.

$$(2.77)\quad \begin{array}{|c|}\hline f \quad x \quad z \\ \hline \text{Fred}(f) \quad |x| = 1 \quad \text{ball'}(x) \\ \text{bag'}(z) \quad C(z) \\ \neg \begin{array}{|c|}\hline Z' \\ \hline Z' \supset z \quad \text{bag'}(Z') \quad C(Z') \\ \hline \end{array} \\ \text{took-out-of'}(f,x,z) \\ \hline \end{array}$$

A principled construction of the DRS for the second sentence would require some further principles. The first of these has to do with ellipsis: 'took out' is somehow short for 'took out of the bag'. But then we also want also to infer that it is the same bag that was mentioned in the first sentence. This can be inferred from the principle that the instance of *the bag* that is recovered by the ellipsis construction invokes the same tacit restrictor C as its overt source in the first sentence. The second issue has to do with the semantics of *other*. But we pass over these tangential issues and simply present the update of (2.77) with the information contributed by the second sentence.

$$(2.78)\quad \begin{array}{|c|}\hline f \quad x \quad z \quad a \quad x' \\ \hline \text{Fred}(f) \quad |x| = 1 \quad \text{ball'}(x) \\ \text{bag'}(z) \quad C(z) \quad \neg \begin{array}{|c|}\hline Z' \\ \hline Z' \supset z \quad \text{bag'}(Z') \quad C(Z') \\ \hline \end{array} \\ \text{took-out-of'}(f,x,z) \\ \text{Andy}(a) \quad |x'| = 1 \quad \text{ball'}(x') \quad x' \neq x \\ \text{took-out-of'}(a,x',z) \\ \hline \end{array}$$

The first step that must be performed as part of the DRS construction for the third sentence involves the plural pronoun *they* that constitutes its subject DP. For pronouns we now also need a way of making the distinction between singular and plural explicit. I confess to having no good idea as to what the syntactic implications and constraints might be that come with taking this matter seriously. So I will stick with the simple-minded syntactic representation format that we have been using so far for both pronouns and proper names, according to which these expressions are directly attached to the DP node itself. The only modification of this that we now adopt is that the 'ϕ-features' of the expression are now explicitly listed behind the word. Thus the feature specifications that come with the pronouns *he*, *she*, *it*, *they* and proper names like *Freddie* are as in (2.79).

(2.79)a.
```
    DP
    |
he[3rd,masc,sing]
```

b.
```
    DP
    |
she[3rd,fem,sing]
```

c.
```
    DP
    |
it[3rd,neut,sing]
```

d.
```
    DP
    |
they[3rd,plur]
```

e.
```
    DP
    |
Freddie[sing]
```

For now only the number feature on the pronouns is relevant. In conjunction with the information that the DP is a pronoun and that it is to be interpreted anaphorically – as we have seen, that is a decision which depends on factors we cannot analyze here; so we assume that it is made by fiat – the number feature tells us whether the antecedent dref should be one representing an atomic or a non-atomic individual.

Given that the subject DP of the third sentence of (2.75.a) is the plural pronoun *they*, the discourse context should provide as antecedent for it a dref that stands for a non-atomic individual. But obviously the DRS in

(2.78) doesn't supply any such dref. That *they* can be given an anaphorric interpretation, viz. as referring to the two balls that were taken out of the bag, indicates that it is possible in this case to 'build' an antecedent from the material that the context DRS makes explicitly available. The reconstruction of this interpretation process that has been proposed in DRT (cf. (Kamp & Reyle 1993)) is that it is possible to interpret a plural pronoun by 'forming sums out of drefs occurring in positions accessible from that of the pronoun'. To this end we add a 'merreologic sum operator' to the vocabulary of our DRS formalism which enables us to specify drefs as sums of other drefs. We use the symbol '$\oplus$' for this purpose.[26] Thus, in the case at hand wee can form the 'sum X of the drefs x and x' and use this X to interpret the pronoun *they*. If we assume that the pronoun itself introduces the plural dref Y, then the two steps needed for the intuitively natural interpretation of *they* are shown in (2.80.a,b).

[26]This constitutes a slight overload of notation since we are already using this symbol in the metalanguage, in which we are defining upper semi-latices and the mereological models based on them, but the overload is quite harmless.

$$
\begin{array}{|c|}
\hline
f \quad x \quad z \quad a \quad x' \quad X \\
Y' \\
\hline
\\
\text{Fred}(f) \quad |x| = 1 \quad \text{ball'}(x) \\
\\
\text{bag'}(z) \quad C(z) \quad \neg
\begin{array}{|c|}
\hline
Z' \\
\hline
\\
Z' \supset z \quad \text{bag'}(Z') \quad C(Z') \\
\hline
\end{array} \\
\\
\text{took-out-of'}(f{,}x{,}z) \\
\\
\text{Andy}(a) \quad |x'| = 1 \quad \text{ball'}(x') \quad x' \neq x \\
\text{took-out-of'}(a{,}x'{,}z) \\
\\
Y' = x \oplus x' \\
\\
\hline
\end{array}
$$

(2.80)a.

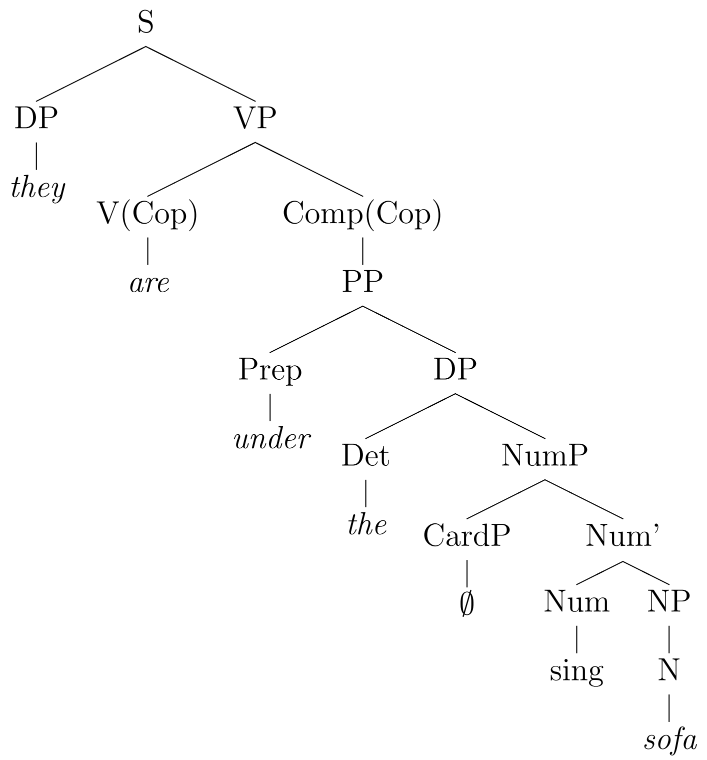

$$
\begin{array}{|c|}
\hline
f \quad x \quad z \quad a \quad x' \quad X \quad Y' \quad Y \\
\hline
\text{Fred}(f) \quad |x| = 1 \quad \text{ball'}(x) \\
\text{bag'}(z) \quad C(z) \quad \neg \begin{array}{|c|}\hline Z' \\ \hline Z' \supset z \quad \text{bag'}(Z') \quad C(Z') \\ \hline \end{array} \\
\text{took-out-of'}(f,x,z) \\
\text{Andy}(a) \quad |x'| = 1 \quad \text{ball'}(x') \quad x' \neq x \\
\text{took-out-of'}(a,x',z) \\
Y' = x \oplus x' \quad Y = Y' \\
\hline
\end{array}
$$

b.

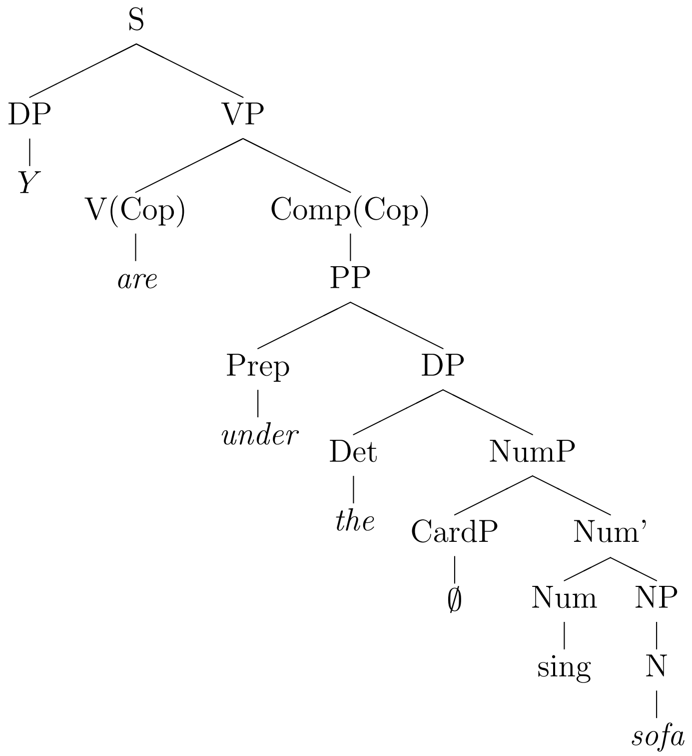

Exercise Complete the DRS construction for (2.75.a).

The DRS construction for (2.75.b) proceeds along the same lines as that for (2.75.a). The only difference is that now the sum needed to interpret the pronoun *they* of the second sentence is this time to be formed from drefs introduced in the same sentence.

Exercise Carry out the DRS construction for (2.75.b).

2.9.1 Conjunctions of DPs (and more Non-Distributivity)

The DRS-construction for (2.75.c) raises two new issues. The first is that of the construction of a semantic representation for the plural DP *Fred and Suzie*, the second is the interpretation of the plural verb phrases *met* in the second sentence and *went on a sailing trip* in the third. As regards the first issue, note that *Fred and Suzie* is a plural DP at least in the morphological sense that it requires plural marking of the verb when it occurs in subject position. But in fact, we will also analyze it as *semantically* plural, viz. as denoting the non-atomic individual consisting of Fred and Suzie (i.e. the mereological sum of the two). Part of such an analysis is an interpretation of the word *and*, when it conjoins two DPs, as expressing the sum operation $\oplus$ – or, what comes to the same thing for our purposes, as representable with the help of the operator $\oplus$ we have added to our DRS language. As regards the syntactic structure of the DP *Fred and Suzie* that we want to adopt for present purposes we have to make a choice between several options. One is to extend the feature marking we have adopted for pronouns and proper names also to conjunctions of such DPs. This leads to the representation in (2.81). An alternative is to adopt a more elaborate structure for DPs of the kind proposed above, which contains a Number Phrase projection between NP and DP. In that case it seems reasonable that the conjunction is formed at or below the Number Phrase projection level, but whether it is formed at that level or below is a further choice point for the theory and it is not entirely clear to me what the best solution to this problem is. But let us not get side-tracked by these further questions and make do with the simpler structure, given in (2.81).

(2.81)
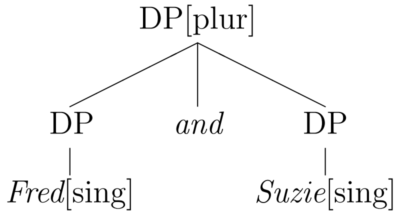

One property that has been built into (2.81) is the feature 'plur' as number value specification for the outer DP. This is another aspect of (2.81) that has been designed to make DRS construction somewhat simpler.[27] The presence of 'plur' has the advantage that when we process the DP (by introducing a dref for it and in serving this dref in the right places), we are entitled to choose right away a dref whose values are restricted to non-atomic individuals. Thus we are entitled to choose the upper case letter Y for this purpose. This dref gets inserted into the direct object slot of the verb *met* and it gets attached to the outer DP node of the DP conjunction.

What into insertion into the object position of the verb comes to in this case has to do with the second issue, which will be addressed momentarily. But our next step is to further analyze the DP structure to which the dref has been attached as argument. It is here that the contribution of 'sum-*and*' comes into play: Y is now represented as the mereological sum of the drefs introduced by the DP conjuncts. (In choosing the drefs representing these conjuncts we can make use once more of the information supplied by the number feature specified on each of the two conjunct DPs. In both instances this feature is 'sing' and that entitles us to choose drefs for the representation of these DPs that are restricted to atomic values. The decompositional analysis of the conjunct DPs then proceeds as before, but of course in this case (where both DPs are proper names) there isn't much to unfold for our construction algorithm.) The structures in (2.82.a-c) show the results of these successive construction steps. For easier reading the condition representing the verbal predication, as well as all else that is involved in the representation of that predication, has been omitted.

(2.82)a.

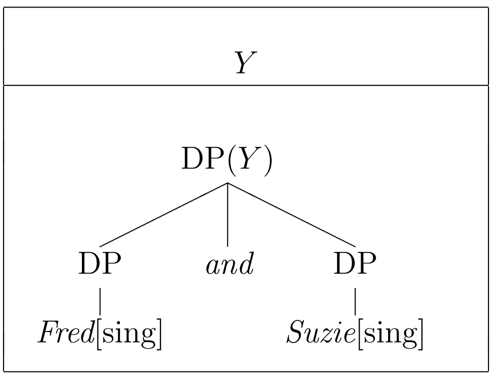

[27]This is part of the general simplification adopted here and we should not get embroiled in the question whether or how this is syntactically justified (just as we should not get hung up at this point on the question whether it is right to dispense with the Number Phrase projection level).

b.

$Y \quad f \quad s$	
$Y = f \oplus s$	
DP(f)	DP(s)
|	|
Fred[sing]	*Suzie*[sing]

c.

$Y \quad f \quad s$		
$Y = f \oplus s$	*Fred*(f)	*Suzie*(s)

The second issue concerns the plural verbal predicates *met* and *set off on their sailing trip*. The verb *meet* is a verb that allows for non-distributive interpretations and for which such interpretations are often the preferred ones. This is so also for the past tense occurrence *met* of *meet* in the second sentence of (2.75.c). Salient among the possible interpretations of the verbal predication in this sentence is the one where it is the sum of John and Mary that stands in the *met* relation to the sum of Fred and Suzie.[28] Formally, the semantic representation of the predication involving *met* in (2.75.c) is unproblematic: we insert the (plural) dref for the subject phrase of the sentence into the subject slot of the predicate 'met" expressed by English *met* and the (plural) dref for the direct object phrase into its direct object slot. But DRSs with such conditions can be correctly evaluated only in models in which the extensions of the predicate 'met" contain pairs involving non-atomic individuals as well as pairs in which both members are atomic. This raises various questions about the internal structure of such extensions. For instance, when the extension of 'met" contains the pair $< a \oplus b, c \oplus d >$, must it also contain the pairs made up from the atomic individuals belonging to these sums, such as $<a, c>$, $<a, d>$ etc? A large part of the semantics of

[28]A doubly distributive interpretation of this predication, according to which each of John and Mary met each of Fred and Suzie, is one of the formally available interpretations in this case. But this certainly isn't the interpretation that leaps from the page. In fact, a more systematic treatment of distributive and non-distributive readings of transitive verbs with plural DPs in both argument positions would have to make clear exactly what the different interpretations are that a predication like that in (2.75.c) makes available. This is a notoriously hairy problem, something of which we have been aware since the early pioneering work of Scha (see (Scha 1981)).

plural arguments has to do with such questions about the extensions of the predicates in which they occur as arguments.

A question similar to the one about *met* can be raised about *set off on their sailing trip*: How are the interpretation or interpretations of this verbal predicate in combination with a plural subject phrase related to its interpretation or interpretations when it occurs in combination with singular subjects? Here too, there is a strong intuition that plural subjects allow for interpretation of the predicate that cannot be reduced to predications of atoms in any straightforward manner. In particular, the natural way of understanding the third sentence of (2.75.c) seems to be one on which the going off on a sailing trip is a 'collective' action on the part of those referred to by the subject phrase: an action that is the execution of a joint intention to make this trip together, and not, say, an occasion where each of them decided to make such a trip on his or her own, but where they accidentally they ended up in the same boat. As in the case of *met*, the analysis of such collective interpretations will have to be salt with within the setting of a general theory of non-distributive interpretations. Such a theory is not among our goals here and so won't do more than draw attention to the issue.

In the absence of a proper account of non-distributive predications the best we can do at the level of their semantic representation is to annotate the DRS conditions that represent non-distributive interpretations of such predications to indicate that this is how the DRS interprets the predication. For instance, we can mark predicate argument combinations that are given a non-distributive interpretation with the feature '-distr'. Since more-place predicates can get a non-distributive interpretation with respect to some argument places and distributive interpretations with respect to the others, the annotation has to make explicit to which argument position it pertains. The simplest way to make this explicit is to add the feature to the occurrence of the dref that fills the argument position in question in the representing DRS condition. Thus a non-distributive interpretation of the last sentence of (2.75.c) – more precisely: of a simplification in which the PP *the next day* is omitted – in which the predication expressed by the VP is marked in this way as non-distributive will have the form in (2.83). (We assume that X represents the subject *they* and leave the anaphoric resolution of X unresolved.)

(2.83)

X
set-off-on-a-sailing-trip(X_{-distr})

Obviously the use of the feature '-distr' doesn't explain anything; it can only be understood as a signal that a proper analysis is missing. But not only that, even as a means of indicating what interpretation is intended '-distr' doesn't provide us with the semantic differentiation one might want. This is because some predications of non-atomic arguments seem to allow for more than one non-distributive interpretation. An example is a sentence like that in (2.84)

(2.84)The Republicans voted for the proposal.

Like most sentences in which a verb that allow singular arguments occurs with a plural argument phrase, (2.84) has a distributive reading accordion to which every member of the set of Republicans voted for the proposal. But besides this interpretation we can discern two others, both of which are non-distributive. According to he first a majority of the relevant set of Republicans viz. those who were in a position to vote on the proposal, voted for it. The second interpretation is one according to which the relevant set of Republicans acted as a single agent. As a group they had just one vote, by some means or other they determined what that vote should be (viz. that it should be in favor) and then communicated to the relevant authority that that was their vote. If we consider these two interpretations as genuinely distinct and want to have semantic representations that make this distinction explicit, then the one feature '-distr' evidently won't be enough.

Since we are not aiming for a proper account of non-distributive predication, there is not much point in continuing this discussion. So I leave this difficult issue as a blank our account that still needs to be filled.

<u>Exercise</u> Construct a DRS for the second sentence of (2.75.c) in which the verb *met* is interpreted as non-distributive in both its subject and its direct object position.

<u>Exercise</u> Construct DRSs for (2.75.c) which represent all possible interpretations of the pronoun *they* in the third sentence. Give reasons for why the readings your DRSs acknowledge are all the interpretations of the pronouns

there are.

2.9.2 Other Ways of constructing Antecedents for Plural Pronouns

In the examples discussed above the antecedents we needed for the interpretation of plural pronouns could be constructed from drefs present in the discourse context by applying $\oplus$ to drefs that were already present in the DRS (and accessible from the position of the pronoun). But there are also cases where the intended antecedents of plural pronouns cannot be constructed in this way. Examples are the pronoun occurrences in (2.85).

(2.85)a. Susan has found every book/most books/only few of the books that Bill needs. They are on his desk.

b. Susan has found only few of the books Bill needs. He is disappointed, for he badly needs all of them.

c. You never see a hedgehog in winter. They hibernate.

The natural interpretation of *they* in (2.85.a) is that in which it refers to the books that Bill needs and Susan found. (2.85.b) is different in that here it is more plausible that *they* is meant to refer to the books that Bill needs, and not to the disappointingly small subset of those that Susan has managed to find. Apparently, the first sentence that (2.85.a) and (2.85.b) share can supply both of these antecedents.

In DRT the two anaphoric antecedents that the contexts of (2.85.a) and (2.85.b) make salient are assumed to be obtainable through application of what is called the *Abstraction Principle*. The Abstraction Principle applies to duplex conditions and forms the sum of all the values for a dref that occurs either in the Universe of the restrictor or in that of the nuclear scope DRS and that satisfy a constraint associated with the duplex condition. The constraint can take two forms: It can be given either (i) by the restrictor DRS on its own (in which case the dref over which we sum belongs to the Universe of that DRS) or (ii) by the merge of that DRS and the nuclear scope DRS, in which case the dref can be from either universe. The first form of Abstraction is needed for the interpretation of (2.85.b), the second for the interpretation of (2.85.a). In both cases the dref over which Abstraction sums is that introduced by the noun phrase containing the noun *book(s)*.

We need notation to express these Abstraction operations. We use Σ for this purpose. Σ binds the dref whose values are being 'summed' and its operandum is always a DRS. As an illustration of what this looks like consider the simplification in (2.86) of the first sentence that (2.85.a) and (2.85.b) share.

– this is the case of (2.85.b) –

(2.86) Susan found most books that Bill needs.

Let us assume that the DRS for this sentence is the one in (2.87).

(2.87)

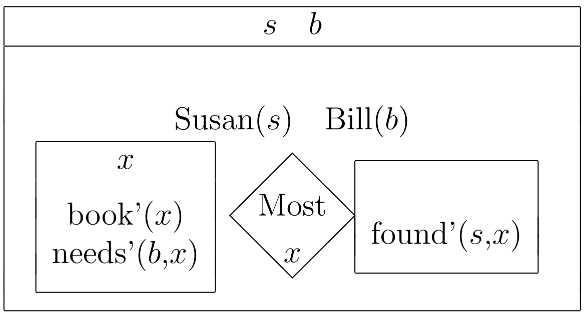

The two drefs η and ζ that can be obtained from the duplex condition in (2.87) by Abstraction (with respect to the one and only dref x that is found in the Universes of restrictor and nuclear scope) dref are given in (2.88).

(2.88)

$$\begin{array}{|c|}\hline s \quad b \quad \eta \quad \zeta \\ \hline \\ \text{Susan}(s) \quad \text{Bill}(b) \\ \\ \begin{array}{|c|}\hline x \\ \hline \text{book'}(x) \\ \text{needs'}(b,x) \\ \hline \end{array} \; \left\langle \begin{array}{c} \text{Most} \\ x \end{array} \right\rangle \; \begin{array}{|c|}\hline \\ \hline \text{found'}(s,x) \\ \hline \end{array} \\ \\ \eta = \Sigma x. \begin{array}{|c|}\hline x \\ \hline \text{book'}(x) \\ \text{needs'}(b,x) \\ \hline \end{array} \\ \\ \zeta = \Sigma x. \begin{array}{|c|}\hline x \\ \hline \text{book'}(x) \\ \text{needs'}(b,x) \quad \text{found'}(s,x) \\ \hline \end{array} \\ \\ \hline \end{array}$$

(2.85.c) exemplifies yet another way in which a sentence may make the antecedent for a plural pronoun available. This last possibility arises when the pronoun is interpreted as referring to a kind. Many kinds are denoted by nouns, and when a noun is understood in this capacity, it functions very much like an ordinary proper name. But on the whole nouns are far more often used as nominal predicates than as kind names. In fact, all uses of nouns in the examples we have so far considered are arguably of that sort and have been treated as such in our DRS constructions; and treating them that way does not involve, as we have seen, the introduction of a dref for the kind that the noun could be taken to name. But even if there is no call for the introduction of such a dref for an occurrence of a noun in a sentence in order to represent the semantic content of the sentence correctly, it harbors the potential for such an interpretation nonetheless, and this potential may be activated by a subsequent sentence when it contains a pronoun that needs the kind-representing dref for its interpretation.

It doesn't follow from this of course that we couldn't introduce a dref for the kind named by the noun as well, but in the examples we have considered the discourse referent wouldn't have made any significant contribution. It would have contributed nothing to the truth conditions of the DRS (it would only have added the usually disconnected, and in any case trivial information that there is a kind named by the noun) nor would it have served any purpose, as anaphoric antecedent. But in this last respect (2.85.c) is different. If, as we have been assuming throughout for pronouns in this presentation of DRT, the pronoun *they* of this example is to be resolved by identification with a dref, then it seems clear that this dref must represent the kind 'hedgehog', for that is what intuitively speaking this occurrence of *they* stands for. Note well, however, that it is only the presence of *they* in the second sentence that provokes the need to interpret the word *hedgehog* occurring in the first sentence as introducing the kind it names. So, given our policy for not introducing kind drefs each time a common noun makes its appearance, the conclusion should be that it is the need to interpret *they* in (2.85.c) that leads to the introduction of the kind dref 'retroactively'.

One feature of the kind name interpretations of common nouns is that the kind drefs nouns are capable of introducing can serve as antecedents for pronouns irrespective of the positions of noun and pronoun; no matter how deeply embedded the position of the noun and no matter what position is occupied by the pronoun, the pronoun always has access to the kind dref that can be introduced on the strength of the presence of the noun. In fact, (2.85.c) provides an illustration of this insofar as the noun *hedgehog* is part of the indefinite *a hedgehog* which is in the scope of the quantifier *never*. Because the indefinite occurs in this logically embedded position the dref representing it will end up in the restrictor DRS of the duplex condition that *never* introduces (I am assuming that that is what *never* does). So that dref will *not* be accessible from the next sentence.[29] But the kind dref that can be introduced because of the presence of the noun *hedgehog* that is part of this noun phrase *is* accessible from the following sentence. And that is because as kind name the noun functions as *name*; in this regard these 'names' are just like proper names.[30] For us this means that the kind drefs introduced by occurrences of common nouns always go into the Universe of the main DRS.

[29] Exercise: Show this by looking at examples in which the first sentence of (2.85.c) is followed by a sentence in which the pronoun *it* is trying to resume the indefinite *a hedgehog*, but where this doesn't work.

[30] The term 'common name', which is found for instance in John Stuart Mill (e.g. his *Logic*) and other writers from that time (and well after that) gets things just right. Names, whether proper or common, name particular entities and the naming relation is impervious to scope.

One curious aspect of the kind referring use of pronouns we are discussing is that it is *plural* pronouns that are used for this purpose. Why should it be plural rather than singular pronouns? As a matter of fact, singular pronouns can also be used to refer to kinds. But the conditions under which they can do this seem to be different from those that license the use of kind referring plural pronouns. (I at least get a clear contrast between (2.89.a), in which *it* replaces *they* in (2.85.c), and which for me is not good, and (2.89.c) which for me is perfect. And (2.89.b) seems even worse than (2.89.a).)

(2.89)a. ? You never see a hedgehog in winter. It hibernates.

b. ?? You never see hedgehogs in winter. It hibernates.

c. The hedgehog is never seen in winter. It hibernates.

The difference between the two types of case can be described as follows. The *it* of (2.89.c) is an instance of the familiar pattern in which a definite description is used to introduce a dref that represents the definite or indefinite term. That what is being referred to by definite description and pronoun is a kind (rather than a person, a hedgehog, an apple, a color, a virtue or whatever) is irrelevant to the anaphor-antecedent relation. The dref introduced for the definite description is an *individual* dref; what sort of individual it presents doesn't matter. The case illustrated by the good (2.85.c) and the bad cases (2.89.a) and (2.89.b) is different. Here the dref introduced by the common noun is a dref of a special type; it is a *kind dref* and not an individual (= atomic) dref that stands for an individual that in this instance happens to be a kind. In one sense a kind dref introduced by the common noun *hedgehog* and an individual dref introduced by the DP *the hedgehog* (when that DP is interpreted as referring to the species 'hedgehog') are standing for the same thing. But there is also a sense in which they do this in different ways, something that has to do with the fact that one is representing an argument (i.e. a DP) while the other is not.

That difference doesn't explain why kind drefs want plural rather than singular pronouns anaphors. The negative part of the story that needs telling can be told with the means we have in hand: Singular anaphoric pronouns, we have been saying all along, must find their antecedent ready made in the discourse context. That explains why they cannot serve as anaphors to kind drefs if we assume that such drefs are never introduced into the representation just on the strength of the occurrence of a kind denoting noun by itself, but only when there is something else that provokes their introduction. A singular pronoun lacks this provoking power. It tries to find an antecedent in the discourse context as is,and when it doesn't succeed in finding what it needs, then interpretation aborts.

But this is only half of the story. It doesn't explain why plural pronouns *can* serve as anaphors to kind drefs; and I have no compelling story to tell about that.

A further complication to kind anaphora has to do with the predicates in which the kind anaphors appear as arguments. *hibernate* is one of many which impose no constraints on the grammatical number of their arguments: It takes both plural and singular pronouns. But this appears not be true for all predicates. One exception is *be extinct.* (2.90.a) is fully acceptable, but quite a few speakers feel there is something peculiar about (2.90.b) and (2.90.c) seems to be out completely. When *is extinct* is replaced by *hibernates* – as in (2.90.d-f), with *dodo* replaced by *hedgehog* to keep absurdities to a minimum – all sentences are fine.

(2.90)a. $\surd$ The dodo is extinct.
b. ? Dodos are extinct.
c. * A dodo is extinct.
d. $\surd$ The hedgehog hibernates.
e. $\surd$ Hedgehogs hibernate.
f. $\surd$ A hedgehog hibernates.
g. $\surd$ Dodos are rare.
h. * A dodo is rare.

The difference in acceptability between (2.90.c) and (2.90.f) is due no doubt to the fact that hibernating is something that an individual can do on its own. Obviously that is not true of being extinct. The contrast between (2.90.b) and (2.90.e) doesn't seem quite as dramatic, which suggests that plurals can play a role that comes close to that of kinds. Note also that plurals certainly can be used with some inherently collective predicates, which is never possible for singulars, as shown by the sentences (2.90.g,h). (2.90.g) is false or misleading, but it is perfectly well-formed. (2.90.h) is gibberish.[31]

[31] There might be a way of interpreting this sentence that makes it acceptable, viz. by taking *a dodo* as short for something like *finding a dodo.* The intention of the present

The differences illustrated by the examples in (2.90) suggest that even when all three variants are possible and seem to express the same proposition, as in (2.90.d-f), there may be deep differences in logical form lurking below the surface. I suspect that (2.90.d) has the logical form of a subject predicate sentence in which the VP is used to predicate something of the kind denoted by the subject term. But that cannot be true of (2.90.f) for if that were the case, then the sentence ought to say that there is some hedgehog that hibernates, which may be true but is clearly not what the sentence means. But then, what *is* the logical form of (2.90.f)? One answer that comes to mind in the light of what we have been saying about indefinites is that (2.90.f) involves a tacit *generic quantifier*, which says of the relevant type of thing that all its typical or normal instances have a certain property. In our terms, this generic quantifier would give rise to a generic duplex condition, with the quantificational force just hinted at. Presumably the generic quantifier of this duplex condition will not be accompanied by a dref it binds, like the duplex conditions for adverbial quantifiers of which an example was given in (2.41). However, if we assume that the indefinite DP *a hedgehog* is placed in the restrictor of this duplex condition, then its dref will end up in the Universe of the restrictor DRS, so that the quantification will be *de facto* over hedgehogs. (I am assuming here that the subject DP should go into the restrictor and the VP into the nuclear scope of the duplex condition.[32])

If this is the right analysis of a sentence like (2.90.f), however, then we face a further problem. As we have so far described it, a sentence sequence like that in (2.91) should not be acceptable because the dref for *a hedgehog* is not accessible from the position of *it*. But apparently there is nothing wrong with the anaphoric connection between *it* and *a hedgehog*. How can we account for this?

(2.91) In the fall a hedgehog eats a lot. Then it finds itself a hole

discussion is that such reinterpretations, in which a DP is reanalyzed as a constituent of some large infinitival clause all else of which is silent and in which it plays the role of argument phrase, should be set aside. Of course, it may not always be quite so easy to identify an interpretation as involving such a reanalysis.

[32]Note well that this is an assumption which doesn't follow from anything we have actually said. As a matter of fact the question what goes into the restrictor of a duplex condition introduced by a generic or adverbial quantifier and what goes into its nuclear scope is a notoriously difficult question. Here too, information-structure plays an important role, as first observed by Rooth in this dissertation (Rooth 1985). (This point wasn't mentioned in our discussion of the examples in (2.42) in Section 1.9, where arguably it would have fitted just as well.)

and sleeps through the winter.

Superficially, the problem we are observing here with the basic DRT account of pronominal anaphora to indefinites seems to be of the same sort as the one we encountered in connection with the examples in (2.50) of Section 1.10. But there is a fairly broad consensus that for cases like that in (2.91) the solution to the apparent problem must be quite different from those we sketched for the earlier examples.

The phenomenon exemplified in (2.91) occurs with quite high frequency. It is known by the somewhat misleading name of *modal subordination*. Examples of the general form of (2.90.f) are easy to come by. Here is another.

(2.92) A lion eats meat. If you try to feed him a carrot, he'll rip off your hand.

(I feel a slight preference here for *he/him* over *it*. Perhaps we see this as distinctly masculine behavior.)

Other examples that are generally seen as illustrating the same general phenomenon are also from the mid-eighties. These are from Peter Sells.

(2.93)a. Every Korean rice farmer has a cart. He uses it to harvest his crop.

b. Every chess set comes with a spare pawn. It is taped to the inside of the lid.

These examples are telling in that here the indefinite DP that serves as antecedent to a pronoun in the next sentence is within the scope of the quantifier *every*. And moreover, in (2.93.a) the pronoun *he* picks up on the dref that is directly bound by *every* itself. The examples are telling because they seem to make it quite clear that in order to get their interpretation right we must allow the scope of *every* to somehow extend over the second sentence. But the question is exactly how this works. In a theory like the one we are discussing, we might be tempted to account for these cases by introducing a rule that allows for the placement of the second sentence into the nuclear scope box of the duplex condition introduced by the quantifier of the first. If such a rule was admitted it could also be used to take care of examples like (2.91) and (2.92). The only difference with the examples in (2.93) is that now it is the nuclear scope of the tacit generic quantifier that the follow-up sentence is placed in.

For the examples in (2.91) - (2.93) the rule seems to produce the right results. But without further justification the rule may rather feel like a hack. And moreover, we will have to make sure that it doesn't lead to serious overgeneration, predicting grammaticality where none exists and readings that nobody gets.

The term *modal subordination* was first introduced by Roberts in her 1987 dissertation, and was chosen by her because of examples like those in (2.94).

(2.94) A wolf might come in. He would eat you first.

What we see here is reminiscent of the examples we looked at above in that the indefinite *a wolf* seems to be within the scope of the modal operator *might*. So in order to be able to get the intended anaphoric link between it and the pronoun *he* we must get the latter somehow within the scope of this operator as well. This example poses additional complications that have to do with the interaction between two different modal operators, viz. the *might* of the first sentence and the *would* of the second, but it seems clear that here too a central part of the solution must be to find the mechanism which extends the scope of the operator *might* of the first sentence over the second.

Although there has been quite a bit of discussion and work on 'modal subordination' since Roberts brought these phenomena to the world's attention in (Roberts 1987) and (Roberts 1989), we seem to be still at some distance from a generally agreed treatment. For an approach like DRT the phenomena are important insofar as they show how hard it is to capture the anaphoric possibilities of pronouns in the form of simple principles like those which we have been using to account for their behavior.

These last few pages may seem to have side-tracked us from the topic of this section, viz. the possibilities for constructing antecedent drefs for plural pronouns. But much of what we have just gone through in our discussion of options for singular pronouns is relevant to plural pronouns as well, and in particular to their possibilities as anaphors to kinds. The first point to note is that in many cases the choice between a singular and a plural pronoun for the purpose of expressing an anaphoric relation is strictly a matter of morphological agreement: singular antecedents want singular pronouns and plural antecedents want plural pronouns. (More on the role of morphological agreement in the next section.) Thus, in analogy with (2.91) and (2.92) we have the felicitous examples in (2.95).

(2.95)a. In the fall hedgehogs eat a lot. Then they find themselves a hole and sleep through the winter.

b. Lions eat meat. If you try to feed them a carrot, they'll rip off your hand.

We already noted (in connection with (2.89.a,b)) that singular pronouns are no good as anaphors to plural antecedents. this is confirmed by the examples in (2.95): if you replace *they/them* by *it/him/he* the result is ungrammatical (or extremely marked). However, hesst parallels between (2.91) and (2.92) on the one hand and (2.95) on the other make it all the more remarkable that substitution of *they/them* for *it/him/he* in (2.91) and (2.92) does *not* result in ungrammaticality.(Note that this underlines our first and central observation that plural pronouns can pick kind references that are retrievable from an utterance or text just because of the presence of noun that names them.[33])

There is much about kinds, reference to kinds and genericity that is still a topic of debate among linguists, and I do not know of anybody who thinks that we have got to the bottom of these phenomena. The problems that are still being discussed range from the ontological question what kinds are and what the exact nature is of their relations to their extensions and their individual instantiations, to the more linguistically oriented question what the logical forms are of the various sentences in which reference to kinds appears to play a role. We have done very little here to clarify these fundamental questions. All that we were concerned about is to get a clear (and hopefully complete) picture of the different ways in which antecedents for plural pronouns can be constructed in cases where no suitable antecedent is already present. For a quite comprehensive record of what was known about genericity and kind reference until about two decades ago see (Carlson & Pelletier 1995*a*).

A striking feature of the mechanisms that are available for the construction of antecedents for plural pronouns is the way they are limited. From a logical point of view Summation and Abstraction are comparatively simple

[33]One implication of the observations we have made is that presumably the pronouns in (2.95) allow for two different construals, (i) as anaphors to the plural DPs *hedgehogs* and *lions* which can pick up the drefs introduced by those DPs because of number agreement (and which require some kind of modal subordination analysis, just as the singular cases in (2.91) and (2.92)), and (ii) as kind anaphors, in which case it is the common nouns *hedgehog* and *lion* that serve as their antecedents and in which no modal subordination is involved.

operations, and so is the introduction of kind drefs (if we ignore all the intricacies relating to the sorts of entities they represent, but these, it seems to me, shouldn't be seen as *logical* intricacies). But there is another logical operation which on the face of it doesn't seem much more complicated than the ones we have discussed and which does *not* appear to be available for the construction of antecedents for plural pronouns. This is the operation of set complementation, or set subtraction, which when applied to two sets A and B results in the set A $\setminus$ B (the set consisting of those members of A that are not in B). When discussing Partee's ball examples we saw that the antecedent we would like to have for *they* in (2.74.b) is the set of two balls that we get by subtracting the set of eight balls from the set of ten balls. But that isn't a possible way of interpreting *they.* At the time when we first confronted this example, it looked just like a replica of the example in (2.52.b), with just a change of the numbers – 2 instead of 1 and 8 instead of 9. But now that we have seen how many more options there are for the interpretation of plural pronouns than there are for singular pronouns, the impossibility of interpreting (2.74.b) the way we would like to gains a new significance. It isn't just that no dref for the set of the two missing balls isn't present in the discourse context, it isn't even legitimate to *construct* such a dref from the material present when the need for it arises.

What precisely is this constraint on the repertory of operations that are available for the construction of plural pronoun antecedents, and why should there be such a constraint? The answers to these questions are, I believe, still unclear. An answer to the first question may be that only 'purely positive' operations are admissible, with Summation, Abstraction and Kind Extraction qualifying as 'purely positive' while set subtraction does not. But what is 'purely positive'? I envisage a model-theoretic definition of 'purely positive', but haven't worked pout the details. An answer to the second question – why pronoun interpretations should be subject to such a constraint – may be rooted in the structure of human cognition. On the whole, most humans do not seem very well equipped to process operations like negation and complementation. (I myself am bad at this to the point of causing acute embarrassment to myself and to the members of my trade, which I would like to think to be that of a logician; but perhaps that is the explanation of why wanted to become a logician in the first place, much like people with speech defects have been said to be attracted to linguistics). Perhaps then that operations which involve negation are simply cognitively too costly for the kind of online processing that listeners must engage in when they have to keep abreast of the flow of speech that is coming at them and that they have to digest.

If I am not too far off in my estimate of the current state of our understanding of the human cognitive system, this can only be speculation. But whatever the answer to this question may ultimately turn out to be, and whatever the exact demarcation may be between what is permitted in the construction of antecedents for plural pronouns and what is not, the fact remains that these constructions are restricted in the way we have seen; and that by itself is a truly remarkable fact, with an important methodological moral for linguistics. To see this, note: (i) that the restriction we are talking about is specific to *plural pronouns* – there is no such constraint on the interpretation of other anaphoric expressions, such as for instance definite descriptions, whether plural or singular; see (ref100.8) below, where such examples are put side by side with corresponding pronoun examples – and (ii) the constraint manifests itself at least as much inter-sententially as intra-sententially.

(2.96)a. Nine of the ten balls are in the bag. The missing ball/the last ball/the remaining ball is under the sofa.

b. Nine of the ten balls are in the bag. It is under the sofa.

c. Eight of the ten balls are in the bag. The missing balls/the last two balls/the remaining balls are under the sofa.

d. Eight of the ten balls are in the bag. They are under the sofa.

The implications of (i) are, I think, clear. Whatever the deeper cognitive reasons may be for the restrictions on the construction of antecedents for plural pronouns, the linguistic system will have to contain the information that this is a restriction that applies to one particular kind of expression (the plural pronouns). Whether this information is located in the syntax of the language – in its interface with the semantics or in the lexicon – may depend on further theoretical assumptions the linguist may feel compelled to make. But either way it must be part of the *grammar* of the language, at least when we take the notion of grammar in a sense broad enough to encompass at least the just mentioned components.

When these implications are combined with (ii), the conclusion emerges is that there are aspects of the grammar of English that do not just concern its individual sentences but also the ways in which two or more sentences can be connected in form and interpretation. Thus so long as the grammar of the language is assumed to include its (syntax and) syntax-semantics interface, it cannot be confined to the analysis of its individual sentences. Today, this finding has probably lost much of the edge that it had when it first emerged from the DRT based analyses like the ones presented in this and the preceding sections. (For one thing this is because work on languages very different from English is revealing that in those languages sentence boundaries – to the extent that they can be made out at all – play a quite different role than they do in Indo-European languages.) But even now it is a point that we do well to keep firmly in mind when evaluating the often fundamentally different proposals for theories of linguistic structure with their often strikingly different architectures.

2.9.3 When Plural Morphology and Plural Semantics come apart

We conclude this survey of aspects of the use of plurals in English with a few examples which show that plural morphology does not always pair up with semantic plurality (in the sense that the semantic value must be a non-atomic individual). For a first set of examples consider (2.97).

(2.97)

a. All my friends own a car.

b. All my friends own cars.

c. Every friend of mine owns a car.

d. Every friend of mine owns cars.

e. All my friends own two/at least two/several cars.

f. Every friend of mine owns two/at least two/several cars.

At first blush (2.97.a) and (2.97.b) seem to have the same truth conditions: I can say both of them truly when each of my friends owns a car and also, it seems, in a situation in which some or even all of my friends have two or more cars. (Both sentences can be followed consistently with 'In fact, they each have just one car.' as well as with "In fact, they all have more than one car.') This is already a noteworthy fact in itself. Apparently the plural DP *cars* does not entail in this case that its denotation is non-atomic. The force of this conclusion is particularly compelling when we think of the semantic representation of (2.97.a). In that representation the DP *cars* will have to be represented by a single dref (in the nuclear scope DRS of the duplex condition introduced by the quantifying DP *all my friends*) and this dref – let it be η – must be able to take atomic entities (i.e. single cars) as values.

Someone intent upon saving the assumption that plural morphology in DPs must reflect semantic plurality in some way might be tempted by the hypothesis that we get a non-atomic value when we sum over all the values that the dref η representing *cars* must take on for the different friends referred to in the subject DP. (In other words, summing over the merge of restrictor and nuclear scope DRSs of the duplex condition with the Summation operator binding η.) But that cannot be the right explanation on its own, as it would predict that the morphological plurality of *cars* in (2.97.d) could also be justified in terms of the non-atomicity of the totality of the cars owned by my friends. But (2.97.d) clearly does not have an interpretation that neutral about the number of cars that each of my friends must have. The sentence can only be understood as saying that each friend has more than one car.[34]

The difference between the semantic contributions made by *cars* in (2.97.a) and (2.97.d) is accounted for by the assumption that the morphological plurals of some DPs are justified by a relation in which those DPs stand to other plural DPs. In particular, so the assumption goes, this relation holds between the subject argument and other arguments of a more-place verb. The relation between subject and direct object of the verb *own*, as in (2.97.a), is one instance of this: if the subject DP has plural morphology, then the direct object DP may have plural morphology without this carrying the non-atomicity implications that we have been assuming up to now. But when the subject DP has singular morphology, then plural morphology on the direct object DP cannot be accounted for as a case of agreement, which means that the plural must signify non-atomicity.

[34]In fact the sentence seems somewhat awkward, as if one felt under some pressure to come out on either side of the divide between one and more: if one does want to say that each friend owns at least two cars, then it is much more natural to use a phrase that makes this explicit, such as *more than one car*, *two or more more cars*, *several cars*. If one wants to convey that each friend owns a car, while keeping the question whether she owns just one car or more than one out of it, then *a car* seems appropriate; and if one wants to make explicit that the number of cars owned could be one or more than one, then a natural DP choice would be *one or more cars* or *at least one car*. But in any case the semantic 'number neutrality' of *cars* in (2.97.a) is not shared by *cars* in (2.97.d).

In the light of this hypothesis it is also of interest to compare (2.97.a) and (2.97.d) with (2.97.e) and (2.97.f). In these last two sentences the direct object DP is also plural, but it contains a cardinal-like expression that requires plurality for semantic reasons: the cardinal-like expression is such that the NumP of which it is part can only be satisfied by non-atomic entities. So here the semantic neutrality of number agreement, if it plays any part at all, is overwritten by the constraints that these cardinal-like expressions impose. For this reason the version with plural subject ((2.97.e)) and that with singular subject ((2.97.f)) express the same truth conditions.

The question that these considerations do not yet settle is whether the dref introduced by a dependent plural should be neutral (a Greek letter without atomicity constraints, in our notation) or be treated as having singular semantics (in our notation: a dref that can be given in the form of a lower case Roman letter). The question has not yet been settled because even when *cars* in (2.97.a) is represented by a 'singular' dref, this does not exclude that some of the friends the speaker is talking about own more than one car. In fact, this seems the more plausible story about (2.97.b), and it is the one that our earlier assumptions about the treatment of singular indefinites like *a car* commit us to: The DP is represented by a singular dref such as, say, y. But the truth conditions that this confers upon (2.97.b) – for each friend x there is a car y such that x owns y – is compatible with there being more cars that some or all of the friends own. If we were to assume that the DP *cars* in (2.97.a) makes this same contribution, then we would still get the same truth conditions, those that (2.97.a) and (2.97.b) seem to share.

But here is an example which shows that at least some dependent plurals must be represented by atomicity-neutral drefs.

(2.98)a. All students brought books that would keep them fully occupied during the next two weeks.

b. All students brought a book that would keep them fully occupied during the next two weeks.

Consider the following scenario. Some of the students brought one book. But those books were thick and difficult and the students who brought them had their hands full getting just through their one book in the two weeks they had. Other students brought several books, more easily digestible and shorter, and if they had brought fewer (and certainly if they had brought just one of them), they would have run out of things to read before the two weeks were over. In this scenario (2.98.a) would be true, whereas (2.98.b) would not be. In the light of all we have already said the reason for this difference is easy to come by: The dref introduced by *a book* in (2.98.b) introduces a singular dref and this imposes on the representation of the sentence truth conditions that are not compatible with what (2.98.b) says. On the other hand, *books* in (2.98.a) introduces an atomicity-neutral dref and this dref can take for the different students the values – some of them atoms, some of them not – that are needed to verify the representation in a model which reflects the scenario as we have outlined it.

Conclusion: In (2.98.a) the dependent plural *books must* be represented by an atomicity-neutral dref. This does not prove conclusively that this must also be so for other dependent plurals, such as, for instance, the one in (2.97.a), but it certainly strengthens the case for this assumption. And since I do not know of any examples for which this assumption leads to counterintuitive results, this is the general assumption I am prepared to make:

(2.99) Dependent plurals are represented by atomicity-neutral drefs.

How dependent plurals should be semantically represented is one question. Another question is which plural DPs can be interpreted as dependent on which. The examples we have so far considered were of direct object DPs that are dependent on the corresponding subjects. But the subject argument-object argument relation is not the only one that establishes number dependence. (2.100) gives some other examples.

(2.100)

a. All women bought cars that had automatic transmissi ons.

b. All women bought cars with automatic transmissions.

c. All women bought cars that had an automatic transmission.

d. All women bought cars with an automatic transmission.

e. All women bought a car that had automatic transmissions.

f. All women bought a car with automatic transmissions.

g. All women bought a car that had an automatic transmission.

h. All women bought a car with an automatic transmission.

i. All women bought cars that friends had recommended to them.

j. All women bought a car that friends had recommended to them.

In (2.100.a) we see two dependence relations, between the subject and the direct object of the main clause and between the direct object of the main clause and the direct object of the relative clause of which the main clause direct object is the subject. Similarly, there are two dependence relations in (2.100.b), between the subject and the direct object and between the direct object and the DP of its prepositional adjunct. That is why (2.100.a) and (2.100.b) can both be understood as true in a situation in which each woman bought one car with one transmission. (2.100.c) and (2.100.d), in which the 'most deeply embedded' DP is in the singular have closely similar truth conditions – not exactly gthe same, but the truth conditions coincide on the assumption that cars have just one transmission – and that is clearly because these sentences stand in the same relation to (2.100.a) and (2.100.b) as (2.97.a) stands to (2.100.b). What we see in these examples is a kind of chain of dependence relations rather than two independent dependencies on the subject DP. That this is so is shown by (2.100.e) and (2.100.f). Here it is the DP 'in the middle', the direct object DP of the main clause, that is in the singular and the effect of this is that the DP *automatic transmissions* can no longer be interpreted as a dependent plural, with the bizarre implication that the cars that were bought had more than one automatic transmission each. This implication can be undone again by making the transmission DP into a singular as well, as in (2.100.g) and (2.100.h).

(2.100.i) and (2.100.j) show that the dependence relation between a DP and some other DP occurring in a relative clause adjoined to the first does not require that in the relative clause the first DP plays the part of subject. That was the case in (2.100.a), but it is not so in (2.100.i). In (2.100.i) the main clause DP *cars* is the direct object of the relative clause of which the DP *friends* is the subject. Nevertheless *friends* is dependent on *cars*, as can be seen from (2.100.j), where *cars* has been replaced by *a car* with the effect that the sentence now says that each car that was bought was recommended to the buyer by several friends (that is, it excludes purchases of cars recommended by a single friend).

This concludes our excursion into the topic of grammatical and semantic plurality, and with that we come to the end of this introduction to early DRT and the ways in which it differs from classical Montague Grammar and, more particularly, from the modern incarnation of it presented in H&K. Our focus in this excursion has been for the most part on pronouns, which have been the red thread through this entire presentation of DRT; and the reason for paying plural pronouns the attention they have been getting in Section 2 is that their semantic behavior differs in striking ways from that of singular pronouns, and ways that at a first encounter may seem quite surprising. Even as an exploration of plural pronouns our excursion has been far from complete – it has covered only some of the known parts of this territory, and there are undoubtedly parts of the territory that so far haven't been charted at all.[35]

But plural pronouns form only a small part of the many issues that the morphology, syntax and semantics of plurals raise more generally. We have seen glimpses of some of the issues that have to do with plurality but not directly with pronouns – e.g. question having to do with the plural forms of other DPs, and those that concern the predicates – in particular verbs and verb phrases, but also, although not mentioned earlier, prepositions – that take plural DPs as arguments. A substantial amount of work has been done to deal with many of these issues, but none of which will be dealt with in the core part of this course. To my knowledge there is at present no truly satisfactory overview of what we know about plurals if even just one language. (Here I am once again thinking of English, the most extensively studied language of all by both syntacticians and semanticists). A large part of the difficulty, and one of the reasons for the absence of such an overview is that there still is a lack of agreement on fundamental equations such as the ontology of plural reference – sets and their members in the sense of Set Theory or parts and wholes in the sense of Mereology? – and about the logical forms of sentences containing various types of plural noun phrases – which of these sentences involve quantification, and what do they quantify over when quantification is assumed to be involved? For lack of a better way to conclude this section here are a few well-known references from the semantics literature on plurals, that present well-established views on some of the major issues: (Schwarzschild 1996), (Landman 1989*a*), (Landman 2000). The most recent survey that I am aware of is that of Nouwen in the Cambridge Handbook of Semantics, which should appear later this year (Nouwen 2015).

2.10 Top Down, Compositonality, Bottom Up

Up to now DRSs have been built top down, by decomposing sentences according to their syntactic structures and in the process converting them into the logically transparent notation of DRT. In what follows we are going to build DRSs bottom up, starting with conversions of lexical predications down the bottom of the syntactic tree and working our way towards DRSs for complete

[35]For some more information on issues that arise with plural pronouns see (Kamp & Reyle 1993), Ch. 4 and the much more recent and more extensive (Zweig 2008).

sentences by integrating DRS-like semantic representations for the syntactic daughters of a given mother-daughter(s) node configuration into a semantic representation of the mother.

Proceeding in this second way may have a more compositional flavor to it, and it may well be that those who first proposed bottom-up construction methods for DRSs were in part motivated by criticisms of the original top down method of DRT that we have been using here as 'non-compositional' (see e.g. [Asher?], (Zeevat 1989)). In the early days of DRT there were quite a few criticisms in this spirit. Some of them were right, putting their finger on sore spots, but others seemed to reflect either a lack of understanding of how DRT exactly works or, worse, various misunderstandings of what 'compositionality' is, or could be, or should be. Our first task must therefore be to explain what the principle of compositionality, which Top Down DRT was said to violate, actually is. Only then will it be possible to say precisely in what way or ways the Top Down method is in breach of compositionality and address the question whether these breaches ought to be avoided, and, if that is what we want, how.

On the face of it the Bottom Up method to which we are now switching might promise a remedy against these supposed shortcomings of the Top Down method. But the Bottom Up method we will adopt introduces new complexities into the representations and their construction (because of the way it handles the representation and justification of presuppositions, which will be one of its central features). These complexities create new tensions with compositionality in its most popular form; but as a matter of fact, they raise the question how useful that concept can really be.

The criticism that DRT isn't 'compositional' must be distinguished from the objection that its DRSs introduce an undesirable element of 'representationalism' into semantics. The representationalism issue doesn't really belong here, where our concern is to motivate the transition from Top Down to Bottom Up DRS construction, since in both cases the final results of the construction are DRSs, so this criticism applies to either method. But it is important to distinguish between the two criticisms – against DRT's representationalism and against its failure to be compositional. Moreover, the question whether 'representationalism' should or shouldn't be allowed a place within semantics is an absolutely central one, which ought to be given some attention somewhere in a course like this one. And this, if not an optimal one, is as good place as there will anywhere in these notes.

We will proceed as follows. We start with a critical discussion of compositionality and its significance for the semantics of natural languages. We will then look at the question to what extent and in what ways Top Down DRT might be accused of violating compositionality. The next subsection of this section deals with the concept of semantic representation and DRT's 'representationalism'. And the final subsection outlines the Bottom Up approach that we will adopt and, in the light of that outline, has a first stab at the question whether a form of Bottom Up DRT can meet the compositionality criticism, and whether it should.

2.10.1 Compositionality

Perhaps the most obvious wisdom about the way in which languages work is that they are governed by two sets of principles, (a) 'syntactic' principles that allow the building of complex expressions from a given finite stock of basic expressions (to which we will refer as the *vocabulary* of the given language) and (b) 'semantic' principles that assign to each correctly built complex expression a 'semantics'. (What the 'semantics' of an expression is supposed to consist in varies between theories. We will come to that shortly.) The construction of a complex expression according to the syntactic principles is piece- and step-wise, by combining expressions already built into larger expressions; and it leaves an imprint (invisible to the naked eye) on the expression that has been built, in the form of a record of how the expression has been obtained from the vocabulary items contained in it. (This record is usually given in the form of a bracketed string or of a tree. The syntactic trees we have been using in the first part of the Notes that is now behind us are instances of this, although in our review a grammar to generate them was never explicitly spelled out. In what follows we will continue this loose practice.) The semantic principles assign a semantics to all well-formed expressions by determining how the semantics of any mother node in the syntactic tree of any well-formed expression is determined by the semantics of its daughters.

Formal languages such as predicate logic and the typed lambda-calculus are compositional in the following sense: There is a recursive definition of their well-formed expressions, and in particular of their well-formed formulas. And given any model for a language of predicate logic or the lambda-calculus there is a recursive definition of satisfaction (of predicate logic formulas in the model) or of semantic values (those denoted by the terms of the lambda-calculus in the model). The clauses of the first definition identify the syntactic principles of the given formal system, and the clauses of the second definition the semantic principles.

It is nowadays widely assumed that natural languages can be described in an analogous way; in fact, many would see this as self-evident to the point of triviality – as a point that you shouldn't make in front of your colleagues if you do not want to come across as painfully naive. But about half a century ago, when formal semantics of natural language as we know it today was first put forward (through the work of Montague) as a genuine possibility, this wasn't seen as self-evident at all. That substantial fragments – substantial in the sense of having expressive power that included that of first order predicate logic – could be described in this way was generally felt as a revelation. Almost everybody at that time thought of natural languages as much too unruly and idiosyncratic to allow for such rigorous streamlined descriptions.

In the time that has passed since the onset of Montague Grammar (and perhaps even earlier) it has become increasingly clear, and obvious, to a growing number of scientific communities engaged in the study of language that there must be some sense in which human languages are 'compositional': The repertoire of expressions we can form in the languages we speak is open-ended. It is in fact fairly rare for us to use the same sentence more than once, and doing so in a single conversation or short text is usually felt to be awkward (unless there is a special point to it. like verbatim repetition for emphasis). Moreover, the sentences of which we make repeated use tend to be formulaic, crafted for particular conventionalized purposes. For the most part, the sentences we produce are ones we never said or wrote before, and the same is true for the sentences we hear or read. On the whole we manage to deal with all this novelty remarkably well, and how would that be possible if our knowledge of language wasn't based on a finite set of principles that relate information content to linguistic form? As listeners and readers we need the right tools to take the novel sentences that reach us apart into their components in order to unveil their message, and something like the reverse must be going on in sentence production, when we look for the words that will accurately express what we want to say. These intuitive reasons for the assumption that knowledge of language must somehow take the form of knowing the principles according to which well-formed expressions can be built from lexical pieces and other principles that endow these expressions with meaning in some systematic, tractable way are reinforced by reflections on how languages can be learned. Language acquisition must happen in a finite amount of time and the result must fit into a finite brain. So only a finite amount of information about the language can be acquired and retained, and this information must put the language learner in a position to deal, as producer and as recipient with an open-ended set of novel sentences.

To what extent work in Formal Semantics like Montague's has helped to appreciate these points I do not know. All that can be said – I do believe that this much is true – is that these arguments from the use and acquisition of language came to be taken for granted as Montague Grammar was becoming established within the linguistic and philosophical communities. But note well that the arguments do not tell us very much about the finite stock of principles that determine what the form and corresponding meaning of the expressions of a given language can be. In fact, as I have stated the argument it does not tell us anything at all about this. All it allows us to conclude is that there must be SOME way in which syntactic and semantic information is finitely encoded and effectively accessible to the language user when she engages in linguistic production or interpretation. There is no reason to assume – not, that is, until much more is known about the mental processes involved in production and interpretation – that the knowledge of language is coded in the form of a recursive definition of the syntax together with a matching definition of how the syntactic forms determine their meanings, along the lines known to logicians from the definitions of formal languages. Looking at the matter just from the perspective of interpretation: the interpreter must have some method or algorithm, which embodies his linguistic knowledge or is derived from it, that assigns the incoming sentences their intended meaning. But that seems to be about all that can be inferred from the argument as is. It may seem plausible to us that such a method can be seen as consisting in doing two separable things: (i) assign a syntactic analysis of the incoming string of words and (ii) compute the semantics of the string from this syntactic analysis. But even that doesn't follow from the conclusion that we must be equipped with some method or algorithm for interpretation. That we have come to find it natural to think of language interpretation as involving these two stages may be no more than the reflection of a practice that has firmly established itself among those who sought to transplant their conception of formal logic and metamathematics to the study of natural language.

Speculative as they are, these considerations should not be forgotten when looking at the formal notions of compositionality that have been playing a prominent part in methodological discussions of natural language semantics over the past decades. A theory can fail to be compositional in the precise sense of such a formal notion, but that need not mean that it fails to be compositional in a sense that would do well enough to satisfy the intuitive requirements which the possibilities of acquisition and use impose on human languages.

So much for these intuitive reasons why the semantics of natural languages is 'compositional' in some way. It is time to turn to the more concrete and precise proposals that have been made about what it is that makes a language 'compositional', and the notion that we will concentrate on is that of *strict compositionality.* Strict compositionality is in the first instance a notion applying to *descriptions* of languages or language fragments. In order to qualify a description must, to begin with, be of the form indicated above; that is, it must consist of (i) a recursive definition of the set of well-formed expressions of the language, which assigns to each well-formed expression a syntactic analysis (which, as said, we assume can be given in the form of a syntactic tree), and (ii) a recursive definition that assigns each well-formed expression a 'semantics' on the basis of its syntactic analysis and does this by specifying, in the way alluded to above, how the semantics of mother nodes depends on the semantics of their daughters. And furthermore, these principles have to satisfy certain constraints. Of primary interest are the constraints on the semantic principles, but these can be stated only (or at least more easily) when the syntactic principles are given in a certain form. So it must be our first task to spell out the form of these syntactic principles.

In stating the characterization of a form of syntax that can serve as the basis for a definition of strict compositionality I will proceed in two stages. At this point I will characterize the syntax needed for that definition in a concise and streamlined form, which allows for a simple and straightforward definition of strict compositonality. But it may not be obvious that a syntax satisfying these restrictions will be capable of generating many of the structures that we have been using as inputs to DRS construction in our DRT survey in these Notes. If that were so, then that would void our discussion of the strict compositionality constraint of any interest: the criterion simply couldn't be applied to DRT (or at lest not to the version presented here). As a matter of fact, any worries that the type of syntax defined below is too restrictive for our present purposes can easily be set aside. I will end this subsection with a few words to that effect.

Here then, under (i), is the characterization of the syntax on which our characterization of strict compositonality will be based. The semantic and central part of this characterization follows under (ii).

(i) (The form of Syntax and Vocabulary)

(i.a) (Syntax) There is a finite set CL of 'Category Labels' that the syntactic principles may refer to. That is, the principles may license mother-daughter node configurations only in terms of the Category Labels involved. And that is, in yet other words, a mother-daughter(s) configuration can be part of the syntactic analysis of a syntactically well-formed expression if and only if there is a syntactic principle that licenses it. (So, for present purposes we can think of the syntactic principles as just specifying a set of licensed mother-daugher(s) category combinations, i.e. a set of tuples $<\mathrm{M},\mathrm{D}_1,.., \mathrm{D}_n>$ with $\mathrm{n} \geq 1$.)

(i.b) (vocabulary) The *vocabulary* of the language is a set of pairs $<w, l_w>$, each representing a word w (identified as a string of letters), and a Category Label, indicating the syntactic Category to which w belongs.

I assume that it is clear what it means to be a well-formed expression according to a syntax specified in this way: A string $w_1, ..., w_k$ of words from the Vocabulary is a well-formed expression if it has a syntactic analysis in the form of a tree T, whose non-terminal nodes have labels from CL and whose terminal nodes are words w, each dominated by a single non-terminal node N_w with the same Category Label l that is also part of an entry $<w, l>$ in the vocabulary.

(ii) (Semantics) For each mother-daughters configuration $<\mathrm{M},\mathrm{D}_1,.., \mathrm{D}_n>$ licensed by the syntax there is a semantic principle $\mathrm{Sem}_{<\mathrm{M},\mathrm{D}_1,..,\mathrm{D}_n>}$ which maps the n-tuple consisting of the 'semantics' of each of the daughters onto the 'semantics' of the mother node.

But what are we to understand by the 'semantics' of a terminal or non-terminal node? Here the convictions of semanticists seem to quite radically diverge, at least on the face of things. We can see this divergence when comparing Montague Grammar (and, more particularly, H&K's system, as one version of MG), with DRT (in particular the version we have been reviewing; but on this point all versions of DRT are alike). According to MG the task of a semantic theory is to assign the right *semantic values* to well-formed expressions of the language. More specifically, any model M for the language will come with an ontology, consisting of (a) the objects in the Universe $\mathrm{U_M}$ of individuals in M as well as (b) a range of Domains formed on the basis of $\mathrm{U_M}$, such as subsets of $\mathrm{U_M}$, functions from $\mathrm{U_M}$ to $\mathrm{U_M}$ (i.e. functions from individuals to individuals); and in case M is an intensional model, also functions from the possible worlds of M to truth values (the *propositions according to* M), functions from the possible worlds to individuals of M (the *individual concepts of* M) and so on. In particular, MG usually assumes that the ontology includes Domains for all the types that can be built from the basic types e, t (or, for intensional models, e,t and s, with the restriction that the entities of type s, (the possible worlds), can never be the type of the ranges of the functions of any functional Domain). To specify what the values in M's ontology are for the well-formed expressions of the language does of course require some clear and precise means of specifying those values, and in practice this means that the theory must make use of some kind of strictly regimented or formal language to do this. Montague used his 'Higher Order Intensional Logic' for this purpose (an intensional version of the lambda-calculus of which he believed that it provides the right framework for doing philosophy and conceptual analysis in general, with natural language semantics as one from a long list of actual and possible applications).[36]

The opposing view of what the central task is for a theory dealing with the semantics of a natural language or language fragment is that the theory should provide *logical forms* for the well-formed expressions of the language. These logical forms should belong to a formalism that is syntactically and semantically well-defined, in the same way that a semanticist committed to the specification of semantic values will demand this of his specification formalism. The 'DRS-languages' we have been using in our review of 'Top Down DRT' are an example of this. They are much like languages of predicate logic. Their well-formed expressions – their DRSs, drefs and DRS conditions – can be defined by means of a recursive syntax, and there is a model-theoretic semantics for these expressions, which specifies which embedding functions verify which DRSs and DRS conditions in which models and, derivative from that, which DRSs are true in which models.

For a semanticist who is committed to the assignment of semantic values it is clear what a strictly compositional treatment of a given natural language or language fragment should come to. Given a syntax and vocabulary of the kind described above, the task consists of two parts: (i) specify a semantic value for each of the items in the vocabulary; and (ii) state for each licensed node configuration $<\mathrm{M},\mathrm{D}_1,..,\mathrm{D}_n>$ a function that maps the n-tuple consisting of the semantic values of the daughters onto the semantic value of the mother node. More precisely: the semantic principle $\mathrm{Sem}_{<\mathrm{M},\mathrm{D}_1,..,\mathrm{D}_n>}$ for the

[36]Most of H&K makes do with a simpler, extensional version of this system, but with an additional provision to denote partial functions, as a way of dealing with certain aspects of presupposition. For what I want to say here, exactly what value specification formalism is employed in a semantic theory committed to the specification of semantic values is not really important.

node configuration <M,D_1,.., D_n> must specify a function f from expressions to semantic values such that whenever this node configuration occurs as part of the syntactic analysis of a well-formed expression, and SV(D_1),..., SV(D_n) are the semantic values that have been assigned to the daughters, then the semantic value to be assigned to the mother is f(SV(D_1),..., SV(D_n)).

But how is this function f to be defined? This depends on the means the theory uses for identifying semantic values. If it uses a specification formalism SVF for this purpose, then the specification of f will have to specify a transformation of the specifications in SVF of the semantic values of the n daughters into a specification in SVF of the semantic value of the mother node. So, de facto, this is what the specification of the 'compositional' part of the semantic value specification provided by the theory inevitably comes down to: a set of descriptions of how tuples of SVF specifications are to be transformed into other SFV specifications.

This is quite close to what a semantic theory committed to the specifications of logical forms is committed to as well. Thus, at the level of actual *execution* of their tasks in the form of an explicit description of the semantics of a certain language or fragment, what is delivered by semantic value theories and what is delivered by logical form theories will look quite similar; and in case the SVF adopted by a semantic value theory is the same as the Logical Form Formalism adopted by a logical form semantics, the deliveries by the two theories may look for all practical purposes the same. Of course, the two theories may come with different commitments as regards the interpretation of the formal definitions they offer, even if these definitions look identical. For the semantic value theorist, the role that the particular SVF he has chosen plays in his theory will be a purely instrumental and ultimately inessential one. If the formalism succeeds in providing specifications of all semantic values that he needs, and if it also satisfies the other requirements for a logical specification formalism (a proper recursive syntax for its own well-formed expressions and a corresponding model-theoretic semantics), then it will be able to do its job; but another formalism with the same qualifications would have done just as well, and his theory could in principle be restated using that other formalism.

For the logical form theorist this will in general not be so. True, the logical forms that his theory assigns to well-formed expressions will also, given a model M, assign to those expressions semantic values in M, viz. those that are determined by the logical forms it assigns to those expressions – values that those logical forms determine by virtue of the given LFF's own model-theoretic semantics. But while these values will be important for the logical form semanticist, no less than trey are for the semantic value semanticist, for him the choice of some particular logical form formalism will have an independent significance and importance beyond the semantic values that its expressions determine. Exactly what that independent significance consists in may vary from theorist to theorist, even in relation to the very same formal treatment of the very same natural language fragment. One motivation for choosing a logical form formalism would be the belief that the logical forms it assigns to sentences (and perhaps also those it assigns to expressions of other syntactic categories) reveal something about the ways in which human users of a language L represent the content of those expressions. Another logical form formalism might fail to do that, even though its logical forms would assign the same semantic values to the expressions of L. But, as should be clear to anyone who has seen as much of DRT as we have at this point, this isn't the only reason for preferring one LFF to an other. (Its logical forms are also needed as discourse contexts, independently from any psychological interpretation.) We will return to the question what reasons there could be for preferring one LFF over another (even when the two are equivalent as semantic value formalisms) in subsection 1.12.2, where we will discuss criticisms of the 'representationalist stance' of DRT.

But at this point our concern is with compositionality. First a warning: Whether a syntax-cum-semantics theory is or is not compositional has, in and of itself, nothing to do with whether the way it is formulated suggests traversal of the syntactic trees top-down or bottom up. The recursive definitions of syntax and semantics that are the central components of a strictly compositional theory are entirely neutral between these two ways of reading them: (i) as recipes for computing the semantic values or logical forms of expressions bottom up, starting from the leaves of the tree and working one's way up by successive steps of value or logical form computation, to the semantic value or logical form of the top node of the tree; or (ii) recipes for analyzing a given word string into smaller and smaller syntactic constituents and to then using the resulting tree to arrive at a fully unpacked specification of the semantic value or logical form, by, again, a top down traversal, this time applying the semantic principles to obtain the stepwise unpacking.

So if the version of DRT we have been looking at fails to be strictly compositional, it cannot be *just* because the algorithm for computing logical forms (and therewith, we saw, also computing specifications of semantic values) is formulated in a top-down fashion.[37] But it doesn't need much of a close look to see that the way in which the top down assignment of logical forms (viz. DRSs) actually proceeds in the cases we have been looking at isn't an instance of strict compositionality as it stands. This is because many sentence parts that the syntax identifies as constituents aren't assigned 'logical forms' in any recognizable way. To take just one particularly blatant case, consider the subject DP of a quantifying sentence like (2.101). (Any of the sentences containing quantifying DPs that we have been looking at in our survey could be used to make this point, but (2.101) makes this particularly easy.)

(2.101) Every guest left.

The one and only substantive construction step that our algorithm specifies for this sentence is the one that deals with the subject DP, creating a universal duplex condition with the DRS-condition 'guest'(x)' placed in the Condition Set of the restrictor DRS, x placed in the Universe of the restrictor DRS and the DRS-condition 'left'(x)' in the Condition Set of the nuclear scope DRS (assuming that x is chosen as representing dref for the DP *every guest*). There are two violations of strict compositionality that we can observe here. The first and less serious one of the two is the bit of 'looking down into the internal structure of the DP', that is needed to determine what operation is to be performed: the algorithm has to identify the DP as one that has *every* as determiner; for as we have seen, DPs of other forms require quite different operations to separate them from their predicates. This is (in my view) not particularly problematic even for one who insists upon strict compositionality. The reason is that we can eliminate the need for looking into the internal structure of daughter nodes just by changing the syntax in a manner that preserves its general architecture. All that we have to do is change the category Label set CL so that the labels encode the information about the different types of DPs which decides the choice between the different possible operations that are required for separating DPs from their predicates. Thus instead of a single Category Label 'DP' we will now have a bunch of them, viz. 'DP_{every}', 'DP_{some}, 'DP_{a}', '$\mathrm{DP}_{\text{proper name}}$', '$\mathrm{DP}_{\text{3dsingpronoun}}$' and so on. Likewise the Category Label 'Det' is to be replaced by the Labels 'Det_{every}', 'Det_{some} etc. The vocabulary will have to be adjusted too. For instance, the item <*every*,Det> should now be replaced by <*every*,Det_{every}>; the item <*Pedro*,DP> by <*Pedro*,$\mathrm{DP}_{\text{proper name}}$> (and likewise for every

[37]We haven't actually *defined* this algorithm in our review here, but only shown how it works in individual cases. For an explicit description of Top Down construction algorithms see (Kamp & Reyle 1993).

other proper name of the vocabulary) and a similar modification is needed for 3d person pronouns: <*she*,DP> becomes <*she*,$\text{DP}_{\text{3dsingpronoun}}$> and the same goes for the other third person singular pronoun forms *her*, *he*, *him* and *it*[38] Finally, this revised set CL and adapted vocabulary induces a revision of the syntactic composition principles. For instance, instead of the node configuration <DP,Det,NP>, the principles must now license a series of such configurations: <DP_{every},Det_{every},NP>, <DP_a,Det_a,NP> and so on, as well as configurations like
<$\text{DP}_{\text{proper name}}$>, which can be instantiated directly by a vocabulary item of the specified category, viz. by any proper name of the revised vocabulary.[39]

With the new syntax there won't be any need for looking down into the structure of a DP in order to determine which construction rule should be applied to a node configuration in which it is one of the daughters; this information is now directly reflected in the Category Label itself. Perhaps the new version of the syntax comes across as clumsy and as missing important uniformities that were captured by the original version. But if that is the only objection to it, then there isn't much to the first criticism that DRT fails strict compositionality. this is a problem that can be dealt with easily and at little if any explanatory cost.[40]

This then is the less serious violation that the Top Down DRS construction method can be accused of. But there is also a more serious charge: In the semantics that this construction algorithm provides no clearly recognizable logical form (and therefore no clearly recognizable semantic value) is assigned to the constituent *every guest*. The contributions that this phrase makes to the semantics of the sentence are spread out over the result of applying the

[38] Of course, similar adjustments are also needed for the various types of plural DPs - plural pronouns, plural definite descriptions and so on. But to simplify the presentation of what is in any case a minor point I am leaving plurals out of the picture here.

[39] I am ignoring the more complex DP structures we were led to adopt in order to account for plural as well as singular count noun DPs. It should be obvious how to translate what is said here about differentiating between different Category Labels to the mnore complex concept of noun phrase structure.

[40] In certain types of syntactic theory the 'inflation' of Category Label sets with additional 'features' has been given a systematic treatment. According to such treatments features are added to the syntax for certain subsets of its Category Label set with general rules for refining the Labels in that set through the addition of features, and with general rules for the corresponding adjustment of vocabulary and syntactic composition principles. In that way the transparency and uniformity of the original syntax can be regarded as preserved while the semantics can take advantage of the more differentiated information that the new Labels provide.

construction rule that deals with it. To make the present discussion easier, the result is displayed in (2.102).

(2.102)

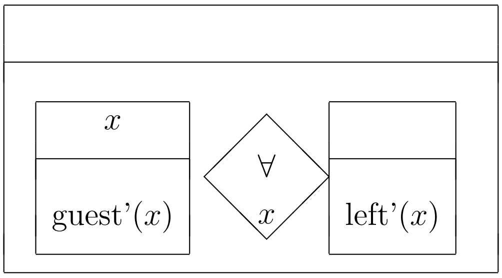

In (2.102) the contributions of *every guest* are spread out over different parts of the DRS: the central quantifier component, the Universe and Condition Set of the restrictor DRS, and through the argument occurrences of the dref x, also the Condition Set of the nuclear scope DRS. This is very different from what we find in MG, and more particularly in its H&K incarnation, where *every guest* gets its own semantic value, and where that value is the result of applying the semantic value of the determiner *every* to the semantic value of the noun *guest*.

For someone who sees strict compositionality as an essential qualification for good semantic theories this ought to be a decisive objection. It is an objection that the bottom up method we will develop in what is to come can be seen to answer in some fashion – though, as I already indicated, the answer may not be to everybody's liking, and that for more than one reason. This is a matter that we cannot say more about until enough of the Bottom Up method will be in place.

The penultimate issue I want to discuss in this subsection is the question how bad it really is if a semantic theory does not meet the demands of strict compositionality. I already noted that there is an informal concept of compositionality that can be derived from general considerations about language use and language acquisition. I do not think there can be any serious disagreement that the formally defined notion of strict compositionality is a way of making this informal notion explicit. But the question is: Is it the only way? And to that may be added a further question: Is it necessarily the best way? What the Top Down DRS construction method we have sampled shows is (a) how a language processor can, by executing a fully systematic procedure, convert the LFs of sentences into logical forms (viz. into DRSs) and thereby assign their logical forms to them. Secondly, the method (b) assigns to each interpreted sentence a semantic value as well (in each model M), viz. the value determined (in M) by the DRS that it assigns to the sentence. And (c) these logical forms can make a systematic contribution to the interpretation of what comes next in the text or conversation. Even without (b) and (c) the algorithm could be seen as meeting the intuitive demands that derive from the general argument from language acquisition and use. And the protests of those who insist that semantics must be 'intentional', i.e. must link the expressions of the language with the subject matter that they are used to describe or refer to, should be silenced by (b). For the semantic values that DRSs determine in models, and that DRS transfer to the sentences (and to multi-sentence discourses and texts) to which the contraction algorithm assigns them, are just that – parts of the actual or some possible world that the sentences text or discourse bits are used to talk about. Moreover, (c) represents an additional benefit, which also has to do with compositionality in its intuitive sense, since it captures some of the systematicity with which donkey pronouns (and other anaphoric devices) make their contributions to the sentences and discourses containing them.

In the light of this, should we consider it a serious problem that the Top Down method leaves sub-sentential constituents without their own logical form and semantic value? That would depend on whether those forms or values are needed for other theoretical purposes than that of determining the logical forms and semantic values of sentences and larger units. Nothing in our presentation of Top Down DRT suggests such a need. That of course doesn't mean that such a need will never arise when the scope of the theory is extended. But as far as I can see, the only sub-sentential constituents for which it is at all likely that their logical forms or semantic values might be independently needed in such an extension of the theory are determiner phrases (as expressions that denote entities or stand for individual concepts) and phrases that have the status of predicates (such as NPs and VPs); and it is quite straightforward to modify our version of Top Down DRT in such a way that it al;so provides logical forms (together with the semantic values those determine). It seems reasonable to conclude from this that from the pre-theoretical perspective which implies that some form of compositionality is a *sine qua non* for semantic theory *strict* compositionality should not be considered compulsory. In fact, repeating a point already made above, from this same pre-theoretical perspective the capacity of DRSs to play the double role of logical forms and discourse contexts can even be considered a 'compositional' advantage over Montague Grammar because of what it has to say about donkey anaphora.

But it is precisely this advantage – or, more accurately, the features of DRT's architecture to which the advantage is owed, viz. the double role played by its DRSs – that has come in for the most emphatic criticisms. We turn to those criticisms in the next two subsections. In conclusion of the present subsection, here are the few words promised earlier about the notion of syntax that we have used as basis for our characterization of strict cxompositionality, so that we won't come away from this discussion with the nagging feeling that it might be irrelevant to a version of DRT like the one we have reviewed.

I want to add that these remarks are marginal to the main arguments of this section. They have what might be described as 'footnote status'. They would make for a very long footnote, however, and one that I wouldn't quite know where to attach. So in the end I decided, somewhat unhappily, to put the remarks in the main text, and right here. But I do it only under the condition that if you do not really want to hear about this, you will skip straight to the next subsection.

If you look at the syntactic structures we have been using as the basis for our DRS constructions (as well as for the few computations of semantic values along the lines of H&K) then you may wonder if a syntax of the sort assumed in our characterization of strict compositionality can generate such structures. If this isn't possible, and it isn't possible either to adjust our specification of syntactic architecture so that it is capable of generating those structures while at the same time retaining its suitability as basis for our characterization of strict compositionality, then the question whether DRT violates strict compositionality can't even be asked; the inapplicability of the notion would already arise at the level of syntax, and there simply wouldn't be any way of putting DRT to the test.

There are two reasons I can see why it might be doubted that a syntactic theory of the kind assumed in our characterization would be capable of generating many of the syntactic structures we have been using in our DRT survey. The first is that the inputs to our DRS constructions have been LFs, syntactic structures obtained from other syntactic structures via movement. Movement is not part in any overt way of the syntax described in this section, and it isn't obvious how something like movement could be simulated with the limited tools that such a syntax makes available. However, even if such a simulation isn't possible, that doesn't disqualify the syntax as provider of the input structures we have been using in the DRSs constructions we have looked at. For it suffices if the syntax can generate these LFs in some way. In particular, it is good enough if the syntax can generate those LFs directly.

Generating the LFs in such a direct way may not make much sense from the perspective of someone interested in the syntactic properties of the natural language in question, in the way that a syntactician in would be; and it may fail to provide us with the means to assign LFs to the sentences whose LFs they are. But none of that matters in this particular context. All we need is that the inputs to DRS construction have the syntactic structure that a syntax of the kind described imposes on the strings it generates.

The second reason has to do with the fact that the structures make use of indexed traces. There are three potential problems here, (i) that the syntactic trees contain traces, (ii) that these traces carry indices and (iii) that the traces are *co*-indexed with other constituents. (i) isn't really a problem at all; we can simply add the trace symbol to the vocabulary with 'DP' (or its various feature-based refinements) as its Category Label or Labels. (ii) is a problem only if we insist that CL must be finite. But there really is no reason to impose this constraint. Assuming that each numeral (from some standard notation for the natural numbers) can serve as 'feature' and that this feature can be used to amplify any Category Label beginning with 'DP' will lead to an infinite set of 'feature-based' Category Labels. But this extension of CL into the realm of the infinite is an extremely conservative one, which neither alters the structure of the individual LFs that can be generated nor the principles that govern its generation. (iii) may seem to be the biggest problem: How can we make sure, in a syntax with the limited resources that we have described, that only those coindexations will occur that can be understood as the result of quantifier raising operations? The answer to this question is: Very likely not, but for present purposes that doesn't matter. As far as coindexation is concerned the syntax may wildly overgenerate. That doesn't matter so long as it generates all the LFs that the DRS construction algorithm may need as inputs. How the parser – the module that assigns sentences their LFs and delivers these to the DRS construction algorithm – deals with its task is of no concern here. All that matters for the applicability of the strict compositionality constraint to our version of DRT is that the LFs on which its construction algorithm operates can be generated by some syntax of the kind presupposed by our notion of strict compositionality.

2.10.2 DRSs as Logical Forms

One criticism leveled at DRT not long after its first versions had appeared in print ((Kamp 1981*b*), (Kamp 1981*a*)) targeted the use it makes of 'representations'. There are a two ways in which this criticism can be taken, and as far as I can tell, both ways were intended, and in some instances by the same critics. The first way is a purely formal one: the role that DRSs play as discourse contexts is one that can be played equally well be much leaner structures, but which nevertheless carry all the information that the semantics needs. (Groenendijk & Stokhof [reference still to be selected]). This criticism was justified in relation to the particular publication it targeted, viz. (Kamp 1981*b*): The phenomena discussed in that paper can be handled with a simpler notion of discourse context than that of a DRS, one which only keeps track of the discourse referents that have so far been introduced, but contains no information about properties of the entities represented by these drefs, or relations between those entities. However, as time went by, it was realized that there are cases of donkey anaphora in which resolution depends on what properties have been explicitly attributed to an individual in the antecedent discourse (as opposed to their actually having those properties). (2.103) is an example that Groenendijk & Stokhof themselves brought up at some later point in time [reference to piece in the Corblin issue of Langages (?)].

(2.103)

a. A man went to see a doctor. The doctor said to the man . . .

b. A man went to see a doctor. The man said to the doctor . . .

It seems plain that in both (2.103.a) and (2.103.b) *the man* is anaphoric to *a man* and *the doctor* to *a doctor*. But what drives these resolutions are the nominal predicates occurring in the antecedent DPs, and not what properties the represented individuals may actually have. In order to take *the man* to refer properly and unequivocally to the person introduced into the discourse by the phrase *a man* we do not have to assume that the doctor is a woman. Nor does the interpretation of *the doctor* as anaphoric to *a doctor* requires the assumption that the man wasn't a doctor too. (Doctors d go and see other doctors when they think they have something wrong with them.)

It is plain that discourse contexts which contain the information needed to resolve the reference of the definite descriptions in examples like (2.103) must not only specify what entities have so far been introduced into the discourse but also what nominal predicates the discourse has applied to them. (Of course, this is not to deny that so long as we make it our task *just* to deal with the phenomena discussed in (Kamp 1981*b*) DRSs can be seen as a sort of overkill.)

2.10.3 Semantic Representations: Mental Representations or just Logical Forms?

There has also been another dimension to the criticisms of DRT's use of DRSs: DRSs import into DRT a psychological element that a semantic theory should be at pains to stay clear of. This criticism brings into play questions about the nature of linguistic meaning and about the brief of semantic theory that go well beyond the aims of this section; and they are questions that do not, as far as I can see, have anything to do with the choice between Top Down and Bottom UP methods for DRS construction. That makes this the wrong place to go into them in any depth, and going into them here to the extent that they deserve is out of the question also because that would take up a considerable amount of space and would seriously divert us from our course. Yet, on the other hand the aspects of DRT that these criticisms touch upon and the views that motivated those aspects of DRT are of such central importance to natural language semantics that something has to be said about them somewhere in these Notes. And besides, some words are needed to explain how these psychology-oriented objections against the use of DRSs differ from those that target their supposed excess of structural detail and information content which we discussed in the last subsection.

The view that semantics should stay clear of any form of psychology can be traced back to two sources. The first is the conviction, still deep-seated at the time when DRT was first presented to a wider audience, that semantics (like logic) has nothing to do with psychology and that bringing psychological considerations to bear on semantic issues leads to no end of confusion. The second, closely related though not identical source is the view, first formulated by Carnap and Morris, that the study of language involves three main parts, syntax, semantics and pragmatics, that the first two of these are concerned with languages as user-independent symbolic systems, which have their syntactic and semantic properties independently from any actual or possible use of them, and that it is only in the third part, pragmatics, that anything having to do with the use of language comes into play. This division of labor also militates against psychological concepts and considerations in semantics – if at all, such considerations should enter the over-all theory only as part of considerations of language use and thus only at the level of pragmatics.

Such a view of human languages is undeniably appealing. Languages can function the way they do only because they are governed by rules that are shared by their users, and the simplest way to explain this fact is to assume that languages have an existence and identity that are independent of those that speak them. It is by acquiring the rules which govern such an independently existing system, and hence by virtue of sharing them with those who have gone through the same acquisition process, that the speakers of a language L can communicate with each other through the use of L: When encoding a thought in words a speaker S will apply the syntactic and semantic rules of L, and by using these same rules (in the converse direction, so to speak) the recipient H of S's words can recover the message that S encoded. If the rules of L weren't shared between S and H, verbal communication would be impossible. it is the independence of these shared rules from the individual users that makes for the reliability with which novel forms can be used to communicate new information to other users (even those about whom nothing more needs to be assumed than that they are also competent speakers of the language).

When Montague did the work that put Formal Semantics on the map this view had the status of something close to dogma. I, a graduate student at the time, was among the crowd of all those who took the view for granted. (That we part of one's 'socialization into the field'.) But in the years that followed my doubts grew. These doubts culminated in the conviction that in the long run psychology cannot be kept out of semantics as radically as the view entails. In my own case it was an exploration of the tenses of the verb that led to this change of heart. There is an anaphoric dimension to tensed, which is reminiscent of donkey anaphora. Both have cross-sentential as well as sentence-internal manifestations and neither can be handled straightforwardly in classical Montague Grammar. But tense anaphora is far more common than donkey type anaphora involving pronouns and other noun phrases. In fact, it is pervasive, especially in narratives and other text and discourse types which include descriptions of complex, multi-event episodes. It was in an attempt to deal with this dimension of the semantics of Tense and Aspect that DRT was originally conceived and developed. At that time others who took a close look at the role of tense in discourse came to similar conclusions: that the best way to explain these phenomena was ti assume that human interpreters build mental models of the episodes described in narrative passages which they extends their interpretation of the passage progresses. The starting point for DRT was the project to develop this idea of mental model building formally, with DRSs playing the part of mental models.

For someone who sees an approach along these lines to the semantics of tense as inevitable the criticism that DRT succumbs to 'representationalism' in that its logical forms are intended to have psychological significance, is just a way of expressing an opposing view of the fundamental nature of language, the relationship between languages and their users and about the general form and purpose of linguistic theory. He will make no efforts to refute the criticism but accept it as a complement.

2.10.4 Bottom Up

To say it once more, from now on we are going to construct DRSs bottom up. As I indicated above, the switch from Top Down to Bottom Up will give us a way of responding to the complaint that Top Down construction algorithms fail the strict compositionality requirement. But to see whether Bottom Up algorithms can meet the objection, we have to first state more explicitly why any one would want a Top Down construction algorithm in the first place.

When the first explicit version of DRT (that of (Kamp 1981*b*)) was given its final form, the main reason for choosing the Top Down method was that it seemed to provide the smoothest treatment of donkey sentences. To make the point fully clear, recall the donkey sentences in (1.44), repeated below, and the construction of the DRSs representing these sentences and in particular the rule applications that dealt with the occurrences of the pronouns *he* and *it*.

(1.44.a) If Pedro owns a donkey he beats it.
(1.44.b) If a farmer owns a donkey he beats it.
(1.44.c) Every farmer who owns a donkey beats it.

The reason why these steps succeeded, giving the pronouns their intended interpretations was that the dref needed as antecedent for the pronoun would already be in place when the step was carried out. This advantage is lost when DRSs are constructed bottom up. But of course, the very fact that the advantage is lost in a Bottom Up construction procedure is, as we noted earlier, an indication that the Top Down algorithm is not reversible. Indeed, when we try to build a semantic representation for a donkey sentence bottom up, then we run into the very problem that the top down method was designed to finesse. Consider for instance the *it* of (1.44.c). If we proceed bottom up, then this pronoun occurrence will have to get a representation before we can get the representation of the VP, because that has to be obtained by combining the representation of the pronoun with the representation of the verb; only when that representation has been obtained, and a representation for the subject phrase *every farmer who owns a donkey* has been constructed as well, will it be possible to combine those two into a representation of the sentence. And only at that point will it be transparent that the dref introduced for *a donkey* is in a position from which it can serve as antecedent for the pronoun. So, if the resolution of the pronoun has to be postponed up to that point, then (a) some special provision has to be made that permits such postponements and (b) however that provision will be made, it doesn't look like we are gaining much headway towards making our construction algorithm strictly compositional.

These difficulties are inherent in the problem that donkey sentences present us with. But they are arguably not unique to donkey pronouns with sentence-internal antecedents. A view that has been gaining support over the past ten or fifteen years is that pronouns, like other definite noun phrases (definite descriptions, demonstratives, proper names) come with *identification presuppositions*, presuppositions to the effect that the interpreter must have a way of 'identifying the referent' of the given phrase.[41] One of the proposals for the What such an identification presupposition precisely amounts to varies considerably with the different types of definite noun phrases. Exactly how it varies is a non-trivial matter that we will look into when we get to the topic of presuppositions in Section 4. But for now only the identification presuppositions for third person pronouns are at issue. And for those the presuppositional requirement is just what we have taken it to be all along: that an antecedent for it can be found in the representation of the discourse context, or – in the case of plural pronouns – constructed from material in the discourse according to certain principles.

One of the general assumptions of the treatment of presuppositions that we will adopt is that presupposition resolution takes place only after a preliminary representation has been constructed for the entire sentence. For certain cases of presupposition resolution this seems the only option. This is because the resolution involves information that is distributed over different parts of the representation and the discourse context and access to the dis-

[41]The classical argument within Formal Semantics for what I am here referring to as the 'identification presuppositions' for pronouns, definite descriptions and other definite noun phrases is (Heim 1982,1988). Heim distinguishes between definite and indefinite noun phrases in terms of *familiarity* and *novelty*: the use of a definite noun phrase conveys to the addressee that he is assumed to be familiar with what the phrase denotes, whereas the most common use of an indefinite noun phrase come with a signal that no familiarity is expected (which can be for any one of a number of different reasons). The account that Heim develops in (Heim 1982,1988), known as *File Change Semantics*, makes the same predictions for the fragment discussed in (Kamp 1981*b*), but the scope of her discussion is wider. In particular it includes definite descriptions as well as pronouns.

tributed pieces of the information is guaranteed only when the preliminary sentence representation is complete. (We will see example of this in the part on presupposition later on in these Notes.) The resolution of donkey pronoun presuppositions is a special (and comparatively simple) case of this phenomenon. For singular pronouns what is needed is always just a single piece of information, consisting of the antecedent dref, but availability of the intended antecedent may also be established only upon completion of the preliminary representation fort the entire sentence.

But what *is* a preliminary representation? It isn't possible to give a precise answer to this question at this point. The answer will become clear when we get to presuppositions, but will have to wait until then. This much, however, can be said right now: a preliminary representation will contain (as it obviously must) explicit representations of the various presuppositions triggered by words and syntactic constructions in the sentence that it represents. Thus the preliminary representation of (1.44.c) will contain as part of it a representation of the identification presupposition for the pronoun *it*; and the preliminary representation for (1.44.a) will contain representations of three identification presuppositions, for the pronoun *it*, for the pronoun *he* and for the proper name *Pedro*. The presence of presupposition representations in preliminary representations makes that such representations look rather different from the DRSs that we have been using as semantic representations up to now. Preliminary representations are in general not DRSs of the forms we have already encountered. They will contain DRSs as parts, but these DRSs will be combined with each other and with other notational material in the representation, in ways that we have not yet encountered.

And that circumstance is decisive for the question whether Bottom Up construction of DRSs for donkey sentences like (1.44.a) can restore strict compositionality. The principal difficulty we face when trying to address this question is that it no longer clear what strict compositonality should come to for a theory that takes presuppositions seriously in the way in which we will do that in the later parts of these Notes. None of the definitions of strict compositionality known to me take presuppositional phenomena into account; and it is far from clear how the notion could be generalized to theories that deal with presuppositional as well as non-presuppositional content.

Perhaps the most that we can demand of a semantic theory that takes presupposition seriously in the way we will is that its construction of preliminary representations should be strictly compositional. But if we want to adopt that as our necessary requirement, then we are facing yet a further question:

Do we want to understand by 'direct compositionality' of the preliminary representations direct compositionality of logical forms? Or do we want direct compositionality of semantic values associated with these forms? If we are content with the former, then preliminary representation constructions we will adopt later on in these Notes will be strictly compositional (simply by virtue of being bottom up and applications of a rule based algorithm). But if the requirement is to be direct compositionality of semantic values, then we have two further problems. First, it is, as we will see, not clear what we should understand by the semantic values of representations that combine presuppositional and non-presuppositional parts. And assuming that a satisfactory answer can be found to this question, we are then faced with the further question whether it will be possible to recast the bottom up construction of preliminary representations as bottom up construction of the corresponding values. For it could be that the combining operations on preliminary representations make use of structural properties of these representations that are lost when we make the move from representations to their semantic values. And when that is so, the principles of representation construction need not be convertible into rules for semantic value composition.

There is more to be said about these different questions, but that requires more concrete information about the form of preliminary representations. Further discussion of these matters will therefore have to wait until we can relate it to actual examples of preliminary representations that show how presuppositions are explicitly represented and how these representations fit within the preliminary representation as a whole. ButI want to conclude this 'preliminary discussion' with a kind of global warning: It is far from clear that passing from our present Top Down algorithm to a Bottom Up algorithm will get us closer to a strictly compositional semantics. But the problem isn't just that the methods we are applying – Top Down or Bottom Up – are in violation of a well defined criterion. Once presuppositional phenomena are taken on board it is no longer obvious how the strict compositionality criterion is to be applied, and that is because it isn't clear what strict compositionality can be in such an altered setting.

Indeed, it was not as an effort to come up with a reply to the charge that Top Down DRT fails to be strictly compositional that we turn to the Bottom Up version that will be developed in what follows. The switch to bottom up processing proves to be a natural one when presupposition is incorporated into the theory. On the one hand the reason for preferring aTop Down algorithm – viz. the way it handles pronominal anaphora – loses its force when pronouns are treated as presupposition triggers, on a par with the presuppositions due to the many different presupposition triggers that are found in English and other languages. And on the other hand, the correct representation and resolution of many other presuppositions strongly suggests (for reasons that we cannot go into at this point) that the construction of preliminary representations must proceed bottom up.

If this is not getting us any closer to strict compositionality, and perhaps even farther away from it than we were, then we shall just have to take that in our stride. Honni soit qui mal y pense.

PART II

Chapter 3

Tense and Aspect I

3.1 Some examples and their challenges

One of the odd things about Formal Semantics of natural language as it has developed over the fifty years since its beginnings has been its attitude towards Tense and Aspect. It isn't that there has been no work on this aspect of meaning in natural languages. In fact, the Formal Semantics literature on Tense and Aspect is by now voluminous and in parts it is remarkably sophisticated. But on the whole little attention is paid to Tense and Aspect in introductory texts. (H&K is one example of this, but other formal introductions to natural language semantics are much the same.)

This is odd because tense is something that you find in almost every well-formed sentence of English (the language that most textbooks focus on, for practical reasons if nothing else). There are very few exceptions to this if any. Candidates are imperatives and perhaps certain forms of ellipsis; but even for such sentences it is debatable whether they are really tenseless. Much the same is true for other languages with a well-developed tense morphology.

What follows in this second part of the notes is an attempt to reverse this trend. In the new start we are going to make, in which the Top Down method for constructing semantic representations is replaced by a method for Bottom Up construction, Tense and Aspect will be the first phenomena on which we are going to focus.

I stated the obvious when I noted that pretty much every English sentence is marked for tense: it contains at least one verb that bears a finite tense form; when the sentence is complex, with one or more subordinate clauses in addition to its main clause, or with two or more main clauses connected by words like *and*, *but*, *or*, there will be several verbs, and each of these will come with its own finite or non-finite tense marker. The temporal structure of such complex sentences and the contributions that are made to it by the tense markings on the different verbs can be quite complicated and there are many unsolved puzzles in this domain. So it seems natural to start our investigation with sentences that consist of a single main clause. (I will refer to simple main clause sentences as 'simple sentences'.)

This is what we will do when we start with our development of DRS construction in the next section. But in this section we will look at more complicated linguistic structures than simple sentence, either seuqnecnes of several simple sentences that are intended as (tiny bits of) cohesive discourse or complex sentences with various kinds of subordinate clauses. It is in such more complex linguistic data that tenses unfold some of their more intriguing and challenging properties. I believe it is good to get a taste of this before we embark on the formal enterprise of laying out the derails of Bottom Up DRS construction, so that we have an ida what is to be achieved in the end. For there will be quite a distance to cover, and for quite a long stretch of that distance our formal project may not seem all that interesting. Si it may help to keep in mind during that extended journey that we eill eventually have something precise to say about the phenomena that we will peruse in this section, and so that something of mor interest will await us upon arrival.

But first, then, the perusal of some of the features of tenses (and certain other temporal expressions) that make the development of a formal theory of Tense and Aspect a hard, but also a fun thing got try.

The tense of a simple sentence will be either a past tense, a present tense or a future tense and in first approximation these forms seem to indicate that the sentence is about the past, the present or the future, respectively. To a large extent this intuition is correct, but there are exceptions. And more importantly, even where it is correct, the assessment of a sentence as being about the past, the present or the future can be no more than part of the full story, since in English (and likewise in many other languages) there isn't just one past tense form, one present tense form and one future tense form, but several; and the choice between, say, one past tense form and another usually makes a semantic difference. (It would be odd if it didn't.) The difference in meaning that is conveyed by different past tense forms is a particularly striking and much studied topic in work on Tense and Aspect. And it is one of the many issues in this domain that can be best appreciated and properly studied when we take account of the role that the tenses of individual sentences play in the context of a larger text or discourse.

Here are some examples to illustrate this point.

(3.1)

a. John proved the theorem in twenty lines. Mary proved it in ten lines.

b. John proved the theorem in twenty lines. Mary had proved it in ten lines.

(3.1.a), with the verbs of both sentences in the simple past tense, does not tell us anything about the temporal relation between the events described in the first and the second sentence. But in (3.1.b), where the second sentence is in the past perfect, the temporal relation seems clear: Mary produced her 10 line proof before John found his proof of 20 lines.[1]

A different phenomenon is demonstrated in (3.2), which shows two possible continuations of the first sentence, (i) by the simple past sentence 'She smiled' and (ii) by the past progressive sentence 'She was smiling'.

(3.2)

When Alan opened his eyes he saw his wife who was standing by his bedside.

(i) She smiled.

(ii) She was smiling.

The follow-up with (i) suggests that Alan's wife smiled at him after he opened his eyes and saw her, presumably as a reaction to that. In contrast, the past progressive continuation in (ii) seems to say that Alan's wife's smiling was

[1]The role that the past perfect plays in (3.1.b) was first recognized by Reichenbach in (Reichenbach 1947). Reichenbach noted that in an example like this one the past perfect locates the described event in the past of a time that is itself in the past of the utterance time, and thus involves two temporal relations, that between the event and the intermediate time (to which he refers as the 'Reference time') and between the Reference time and the utterance time. Reichenbach took this observation as the point of departure for a general analysis of all the tense forms of English (and implicitly of tense more generally). His account of the tenses, although very brief, has been extremely influential in the work on Tense and Aspect of the past seventy years.

already going on at the time when he opened his eyes; the first thing he saw was a smiling wife.

(3.3) can be seen as three variants of the same general pattern, each consisting of two event-describing sentences in the simple past tense.

(3.3) ((Webber 1988))

a. Fred went to Rosie for dinner. He came home in a state of euphoria.

b. Fred went to Rosie for dinner. He put on clean trousers and his nicest shirt.

c. Fred went to Rosie for dinner. He bought flowers on the way.

Intuitions as to the temporal relations between the events described by the first and the second sentence in each of these three examples seem pretty clear. It is hard to read (3.3.a) in any other way than as a sequence of two events of which the second came after the first. The impression left by (3.3.b) is that the event of the second sentence, Fred putting on clean trousers and his nicest shirt, preceded his going to Rosie for dinner. And a natural interpretation of (3.3.c) locates the event of the second sentence, the purchase of the flowers, within the event of Fred making his way to Rosie's. These judgments are consistent with the impression left by (3.1.a) that from the mere fact that two successive sentences are both in the simple past nothing can be inferred about the temporal order of the events they describe.

(3.4) amplifies the observation we made in connection with (3.1.b). The event described by the second clause, 'he had got up at six thirty', with its verb in the past perfect, is understood as having taken place before that described in the first, simple past sentence. And that is equally true of the next three past perfect clauses. Together the series of four past perfect sentences tell a multi-event episode all of which is located before the event of Bill's arrival. Such sequences of past perfect sentences, which are part of a longer narrative text or discourse in which they are flanked at the beginning, and sometimes also at the end, by simple past clauses, are sometimes referred to as *extended flashbacks*. In the case of (3.4) it is natural to assume certain temporal relations among the events described by the clauses of the extended flashback: each next clause of the flashback describes a new event that follows upon the event from the clause immediately before it. In general, the order of the flashback events that are described in a series of past perfect sentences will heavily depend on context, much like we saw for the pairs of simple past sentences in (3.3). But the point about (3.4) to be stressed here is this: Whatever the temporal relations between the flashback events, it is because of the past perfect of the flashback clauses that all these events are interpreted as having occurred in the past of the event introduced by the (3.4)'s opening sentence, the simple past sentence 'Bill arrived at noon'.

(3.4) (Kamp & Reyle, 1993)

> Bill arrived at noon. He had got up at six thirty, had cooked himself a full breakfast, and had washed up after finishing it. He had left the house in time to catch the 7.54 train at the central station.

(3.5) illustrates a different phenomenon. In each of (3.5.a-c) the order of the events is clear and it is the same: the leaving event comes after the arriving event. But the ways in which this information is conveyed in each of the three sentence sequences (3.5.a,b,c) are different. One difference has to do with the perspective point from which the narrated events are viewed: In (3.5.a) this is a point in time after both events have taken place, so that both can (and should) be reported as past events. Moreover, this perspective point can be identified with the time at which the discourse is uttered. In this last respect (3.5.a) is like (3.5.c), but the difference is that in the case of (3.5.c) the perspective point/utterance time is situated between the two events rather than in the future of both, with the effect that while the first has to be reported in the past tense, the second must be described with the help of a future tense.

(3.5.b) differs from both (3.5.a) and (3.5.c) in that perspective point and utterance time no longer coincide. The arrival event is narrated as a past event, which means that the utterance time must be after that event, just as for the other two discourses. But the tense form of the second sentence in (3.5.b), the so-called 'future of the past', implies that the leave event is seen from a vantage point from which it still is future. Presumably this vantage point coincides with the time of Henry's arrival. Being in the future of a past vantage point is in principle compatible with any temporal relation to the utterance time, before it, after it or simultaneous with it. And indeed, (3.5.b) can be used to describe any of these possibilities. First, the sentenced pair could be uttered some time – several weeks, say – after the two events, arrival and departure, took place. but it could also be used, though arguably less naturally, on the Sunday of the departure or on the Saturday immediately before it. Just on its own (3.5.b) would probably come across as rather strange to most recipients. But it would much improve when followed by something like: 'But that is today/tomorrow. So he may still be here.'[2]

(3.5)

a. Henry arrived on Wednesday. He left again on Sunday.

b. Henry arrived on Wednesday. He would leave again on Sunday.

c. Henry arrived on Wednesday. He will leave again on Sunday.

A further point about the examples in (3.5) is that in each of them the inference that the arriving preceded the leaving is justified by the occurrence of *again* in the second sentence, irrespective of the tense forms in the respective second sentences. This is because the only justifiable interpretation of *again* in these discourses is the 'restitutive' one, according to which the event described by the verb phrase of which it is part can be understood as reversing the effect of some previous discourse-salient event, 'restituting' the state of affairs that obtained before that other event occurred and which that event then put an end to. (Here the 'restituted' state is that of Henry being in some place other than the one where he arrived, as described in the first sentence.) For our exploration of temporal reference examples containing *again* will prove particularly useful because of the crisp intuitions speakers have about the effects it has on temporal interpretation.

The examples above all show a certain rough similarity with 'donkey discourses', in which a pronoun in one sentence is anaphoric to an indefinite in another sentence. In these examples it is the temporal location of one event, described in a later sentence, that is determined in relation to that of some other event referred to in an earlier sentence. Note well, however,

[2] Interpretations of discourses like that in (3.5.b), which locate the event e_2 described by the second sentence in the future of the utterance time, seem to be dispreferred. The following consideration may be an explanation of this, or part of such an explanation: when the speaker is in a situation where she has already mentioned the occurrence of one event e_1 as located in the past of the utterance time and then wants to present an event e_2 as located in the future of e_1, then there are two forms available to her, the simple future tense with *will* and the future of the past with *would.* If the speaker knows that e_2 is in the future of the utterance time, then the use of the simple future is the more natural choice: it locates e_2 unequivocally in the future of the utterance time, whereas the future of the past would only carry the information that e_2 is in the future of e_1 and thus would be less informative. Given this preference, the recipient of an utterance in which the first sentence in the simple past tense and the second in the future of the past and who expects that the speaker knows whether e_2 is after the utterance time or not, will infer that e_2 is not after the utterance time; for had it been, then the speaker would have known that and would have used the simple future tense.

that while there is some resemblance with pronoun anaphora, the instances of anaphora we have just seen aren't quite the same thing. The anaphoric relation between a pronoun and its antecedent is always one of 'coreference', whereas that between 'anaphorically connected' events involves temporal relations such a temporal precedence, succession, overlap and inclusion (rather than just coincidence, the temporal counterpart of coreference).

Besides such 'temporal donkey discourses' temporal reference also has its instances of sentence-internal donkey pronoun-like effects. And these are not only found in the same grammatical structures where we find sentence-internal donkey anaphora, such as conditionals and quantified sentences – an example of which is given in (3.6) – but also in indirect discourse and attitude reports.

(3.6) Whenever John called, Mary wasn't at home.

A typical example of temporal anaphora in indirect discourse is (3.7.a). Here the verb *felt* of the embedded clause 'that she felt sick' bears the same tense form – the simple past – as the matrix verb *said*. To a good many English speakers (3.7.a) strongly suggests that the time at which Mary said that she felt sick coincided with the time at which she said it. (Her own words at the time would have been: 'I feel sick'.)

(3.7)

a. Mary said that she felt sick.

b. Mary said she ate an apple.

c. Fred and Mary told us of the horrible scene they had watched when coming out of the train station. Mary said she felt sick.

According to some accounts of the English tense system (Stowell, Ogihara) the connection between the tensed verbs *said* and *felt* is different from the temporal relations between events that can be found in our previous examples. These theories see the past tense on *felt* in (3.7.a) as a case of *grammatical congruence*, as if the tense was copied over from the matrix verb *said* to the verb *felt* of its complement clause. Thus the past tense of [*felt* is not making an independent contribution to interpretation, by locating its event in the past of the utterance time, but passes the time of the event of the matrix verb on to the state of its own clause.

But this isn't the only way in which the past tense of *felt* can be construed. For it is also possible to understand the complement clause of (3.7.a) as describing a state that held before the reported saying time. This is the preferred interpretation of (3.7.a) when it occurs as last sentence of the discourse in (3.7.c). (Note that interpreting the state of Mary feeling sick as holding at some time before the saying time is made possible here because the antecedent discourse offers another time – that when Fred and Mary came out of the station – as a possible temporal anchor for the state description *she felt sick*.) Furthermore, in (3.7.b), where the complement clause describes an event, only a non-simultaneous reading seems possible, according to which the eating preceded the saying.[3]

Why the state described by a past tense complement clause can be interpreted as simultaneous with the matrix event while this is not possible when the complement clause describes an event will be explained later. For now we just note that this is one of many respects in which state descriptions and event descriptions behave differently. In what follows the distinction between events and states will be all-important.

Nevertheless we will often want to refer to the state or event that is described by a sentence or clause while leaving it open whether that is an event or a state. In such cases we will make use of the term *eventuality*. *eventuality* is simply a shorthand for 'state or event'. This use of the word *eventuality* was introduced in the eighties by the linguist Emmon Bach. You may feel that it isn't a very good word for this particular meaning. But no one seems to have ever come up with a better word, so this one has stuck, and these days it is used by pretty much everyone in the tense-and-aspect community.

The phenomenon of tense congruence between matrix verb and complement clause is a property that English shares with some languages that have well-developed tense systems, though not with all. It is known as *Sequence of Tense* (or *Consecutio Temporum*, if you want to show off with a fancy term).

[3]One of you (Luke Kundle-Pinette) reported a preference for the non-simultaneous reading even for (3.7.a) when it is offered without context, adding that a better way of expressing the simultaneous reading would be to use the past progressive ('Mary said she was feeling sick'). This judgement may have to do with an understanding of the verb *feel* as an event verb. In that case the simultaneous reading in (3.7.a) would be ruled out for the same reasons as in (3.7.b). Using the progressive instead of the non-progressive form would then restore the possibility of a simultaneous reading, just as it does in the case of *eat*. ('Mary said she was eating an apple' can have a simultaneous reading just as 'Mary said she was feeling sick' can.)

In languages without Sequence of Tense, simultaneity between the eventuality of the complement clause of a past tense matrix clause and the eventuality described by the matrix verb is typically expressed by means of the present tense (and not by the past tense, which in those languages places the complement eventuality unequivocally in the past of the eventuality described by the matrix verb). In a language like English, interpreting the eventuality of a past tense complement clause as preceding the matrix eventuality is tantamount to not construing the tense of the embedded verb as 'congruent' with the tense of the matrix verb in the sense of 'congruence' described above. Such interpretations treat the embedded tense as a tense that makes its own semantic contribution by placing the eventuality described by its clause in the past of the utterance time. Further factors are then responsible for the temporal location of the embedded eventuality w.r.t. the other eventualities around, and in particular for the fact that the embedded eventuality is located before and not during or after the matrix eventuality.

Indirect discourse also includes examples that echo what we noted about the use of simple future and future of the past in (3.5). Suppose that now (some time in the year 2015) is the utterance time for the two sentences in (3.8.a) and likewise for the sentences in (3.8.b). Then the two sentences in (3.8.a) – the one with *in the year 3000* and the one with *in the year 2000* – can both be used; but of the sentences in (3.8.b) only the second one is acceptable, in which the temporal adverbial is *in the year 3000*. The reason for that should be clear at this point: the simple future tense in (3.8.b) carries the implication that the described eventuality is in the future of the utterance time, and that is true for the year 3000 but not for the year 2000. In (3.8.a), where the corresponding tense form is the future of the past, there is no commitment with regard to the temporal relationship between the utterance time and the described eventuality; all that matters is that the eventuality is in the future of the reference time. In the case of (3.8.a) the reference time is the time of the prediction. All that (3.8.a) tells us about this time is that it is before 2017. But how far that time was before 2017 is left open; in particular there is no reason why it couldn't be well before 2000. On the assumption that it was the sentence with *in the year 2000* becomes acceptable too.[4]

(3.8)

[4]There previous footnote is applicable here as well. Uttering the 'in the year 3000' version of (3.8.a) today is marked because we know that the year 3000 is in the future of the utterance time. When that knowledge is available then there is a preference for the simple future tense.

a. It was predicted once that civilization would come to an end through a world-wide epidemic (in the year 3000/in the year 2000).

b. It was predicted once that civilization will come to an end through a worldwide epidemic (in the year 3000/# in the year 2000).

For a final example involving complement clauses consider (3.9).

(3.9) Mary told me last week that she was going to file for a divorce in a couple of weeks but that she would tell Fred only then that she had (filed for a divorce).

This sentence is more complicated than those we have looked at so far in that it has one indirect discourse clause embedded within another. In sentences involving such multiple embeddings it is possible for a complement clause with past perfect tense to refer to a time that is in the future of the utterance time – something that might seem surprising given that normally the past perfect is used to refer to times that are in the past of some other time that is itself in the past of the speech time. (Recall our discussion of (3.4).)

These examples should give a flavor of the complexities involved in the interpretation of tense. Furthermore, they show how important inter-sentential temporal linking is in the interpretation of discourse, and that tenses play an important role in that – though, as we have also seen, other factors play their part in this too. And, finally, examples like those in (3.2) and in (3.7) show that aspectual issues – i.e. whether a sentence or clause describes an event or a state – play an important part in determining temporal relations.

3.2 A first Sample of DRS Construction Bottom Up

The examples of the last section are plenty for us to get our teeth into. In fact, I already noted that it won't be possible to give a proper account of any of the phenomena they were chosen to illustrate until much later. But some components of our treatment of tense and aspect can be put into place without too many preliminaries, and with those we will start.

Even for those components, however, a certain amount of preparatory groundwork needs to be done first. There are three main reasons for this. First, since we are now going to deal with questions of tense and aspect, we will need a finer analysis of syntactic structure, in which the tenses of verbs are represented separately from the verb lemmata to which they are attached as affixes. (We will soon also need separate representations within our syntactic structures of certain aspectual operators, such as the progressive and the perfect, but we will face that further complication when we have to.) The second reason is that we will from now on treat verbs as descriptions of events or states, in the concrete sense that each occurrence of a verb in a sentence will give rise to a discourse referent for the eventuality that occurrence describes. This makes for representations that look rather different from those presented in PART I. The third reason why we need to do some ground clearing before we can start with a systematic investigation of tense and aspect is that we will now be building DRSs bottom up. We already remarked in Subsection 2.12.4 that going bottom up will require special provisions, which are needed in order to deal with anaphoric pronouns: we need some way of keeping the dref for the pronoun on hold until its antecedent has become available and chosen. Something like that will also be needed in connection with the eventuality drefs introduced by verbs. In large part this is because of an assumption that has come to be widely accepted within generative syntax and that we are also going to make, viz. that the information provided by tense is quite high up in the syntactic tree of a tensed sentence, whereas a lexical verb will occur lower down, as a 'leaf' of the syntactic tree for the verb's clause, and often there will be several syntactic layers separating the two. Since we will also assume that lexical insertion for a verb– that is replacement of the occurrence of the verb at its node in the syntactic tree by the semantic representation specified for the verb in the lexicon – introduces a dref for the event or state described by the verb, this dref will have to be kept into an accessible location at least until the information associated with tense provides a temporal location for this event or state.

But the best way to find out about these complications is to look at a concrete example. So we start with the DRS construction of a very simple sentence, (3.10).

(3.10) Frieda smiled.

Our first concern is to determine a suitable syntactic representation for this and other tensed sentences, in which tense has its own, separate representation. We adopt the principle, widely assumed by syntacticians, that the information provided by tense is located in a position high up in the tree, as part of a projection level which also contains the subject DP. The relevant part of the tree configuration is that in (3.11), in which the tense information is given by the node labeled 'T'. [reference(s) to the syntactic literature?]

(3.11)

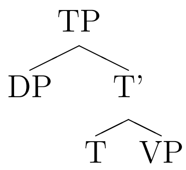

We assume that the simple past tense contributes a feature 'past' and that that feature is attached to the T node. So the syntactic representation we get for (3.10) is the tree in (3.12). This will be the LF from which we are going to build our semantic representation.

(3.12)

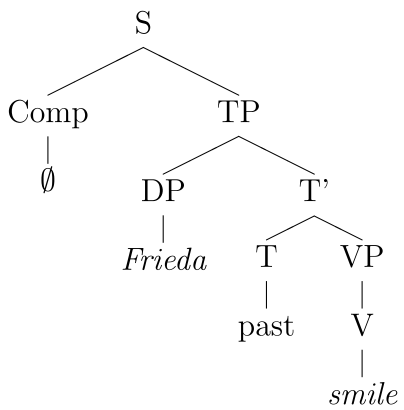

If we are to build our representation bottom up we need to start with entries for the lexical items. These entries should provide the basis for the *lexical insertion* operations that must be performed when an occurrence of a lexical item as one of the leaves of an LF tree is replaced by its semantics. For now we stick with our earlier treatment of proper names according to which their semantics consists of a representing dref x together with a condition 'PN(x) to state that x stands for the given bearer of the name. The entry for the name *Frieda*, in the format we now adopt, is given in (3.13.b) below.

Of more interest is the entry for the intransitive verb *smile*. We treat *smile* as an event verb – that is as a verb whose basic function is to describe events. In fact, we will assume from now on not only that the semantics of all verbs involves eventualities, but also that these eventualities can be treated as arguments of the verb. Even more than that, we will treat the eventuality arguments of verbs as their *referential* arguments, in a sense of 'referential argument' that I will explain momentarily. In addition to its referential argument a verb (or other predicate word, for that matter) may have additional arguments. In fact, verbs (with perhaps a small number of exceptions such as 'weather verbs' like *rain*, *snow* and a few more) always have at least one *non-referential* argument, and often more (one non-referential argument for 'truly' intransitive verbs, and more for transitive verbs and verbs with various prepositional arguments). Thus we will assume that *smile* has two arguments, one for the events that are described by occurrences of the verb and one for the verb's grammatical subject.

In what follows the representation of events, states and times will be of central importance and our DRSs will be replete with discourse referents for entities of these three sorts. It will be convenient to have a convention that allows us to immediately recognize drefs as standing for entities of these respective sorts, and to that end we follow the widely established practice of using 't' to refer to times, 'e' to refer to events and 's' to refer to states. In other words, we will use drefs of the forms t, t', .., t_1, t_2, .. to represent times, e, e', .., e_1, e_2, .. to represent events and s, s', .., s_1, s_2, .. to represent states. And in addition to these three kinds of 'special sort symbols' we will also make use of fourth kind: ev, ev', .., ev_1, ev_2, .. will be used to represent eventualities, entities that could be either events or states. Formally the use of such special dref symbols is like that of the lower case Roman, upper case Roman and lower case Greek letters that we introduced at the point in Part I when we included plurals in the repertoire of natural language expressions for which our DRSs language must provide the correct representations. Greek letters were introduced as number neutral drefs, upper case Roman letters as shorthand for the combination of a number-neutral dref (ξ, say) together with the condition 'non-atomic(ξ)', and the use of lower case Roman letters was reinterpreted as a combination of a number-neutral dref ξ with the condition 'atomic(ξ)'. Analogously we now introduce into our DRS language besides the new special symbols for drefs the predicates 'Event' and 'State' and treat the use of an event dref e as short for the combination of an eventuality dref ev and the condition 'Event(ev)' and likewise for state drefs. (Strictly speaking we ought to also introduce the predicate 'Eventuality', so that the use of eventuality dref ev can be seen as short for the combination of an arbitrary dref x together with the condition 'Eventuality(x)', but the need for this reduction will not arise in anything we are going to do.) In the same way the use of a dref such as t can be taken as the combination of an unsorted dref x combined with the condition 'time(x)' (short for 'x is a time').

One difference between referential and non-referential arguments is that referential arguments are introduced by the lexical entries of their predicates, whereas non-referential arguments are provided by other phrases in the sentence. (In case there is no overt realization of a non-referential argument, it must be reconstructed from the context or accommodated). This means that at the level of the lexical entry of a predicate word the slots that will eventually be filled when the predicate is used as part of a well-formed sentence must be marked in a way which shows that they are still to be filled. In this regard they differ from the referential argument slot of the predicate, which is filled right away with the referential argument that comes as part of the predicate's lexical entry. We mark this distinction – between argument slots that have not yet been filled and slots that have been filled with an argument – by using *underlined* lower case letters to mark slots that are as yet unfilled. Thus note well: underlined letters are **not** drefs; they do not function as representations of entities that are being talked about. They are simply used to mark gaps that wait to be filled and that in a completed representation must have been replaced by gap fillers (i.e. actual drefs) .

There is one further distinction we need to make. The 'drefs' that are part of lexical entries for predicate words cannot in general be the same drefs that get inserted into their referential argument slots when the semantics of the entry is inserted for an occurrence of the word in a sentence tree. For the same predicate may have several occurrences in the same sentence or discourse and these different occurrences may involve different referential arguments. So the drefs filling the referential slots in the semantics that get inserted for those different occurrences must be different as well. We deal with this problem as follows. Whenever the semantics of the lexical entry is inserted into the semantic representation that is constructed for a sentence (or other well-formed expression) a fresh dref is chosen to play the part of the referential argument in the lexical entry. In practice this will usually result in replacement of this referential argument; but it doesn't have to, and in examples we can often make do with the referential argument dref of the entry itself.[5]

[5] Another way to handle this would be to have a special symbol for 'schematic' drefs for referential arguments in lexical entries. Lexical insertion then would require the replace-

Here then are the lexical entries for *smile* and *Frieda*, in the format on which we settle for the time being.

(3.13) (lexical entries for the verb *smile* and the proper name *Frieda*)

		smile (V) e	nom $\underline{x}$
a.	Sel. Restr:	event	human
	Sem.Repr:	$< e \mid \begin{array}{\|c\|}\hline \\ \hline e\text{: smile'}(\underline{x}) \\ \hline\end{array} >$	

		Frieda (PN) x
b.	Sel. Restr:	human
	Sem.Repr:	$< x \mid \begin{array}{\|c\|}\hline \\ \hline \text{Frieda'}(x) \\ \hline\end{array} >$

So that we do not lose too much momentum let us proceed with the construction of the DRS for (3.10) and then follow up with some explanations and motivations for why it is desirable to distinguish between referential and non-referential arguments.

The first operations that need to be performed when going bottom up are the lexical insertions. The results of performing both lexical insertions on the LF in (3.12) are shown in (3.14). The explanation of how these results are to

ment of the schematic drefs by real drefs. But to do this we would have to introduce yet another type of symbol, overloading the notation and making everything just that much harder to digest.

be read is stated below. Since these are the first drefs to be introduced into the semantic representation we are building, nothing speaks against against using the very drefs that are displayed in the Semantic Representations of the lexical entries of (3.13).

(3.14)

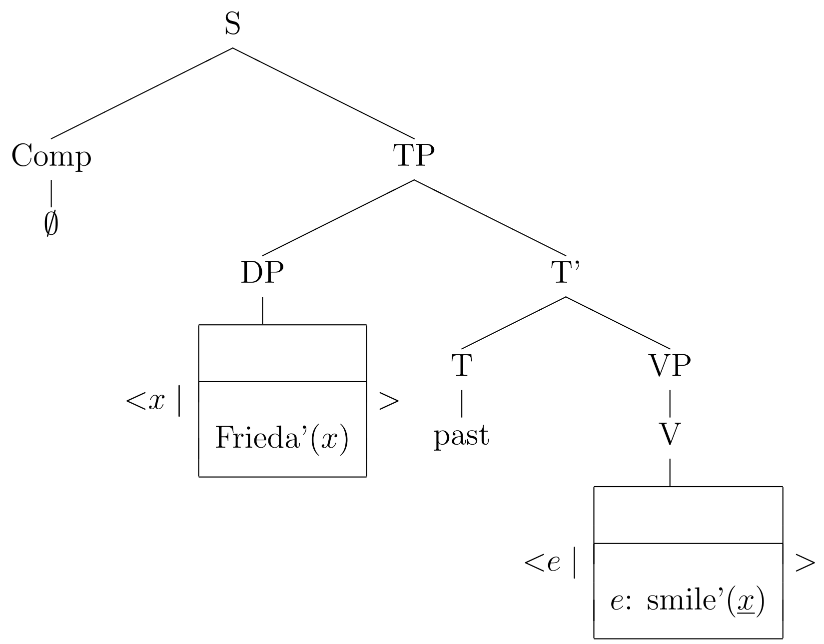

The representations that replace the lexical items in (3.14) are prototypical for most of the representations we will be dealing with from now on. These representations consist of two components, separated by a vertical stroke:

(i) a *store*
(ii) a DRS

The store is a list of discourse referents which still need to undergo binding – either by insertion into a DRS universe or by a quantifier in a duplex condition – at a later stage. The store is thus conceptually similar to 'Cooper stores', introduced into Montague Grammar by Cooper in (Cooper 1983) for what is in essence the same purpose: as a device for delaying the binding of variables until some subsequent stage, which can come quite a bit after they have been introduced. This also gives a clue as to the meaning of such representations: they can best be thought of as expressing relations between the discourse referents in the store, with the DRS following the store expressing the relation. Thus the structure inserted for *Frieda*, with just the dref x in the store can be thought of as expressing the property of x that consists in x being the bearer of the name *Frieda*. The interpretation of the structure inserted for *smile* presents an additional problem in that the DRS contains the argument slot marker $\underline{x}$. It is simplest to think of this element as also marking a term of the expressed relation. So – this is a correction of what has just been proposed – structures like those inserted in (3.14) can be thought of as relations between (i) the elements represented by the drefs in the store and (ii) the elements that will be represented by the drefs that will eventually fill the slots marked by slot markers in the DRS. According to this specification the structure inserted for *Frieda* expresses a 1-place relation (i.e. a property) and the structure inserted for *smile* a 2-place relation, between events and human individuals.

Since V is the only daughter of VP, the semantics of V is simply passed up to VP. As next step the VP semantics must be combined with the information carried by the feature 'past'. But what is the semantics that this feature contributes? That is one of the central questions we will have to answer. We are going to answer it piecemeal and the answer that we will give at this point is only a first step in the right direction. It does nothing to account for all the cross-sentential 'donkey like' effects that are prominent in the examples we have discussed in Section 3.1. The full story will be given only after we have introduced the necessary provisions for dealing with presuppositions.

We will treat the semantics of past and other temporal features (associated with other tense forms) as a kind of operator which puts a temporal constraint on the referential argument of the representation of the verb projection to which it is applied. I will represent this semantics in a format that is as close as I can make it to the lexical entries in (3.13). But before I do this, let us first complete the DRS for (3.10), so that we see what the past tense feature does in this particular case. (3.15) shows the result of combining the VP semantics with that of past.

(3.15)

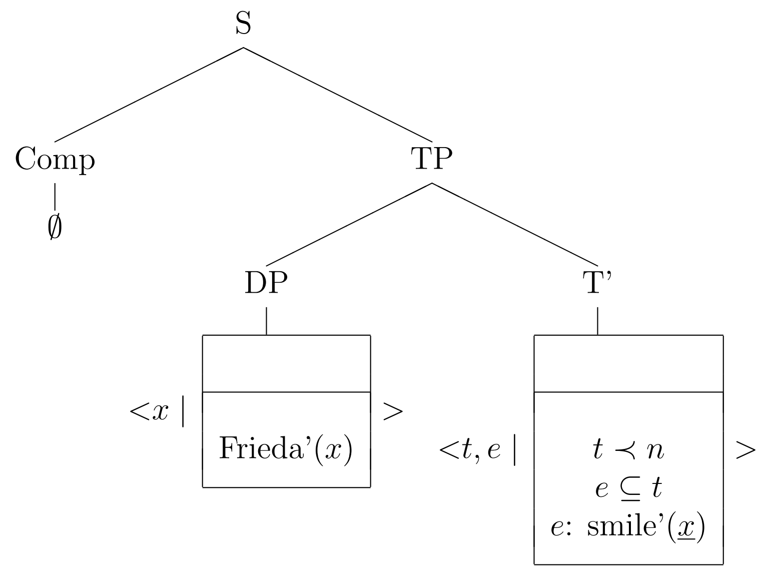

The main contributions made by past to the semantic representation of the T' node are the conditions '$t \prec n$ and '$e \subseteq t$'. The first of these conditions says that the time represented by t precedes the utterance time n. n is an *indexical discourse referent*. It serves to refer to a feature of the *utterance context*, viz. the utterance time – that time at which the utterance in question is made. The sudden emergence of a reference to utterances and their contexts may seem to come out of nowhere here. But this is one of the consequences of our forging ahead with the construction of the DRS for (3.10) before all background issues have been cleared. The central conceptual issue we are facing when dealing with tense is that the semantics of most uses of most tenses has an *indexical dimension*: it involves some relation between the time at which the tense locates the eventuality described by its clause and the time at which the clause (or the sentence or discourse of which it is part) is uttered. This means that from now on we can no longer represent sentences (and longer pieces of discourse or text) in abstraction from particular uses of them, but only as utterances, or parts of utterances, that are made in some particular context and, more specifically, at some particular time. It is to this time (which in the work of Kaplan and others is denoted as 'c_T', as short for 'the time of the utterance context c'), that n must be taken to refer when we evaluate the resulting sentence DRS K for truth or falsity in a given model M. That is, it is now only possible to evaluate K in M *as the semantic representation of a given sentence uttered at a given time* c_T. And that amounts to the following: K (as representation of an utterance made at c_T) is true in M iff there is a verifying embedding f of K in M which maps n to c_T.

Note that this presupposes that the model M contains the time of the given utterance as one of its elements. This imposes certain constraints on the models that are suitable for the evaluation of sentences and larger pieces of discourse in which the semantics of tense is taken seriously. We will address the question about the form of models later on Section 3.6. At this point what matters is to firmly keep in mind that the analysis of tense imposes this constraint on the truth evaluations of DRSs in models and on the models in which such evaluations are possible.[6]

Since n is a discourse referent, one might, in the light of what we have seen so far, expect it to appear not only as an argument of the condition '$t \prec n$' but also in a DRS universe, or else in the store, waiting to be properly bound. As a matter of fact, n is different from the discourse referents we have encountered so far. It always gets its value from the context in which the represented utterance is made. To make sure that it is among the drefs that are mapped into the universe of a model by verifying embeddings it should be included in the main DRS Universe, with the stipulation attached to it that verifying embedding functions always map it to the utterance time. Since the main DRS Universe is the place where n always ends up, we forgo the trouble of putting it there explicitly in our displays of DRSs. But n is always present there nonetheless – a member of the main Universe of every DRS we are going to build from now on.

The second condition, '$e \subseteq t$', should be read as 'the event e is temporally included within the (interval of) time t'. In other words: (i) $\subseteq$ is to be read as temporal inclusion and (ii) this relation holds between e and t by virtue of the 'duration' of e – viz. that part of the course of time during which e is going on – being included within t. (Often in the literature this relation is expressed as '$\tau(e) \subseteq t$' where τ is a function that maps each eventuality onto the portion of time that it occupies. In (Kamp & Reyle 1993), where $\subseteq$ is also used as temporal inclusion relation that is directly applicable to both times and eventualities, the function τ goes by the name of 'dur'.)

In addition to the introduction of the two conditions just discussed, the effect of past on the T' node of (3.14) is the introduction of the time dref t to represent the temporal location that the tense provides for the described eventuality. Both this time dref and the dref representing the eventuality

[6] We will see later [in Section ?] that there is also another way to interpret DRSs like the one that will emerge from the present construction, which corresponds more directly to Kaplan's notion of *character* (see (Kaplan 1989)).

that the time dref serves to locate are put and kept, respectively, in the store at this point, since they may be needed later on for further processing. (Such a need doesn't arise for the sentence with which we are dealing, but it will arise in many others. So our general policy must be to keep these drefs in the store for as long as this need may still arise.)

This is all we will say for now about the effects that the feature past has on DRS construction.

The next construction step combines the semantic representation of the DP *Frieda* with that of the T' node. Here we encounter another problem that is inherent to Bottom Up as opposed to Top Down construction: Where should the dref x that represents the subject Frieda according to the semantic representation of *Frieda* go so that the semantic representation correctly connects DP with its predicate? Intuitively the answer is plain: x should be inserted into the subject argument slot $\underline{x}$ of the predicate smile'. But how can this intuition be captured formally? Here we run up against a non-trivial issue about the syntax-semantics interface. The construction algorithm that we are developing is only one of the several components of a more complex processing procedure of linguistic inputs. Which component of this larger procedure should bear the responsibility for correlating the argument phrases in a given clause with the argument slots of their predicates? Or, focused more directly on the argument slots of verbs and their occupiers, what is responsible for determining which DPs in a clause fill which argument slots of the clause's verb? This is a well-known problem from the linguistic literature, often referred to as the 'linking problem' [reference(s)?]: Given that the verb of a well-formed clause acts semantically as a predicate with a certain number of non-referential argument positions, which of the DPs in the clause belong to which of those argument slots?

We will assume that this is the responsibility of the syntactic parser, the component of the larger procedure that provides the DRS construction algorithm with its inputs. This is not an unreasonable assumption, since a parser, which authenticates a string as a well-formed expression of the language by assigning it an authenticating syntactic structure, must have access to the lexicon in order to be able to do its job: the parser needs to identify the elements of the string it is parsing as words belonging to the lexicon and it needs to identify the grammatical categories that the lexicon assigns to them (since these categories are part of the structure the parser must assign to the string). In particular, the parser must find the verbs in the string and for each verb V it finds it must identify the number of its arguments (i.e. whether it is an intransitive, a simply transitive, a ditransitive verb). And on the basis of that information it will have to determine whether the string has the right set of argument phrases to fill the verb's argument slots. These argument slots occur explicitly in the semantic representation that is part of the verb's lexical entry. Moreover, a properly functioning parser will not only have to check that the number of argument phrases in each clause matches the number of argument slots of its predicates but also which argument phrase goes into which slot. (This is needed for instance to check that the selection restrictions associated with a given argument phrase are compatible with those associated with the slot that it fills.)

Assuming that the connection between slots and their fillers is part of the information that the parser establishes, the next question is how this information is made accessible to the construction algorithm.The technical device that is most often used to this end, and that as far as I can see is as good as any other, is to co-index argument slots with their fillers. But there is a practical problem with the implementation of this device: in the LFs that we are using as inputs to DRS construction – such as for instance the tree in (3.12), the LF for the construction we are involved in right now – there is no way to add the co-indexing because the argument slots aren't overtly represented in these LFs. They are explicitly represented only after the semantic representations from the lexical entries for the predicate words have been substituted for the words themselves, as in (3.14). The quandary, in other words, is that the first opportunity we have for making the linking between DPs and argument slots explicit arises only after lexical insertion.

The solution to this problem would be to let the parser deliver LFs in which lexical insertion has already been executed and the coindexations between slots and fillers have been made explicit. In spite of this we will often start by showing LFs without lexical insertion as starting points for DRS constructions, as a way of showing more clearly what syntactic structure is assumed as point of departure for the construction.

The main upshot of this is that from now on syntactic trees with lexical insertions will also be decorated with slot filler coindexations. In particular, the tree in (3.14) now takes the form shown in (3.16), with the subject DP *Frieda* coindexed with the slot symbol $\underline{x}$ in the non-referential argument position of the predicate 'smile' '.

(3.16)

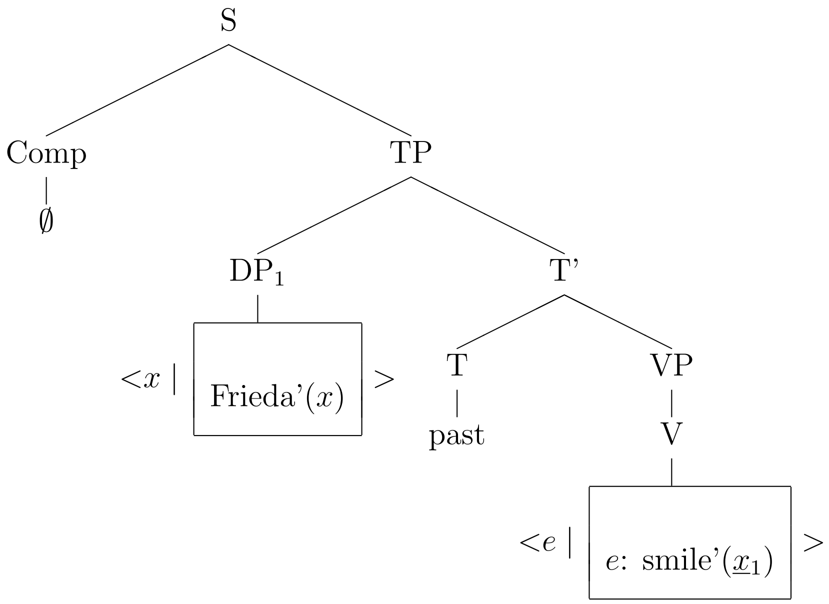

The coindexation in (3.16) is not affected by the steps of the DRS construction for 'Frieda smiled' that have already been performed; so the structure which requires combining the semantic representation of the subject with that of the T' node should be as shown in (3.17).

(3.17)

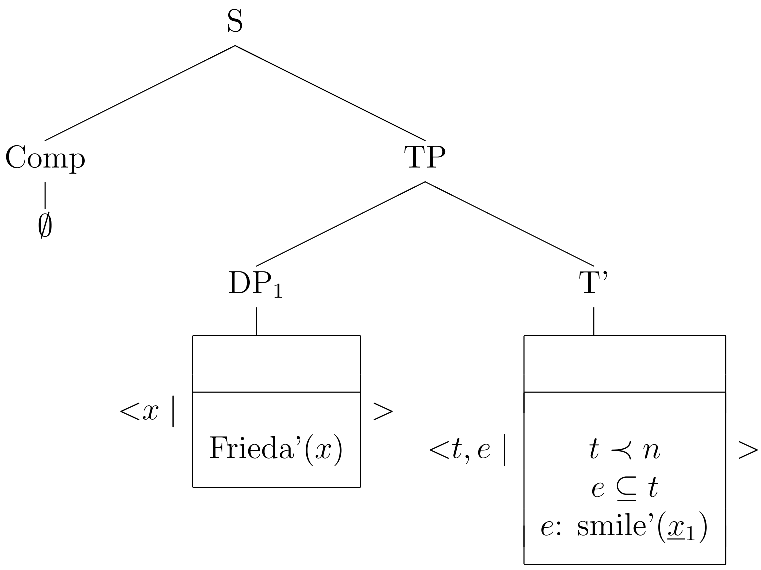

The coindexation can now be used to guide the insertion of the referential argument x of the DP representation into the coindexed slot of the verb.

The further operations that have to be carried out at this point are: merge of the DRSs of the two representations – in the present instance this amounts to adding the condition 'Frieda'(x)' to the Condition Set of the DRS from the T' representation – as well as merging the two stores, which in this case means adding the dref x to the store list from the T' representation. (Since we have defined the stores as lists, the order in the drefs matters in principle. For the time being we assume that in 'argument insertion operations', such as the one we are involved in now, the dref which plays the part of referential argument of the argument phrase (here: the subject phrase *Frieda*) is added to the store of the predicate, sometimes to its end and sometimes to its beginning. In the case at hand the one element of the DP store, x, is appended to the end of the store of the TP representation, in other words after t and e. In connection with the present example it is impossible to come up with a rationale for this decision, let alone with an answer to the general question how the store of the representation of an argument phrase should be combined with the store of the representation of its predicate. This is another point that we will come back to.)

(3.18)

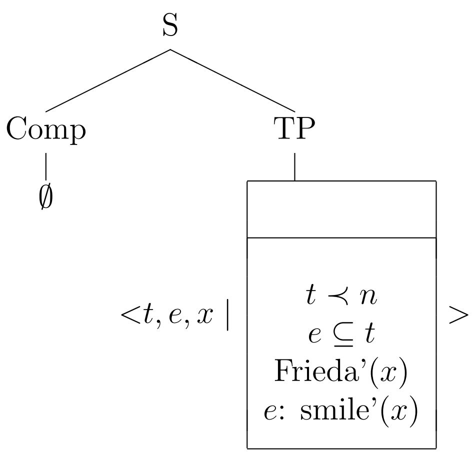

There is only one more step to be performed, which implements the contribution from the node labeled Comp. What contributions that Comp nodes can make to sentence semantics is a quite complicated matter, which has to do with the various ways in which main clauses can combine with subordinate clauses and also with the ways in which clauses can be coordinated through the use of conjunctions such as *and*, *but*, *or*. This is an aspect of the semantics of complex sentences that we will hardly touch on in these notes. By and large we will be concerned only with the case where Comp is the Comp node of a main clause.

For now (that is, until we introduce presuppositions into DRS constructions) the only operation that may be needed to achieve turn a semantic representation attached to a TP node into the DRS for the main clause S node that immediately dominates it is the transfer of remaining drefs in the store to the main Universe of the DRS to its right. The now empty store can then be discarded, so that only the DRS remains. (If all the drefs have already been transferred from the store, then nothing needs to be done at this step.).

(3.19)

$$\begin{array}{|c|}\hline t \quad e \quad x \\ \hline \\ t \prec n \\ e \subseteq t \\ \text{Frieda'}(x) \\ e\text{: smile'}(x) \\ \hline \end{array}$$

3.3 A Lexical Entry for the Simple Past

In view of the primary concern of this part of the notes the most important step in the DRS construction for (3.10) was that in which the information contributed by tense, represented by the feature past attached to the T-node, is combined with the semantic representation of the VP node that is the T node's sister. And that is also the step for which no proper justification has yet been given. Giving a proper account of the semantics of the past tense is not easy and we will do so in two steps. This section presents the first step.

The transition from (3.14) to (3.15) makes plain what contribution the past tense of (3.10) makes to its semantics. The feature past contributes a time dref t with the DRS conditions '$t \prec n$', which says that t is in the past of n, and '$e \subseteq t$' which temporally locates the event e of Frieda's smiling within t. Before we move from this particular case to a general 'lexical entry' for the feature past, we need to see whether the contribution made by past is always of this form. And the answer to that question is an emphatic 'no'. The principal reason for this has to do with aspect. One of the most important aspectual distinctions is that between *perfective* and *imperfective* aspect. This distinction will play a central part in our analyses of tense and aspect generally and various things will be said about it as we go along. All that matters for the moment is the way in which eventuality descriptions with perfective and imperfective aspect relate to their location times. The relevant difference cannot be explained at the hand of a sentence as simple as (3.10); but it can be explained by comparing the sentences in (3.20).

(3.20)

a. At 18.00 Frieda closed the shop.

b. At 18.00 Frieda was closing the shop.

There appears to be a clear difference in meaning between (3.20.a) and (3.20.b). (3.20.a) says that the time referred to as '18.00' was the time at which Frieda closed the shop, i.e. the time at which an event of her closing the shop was located. (3.20.b), on the other hand, seems to imply that at 18.00 a certain process was going on, one which, for all the sentence says, will have started some time before 18.00 and may go on for some time after. ((3.20.b) would have been a natural thing to say for an eyewitness who was passing the shop at 18.00, saw that Frieda was in the process of closing up – something involving a number of different actions, like carrying inside the wares that were on display on the sidewalk, pulling down the blinds, switching off the lights etc – but who didn't wait to see the end of it all and consequently has no information about how much longer it took Frieda to complete the process, or for that matter whether the process was completed at all.)

Here is another way of describing this difference. (3.20.a) describes an event – that of Frieda closing the shop – as located in its entirety at the time described as '18.00'. (That requires us to think of this time as itself being a (short) interval rather than an instant – an indivisible, infinitarily small part of physical time conceived as the real number axis – since otherwise the event could not fit within it; but that is indeed how we often understand time specifications involving clock times in natural language discourse.) The temporal relation between the eventuality described by (3.20.b) and the time contributed by *at 18.00* is understood differently. Here it is the eventuality that is taken to temporally include the time denoted, rather than being temporally included within it. This difference, between (i) the time of the temporal adverb including the described eventuality and (ii) the described eventuality including the adverb time, correlates with the traditonal distinction between perfective and imperfective aspect: for sentences with perfective aspect, such as (3.20.a), the eventuality is included within the adverb time;

for sentences with imperfective aspect, such as (3.20.b), the relation is the reverse.

As the discussion of the last paragraph may already have suggested, the perfective-imperfective distinction also correlates with yet another one, that between *event descriptions* and *state descriptions.* This correlation is hard to prove (even in an informal sense of 'prove'). Our untutored intuitions about what is an event and what is a state only go so far, and our intuitions are even less articulate in relation to the question whether a given clause should be regarded as describing an event, or as describing a state. Nevertheless, over the years a body of observations and theorizing has built up which suggests that this correlation is more than a convenient whim and that it is consistent with what our intuitions seem to tell us.

At this point we cannot do better than postulate this correlation – as a kind of working hypothesis you might say – together with what it implies for event and state location in the light of what has just been said about temporal location.

(3.21)

a. Perfective clauses describe events. Imperfective clauses describe states.

b. Temporal location times t locate events via the condition '$e \subseteq t$', and states via the condition '$t \subseteq s$'.

If (3.21.b) is to yield the intuitively right results when applied to the sentences in (3.20), then the sister node to T of the first sentence should be analyzed as the description of an event and the sister node to T of the second sentence as the description of a state. But what justifies us in saying that the second *is* the description of a state? This is one point where the correlation in (3.21.a) may appear like a stipulation. If application of (3.21.b) is to yield the semantic representation we want, then we must assume that the progressive in (3.20.b) is responsible for turning the event description of (3.20.a) into a state description. But is it plausible that progressives can turn event descriptions into descriptions of states? It certainly isn't intuitively obvious. There appears to be a difference between typical state descriptions like *know*, *love* or *be sick* and a description like *be closing the shop.* It might be thought that it is more natural to classify *be closing the shop* as the description of a process than as the description of a state.

Indeed, some theorists make a three-fold aspectual distinction between events, states and processes. We could do that too, while stipulating that states and processes behave in the same way as far as temporal location is concerned. But for our purposes the binary distinction between events and states – with states and processes lumped together if you like, or with the category 'process' subsumed under the category 'state' – will work well enough, and so we adopt this simplification. A more detailed discussion of the progressive follows in Section 3.5.2.[7]

In our discussion of the sentences in (3.20) we have focused on the way the eventualities they describe are located by the temporal adverb. We did not say anything about location by tense, so the assumptions we have made thus far still do not tell us how to state our 'entry' for the feature past and to justify our DRS construction for (3.10). Here it is even more difficult to argue from direct observations of examples, which is why we switched from (3.10) to (3.20) at the outset of this discussion. And we can't do much more than simply state, as another working hypothesis, our assumptions about the interaction between eventualities and the location times introduced by tense. In a way we have already done that by stating (3.21.b) the way we did. For (3.21.b) speaks simply of 'location times', without differentiating between those introduced by temporal adverbs and those introduced by tense. In other words, location by tense is subject to the same difference we have assumed for adverbial location: events are included within the location times introduced by tense and states temporally include these location times.

To summarize: the semantic representation of the sister node to T is always an eventuality description and this description is either the description of an event or the description of a state. Furthermore, the feature located at T always prompts the introduction of a dref t for the location time of the described eventuality, and that time is interpreted as temporally including this eventuality when it is an event, and as temporally included within it when it is a state. And when the T-feature is past, then t itself is located in the past of n.

One consequence of this is that the semantics of past involves a disjunction, between the case where the sister to T specifies an event description and that where it specifies a state description. This means that there will have to be a disjunction in the lexical specification of past somewhere. That may

[7]The proposal to treat the progressive as turning event descriptions into state descriptions is one of long standing. See in particular (Vlach 1992).

look like a kind of hack, in which a disjunction is used as stand-in for some deeper underlying generalization which hasn't yet been captured – and in a way it is. In fact, there is an alternative way to deal with temporal location which obviates the need for disjunctions in the lexical entries of past and other tense features. But we will not get to this until much later, in Section 4. For now we will make do with disjunctive entries. Although not optimal, this isn't all that bad, especially since the same kind of assumption is made in a good deal of other work on tense and aspect in the current literature.

The disjunction with which we will make do in our entry for past (and, later on, the feature fut for the future tense) isn't quite like the disjunction operator familiar from classical propositional logic – the operator $\vee$ characterized by the principle that the formula $A \vee B$ is true if either A or B or both are true and that that is all there is to it. The disjunction we need here is one that comes with a requirement to the effect that it must be possible to determine which of its disjuncts is the correct one. In other words, it is a kind of binary selection function, which forces the choice of one of its disjuncts, whereupon the other disjunct is discarded. In the case of the feature past we assume that the choice between its disjuncts can be made on the basis of the information that is provided by the input representation (that is, the semantic representation of the sister node to the T node to which the feature past is attached). The input representations to past do provide this information because they will always contain an overt display of the eventuality they describe, and that eventuality will always either be an event dref or a state dref. The symbol that we will use for the kind of disjunction that we need in the lexical entry for past is '$\overset{!}{\vee}$'.[8]

There is another matter that needs to be sorted out before we can present the semantics of the feature past in the form we want for now. This one has

[8] $\overset{!}{\vee}$ has a more general use within *underspecification formalisms*. The point of underspecification formalisms is to provide the possibility of postponing steps in the construction of 'fully specified' representations until all the information that is needed for their execution has come in. Suppose F is a formalism which defines 'fully specified' representations. An *underspecification formalism* UF *for* F is an extension of F which contains besides the fully specified representations of F also 'underspecified' representations, representations that can be turned into F-representations and normally can be turned into more than one. UF will come with its own syntax and model-theoretic semantics; and if it is worth its mettle, it will also come with a logic of its own, which enables transparent and efficient verifications of entailment relations, especially those between underspecified representations and the ugly specified representations into which they can be converted. For an authoritative statement of Underspecification in the context of DRT (Reyle 1993) and for applications to temporal undersepcification see e.g.(Reyle, Rossdeutscher & Kamp 2007).

to do with the fact that, speaking informally, past is an 'operator', whose semantics manifests itself through the way it transforms input representations into output representations. We make its operator role explicit by stating its semantics in the form of an input-output relation, with the input to the left of the 'leads to' arrow '$\leadsto$' and the output to the right of it. The output takes the form of a $\overset{!}{\vee}$ disjunction, with one disjunct containing the condition 'Event(ev)' and the other the condition 'State(ev)'. Since the input representations always decide between these two conditions – they are either the description of an event or the description of a state and so contain, implicitly if not explicitly[9], exactly one of the conditions 'Event(ev)' and 'State(ev) – the choice between the two $\overset{!}{\vee}$-disjuncts will always be determined. It is determined by logical compatibility, if you like; for choosing the wrong $\overset{!}{\vee}$-disjunct would lead to an inconsistent representation, one that would entail the conjunction of the two incompatible conditions 'Event(ev)' and 'State(ev)'. (That these two conditions are incompatible wasn't stated explicitly so far, but is a central part of the eventuality ontology we adopt: 'Event' and 'State' are disjoint subcategories of the category 'Eventuality'.)

Another difference between entries for lexical predicates and lexical entries for operators like that triggered by past is that operators involve neither referential nor non-referential arguments. Their only 'arguments' are their input representations. There is still room for Selection Restrictions, but these Selecvtion restrictions will apply to drefs bel;ongoing to the input representations on the left of $\leadsto$. The entry for past is given in (3.22).

[9]The condition may be implicit in the input representation in that the dref representing the eventuality that the input describes is a symbol of the 'e'-sort, which only represents events, or of the 's'-sort, which only represents states. In the former case we could have chosen the notational variant in which the described event is represented by a neutral eventuality dref ev together with the condition 'Event(ev)'; likewise when the input is the description of a state

(3.22) (lexical entry for the tense feature ‘past’)[10]

past (tense feature)

Sel. Restr: eventuality description

Sem.Repr: $<ev_{ref}, ... \mid K> \rightsquigarrow$

$$<t, ev, ... \mid K \cup \begin{array}{|c|}\hline \\ \hline t \prec n \\ \begin{array}{|c|}\hline \\ \hline \text{Event}(ev) \quad ev \subseteq t \\ \hline \end{array} \\ \overset{!}{\vee} \\ \begin{array}{|c|}\hline \\ \hline \text{State}(ev) \quad t \subseteq ev \\ \hline \end{array} \\ \hline \end{array} >$$

The Semantic Representation in (3.22) should be seen as a shorthand for (i) an explicit description of the kinds of input representations for which the operator triggered by past is defined and (ii) the formal manipulations involved in turning the input representation into the corresponding output representation. But most of what the Sem. Repr. is trying to convey should be clear. The input structures consist of a store and a DRS K, and the store starts with an eventuality dref ev. The store may have additional drefs in it, but need not. (It doesn’t in the example we have gone through.)

What operations are needed to transform the input representation into the output representation should be clear as well: A new time dref is added (to the beginning of the store) and the DRS K is extended with the conditions

[10]The subscripts $_{ref}$ will be explained in the next couple of sections. Just ignore them for now

of the DRS following $\bigcup$. ('$\bigcup$' is used here to denote the operation of *DRS merge*. This is the operation which, when applied to two DRSs K_1 and K_2, combines them into a new DRS whose Universe is the set-theoretic union of the Universe of K_1 and the Universe of K_2 and whose Condition Set is the set-theoretic union of the Condition Set of K_1 and the Condition Set of K_2.) We have already discussed the last of the new conditions, with its special disjunction $\overset{!}{\vee}$. To repeat, a DRS containing such a condition is unfinished: a choice must be made of one of the two disjuncts, so that in the final DRS $\overset{!}{\vee}$ no longer occurs.

To get an idea of how (3.22) works when it is applied to an actual input, let us go back to the critical step in the construction of the DRS for (3.10). (3.23) shows the representation obtained after dealing with the VP node. (In essence this is the same structure as the earlier (3.14).)

(3.23)

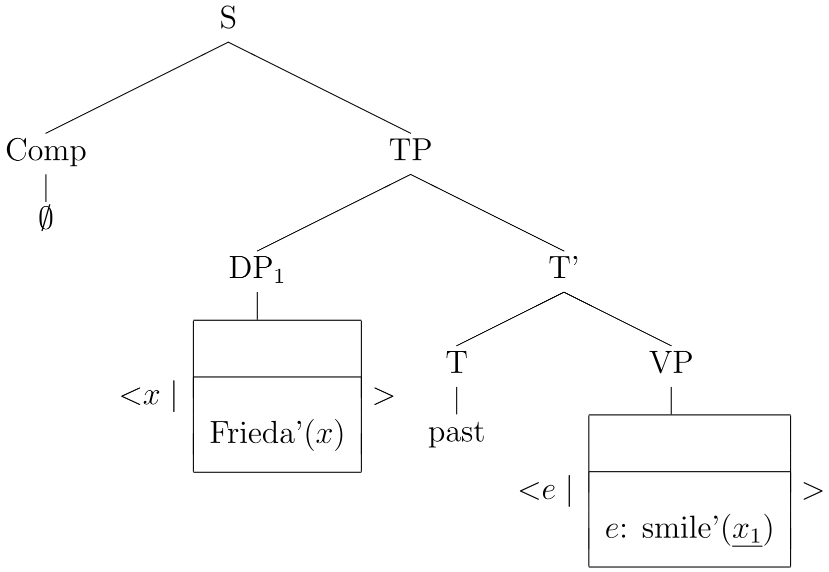

Recall that the use of the event dref e is short for a dref ev of the sort 'eventuality' together with the condition 'Event(ev)'. Furthermore, the dref e in (3.23) is the referential argument of the verb, and therefore can (and should) have been annotated as such. (Once more, the next subsection will provide details.) So the semantic representation of the VP in (3.23) can be seen as a shorthand for the structure in (3.24).

(3.24)

$$<ev \mid \begin{array}{|c|} \hline \\ \hline \text{Event}(ev) \\ \\ ev\text{: smile'}(\underline{x_1}) \\ \hline \end{array} >$$

The VP representation is now in a form that directly matches the description of the input given by the Semantic Representation in (3.22). And according to the transformation shown in this Semantic Representation the representation of the T' node should be as in (3.25).

(3.25)

$$<t, ev \mid \begin{array}{|c|} \hline \\ \hline \text{Event}(ev) \\ t \prec n \\ \\ ev\text{: smile'}(\underline{x}) \\ \\ \begin{array}{|c|} \hline \\ \hline \text{Event}(ev) \quad ev \subseteq t \\ \hline \end{array} \\ \overset{!}{\vee} \\ \begin{array}{|c|} \hline \\ \hline \text{State}(ev) \quad t \subseteq ev \\ \hline \end{array} \\ \hline \end{array} >$$

Note that the condition 'State(ev)' in the second $\overset{!}{\vee}$ disjunct is incompatible with the condition 'Event(ev_{ref})' in the Condition Set of the outer DRS. This disjunct is thus inconsistent, so its selection would not lead to a coherent interpretation. Therefore the choice must fall on the first disjunct (whose conditions are consistent with those in the outer DRS). So we end up with the T' representation in (3.26.a), or, after reintroduction of the shorthand involving the dref e, as in (3.26.b).

(3.26)

a. $$<t, ev \mid \begin{array}{|c|}\hline \\ \hline \text{Event}(ev) \\ \\ t \prec n \\ ev\text{: smile'}(\underline{x_1}) \\ \text{Event}(ev) \quad ev \subseteq t \\ \hline \end{array} >$$

b. $$<t, e \mid \begin{array}{|c|}\hline \\ \hline t \prec n \quad e \subseteq t \\ e\text{: smile'}(\underline{x_1}) \\ \hline \end{array} >$$

The information that the conditions ‘Event(ev)’ and ‘State(ev)’ are incompatible has to be coded somewhere. One option would be to encode it in the description of the application of the past tense operator: if the DRS K of th input operation contains the condition ‘Event(ev)’, then choose th upper disjunct of the $\overset{!}{\vee}$-disjunction; if it contains the condition ‘State(ev)’, then choose the lower disjunct! But that would require that we repeat this as part of the description of other tense operators as well. That wouldn’t just be cumbersome, but also a sign that we are missing an obvious generalization. The more natural way to proceed is to allow for a more general procedure of disjunct selection, which includes finding that one of the disjuncts is contradictory while the other is not as a trigger for selecting the latter.

This is how I motivated the application of past in the DRS construction for ‘Frieda smiled’. But then the information that ‘Event(ev)’ and ‘State(ev)’ contradict each other has to go somewhere else. The natural place is a repository of information about the logical properties of the logical form formalism we are now using – the one that licenses pairs <STORE,K> as logical forms. In our discussion of Montague Grammar we allowed for Meaning Postulates to be added to the Logical Form Formalism we used there, which enable us to state certain semantic properties of (the LC constants representing) certain predicates, or semantic relations between several predicates. Something like this is also needed as a supplement to the logical form formalism we are using now. The mutual exclusiveness of ‘Event(ev)’ and ‘State(ev)’ can be represented as one such Meaning Postulate. In the notation of DRT it can be represented by a conditional DRS condition, as in (3.27).

(3.27)
$$\begin{array}{|c|}\hline \\ \hline \begin{array}{|c|}\hline ev \\ \hline \text{Event}(ev) \\ \hline \end{array} \Rightarrow \begin{array}{|c|}\hline \\ \hline \neg \begin{array}{|c|}\hline \\ \hline \text{State}(ev) \\ \hline \end{array} \\ \hline \end{array} \\ \hline \end{array}$$

3.4 Referential and Non-Referential Arguments

The distinction between referential and non-referential arguments has been mentioned and used, but a proper account of it is still outstanding. The best way to explain the distinction is by having a look at *relational nouns*, like *friend* or *mother*. From a logical perspective it seems plain that these nouns denote binary relations: they are used to refer to individuals that are friends or mothers of other individuals. But note how such relations are expressed when we use these words, and pay attention to the curious asymmetry between the realizations of their two arguments. (3.28) gives a few examples.

(3.28)

a. A friend of Mary died.

b. Mary's mother died.

c. Mary invited two of her friends.

The point of these examples is that the second arguments of the relational nouns – the arguments whose mother or friend or friends is/are said to have died or been invited – are realized by overt noun phrases (the DPs of which they are the referents), *Mary* in the first two sentences and the possessive pronoun *her* in the third. For someone with a formal logician's take on things, who thinks of the way in which predication is expressed in the standard formulations of predicate logic, this is what one would expect: In the sentences of (3.28) the DPs in question function as terms that fill the second argument positions of the predicates *friend* and *mother*. But what about the first arguments of the occurrences of *friend* and *mother* in (3.28)? These are not realized by constituents occurring somewhere within the complex noun phrases of which *friend* and *mother* are the lexical heads – or, for that matter, by argument phrases elsewhere in the sentence. Rather, it looks as if it is the lexical heads themselves – the nouns *friend* and *mother* – that bring these arguments along. (For instance, the first argument of *friend* (3.28.a) is contributed by the noun *friend* and it is this argument that then becomes the 'referent' of the DP *friend of Mary*.) The central claim of this section is that this isn't just the way things look. It is the way things are, and not just for relational nouns, but for nouns of any kind. And not just for nouns, but also for other categories of predicate words, verbs, prepositions, adjectives. And not just for English, but for language after language, and quite possibly for all.

This double role, which the sentences in (3.28) illustrate for the relational nouns *friend* or *mother* – the role of supplying predicate and one of its arguments all at once – goes both against the spirit and the letter of Predicate Logic. In natural languages, however, this is apparently the way that predication is organized. But how are we going to account for this?

For the time being let us continue to focus on relational nouns like *friend* and *mother*. A moment's reflection suffices to see that the foundation of their twofold contributions to the sentences in which they occur has to be laid in the lexicon. Take the case of *friend*. The core of the Semantic Representation of the lexical entry for *friend* ought to be a predication of the form '$friend'(\alpha, \beta)$', with a DRS constant friend' that serves to translate the lexical noun *friend*. (This much should be plain from what we have been saying about the role that Semantic Representations from lexical entries of predicate words are playing in the construction of sentence DRSs containing those words.) But of the two argument terms α and β that occur in this lexical predication the first should be a 'true argument' – a dref x, which will identify the bearer of predications expressed by NPs with *friend* as nominal head – and which represents the 'referent' of any DP that can be formed from such an NP. In contrast, the second argument term, β, which is to be filled by means external to the noun after its lexical Semantic Representation has been inserted into a sentence representation, should not be a dref, but a 'mere argument slot', which will have to be replaced by a dref coming from elsewhere. We continue our practice of using plain lower case letters for drefs and use underlined lower case letters for argument slots (as we did for the argument slot of *smile'* in our DRS construction for ‘Frieda smiled.’).

Combining these stipulations about the argument terms of friend’ with the decisions already made in relation to the Semantic Representations of predicate words, we get (3.29) as entry for *friend*.

(3.29)

	friend (N, rel.) x	nom $\underline{y}$
Sel. Restr:	animate	animate
Sem.Repr:	$< x \mid \begin{array}{\|c\|} \hline \\ \hline \text{friend'}(x,\underline{y}) \\ \hline \end{array} >$	

As emphasized above, the distinction between the first and the second argument of relational nouns is crucial to the syntax and semantics of the higher projection levels of such nouns (the NPs and DPs in which the noun is the lexical head) and as we will see more clearly in a moment, it continues to be equally important at syntactic levels beyond the DP level, at which the DP functions as argument phrase to some predicate that is contributed by other word or phrase. We give this difference a name by calling the first argument of a noun like *friend* its *referential argument* and its second argument a *non-referential argument*.

The importance of the distinction between referential and non-referential arguments can hardly be overstated. To begin with, the observation that nouns bring along their own arguments isn’t restricted to relational nouns. Relational nouns are a good starting point for explaining the distinction, and in particular the double role that the noun plays in relation to its referential argument, as predicate and argument all in one, so to speak. But non-relatonal nouns – surely the bulk of nouns in English, and presumably also in all or most other languages – are like the relational nouns in that they too bring an argument along whenever they occur as parts of sentences: they too have a *referential* argument; the only difference with relational nouns is that they lack a non-referential argument.

For an example consider the noun *cat*. Its lexical entry, given in (3.30), shows it as having a referential argument x and no other arguments. And the entry's Semantic Representation just says that x is the bearer of the property expressed by *cat*.

(3.30)

cat (N)
x

Sel. Restr:

Sem.Repr: $$<x \mid \begin{array}{|c|}\hline \quad \\ \hline \text{cat'}(x) \\ \hline \end{array}>$$

What is most important about the distinction between referential and non-referential arguments for our analysis of tense and aspect is that the distinction applies to verbs just as it applies to nouns. Verbs too have referential arguments, and as with nouns these referential arguments are 'invisible': they too are brought along by the predicates of which they are the referential arguments – that is, by their verbs – and are not realized by some independent phrase or mechanism. In the case of verbs, however, the referential argument is always an eventuality. And a further difference between verbs and nouns is that whereas the typical noun has a referential argument but not others – relational nouns are something of an exception, not the rule – verbs almost always have at least one non-referential argument in addition to their referential argument; and many verbs – simply transitive verbs, ditransitive verbs, various verbs with prepositional arguments – have two non-referential arguments or more.[11]

[11] There is a small group of verbs that have a referential argument but no non-referential arguments. Among them are so-called 'weather verbs', like *rain* or *snow*. From a semantic point of view these seem to have only an event argument. But their atypicality is, you might say, 'grammaticized' in that in a well-formed sentence they still need a 'dummy subject', the *it* of 'It rained last night', 'It has been snowing' and so on. In English, apparently, the principle that verbs have subjects (and thus, on the present analysis, at least two arguments), has been built into its basic architecture.

It is a somewhat puzzling fact about the history of our subject that it should have taken so long for the distinction between referential and non-referential arguments of verbs to become noted. Even after it became widely accepted (for reasons having to do with tense and aspect) that verbs have eventuality arguments at the level of semantic representation, the implications of this for the syntax-semantics interface have often been ignored. The reasons, I suspect, may have to do on the one hand with the feature of the grammar of English and many other languages mentioned in the last footnote – that verbs need a subject even if it is just a dummy subject – and on the other with the dominance of non-relational nouns in the nominal domain and the rarity of 'non-relational verbs' (i.e. the weather verbs). These factors may have conspired in making us think that intransitive verbs are of the same semantic type as ordinary nouns. Traditionally these have been treated as creatures of the same logical type, viz. as 1-place predicates; and the difference between the ways in which the single arguments of these two kinds of predicates are actually realized – 'internally' in the case of ordinary nouns and 'externally' in the case of intransitive verbs – was somehow overlooked or ignored.

Another factor that may have contributed to this traditional perception is that nouns typically occur as constituents of DPs. It may seem quite natural to think of the referent of a DP as emerging only at the point where the DP is formed through the combination of NP and determiner. And this referent (or the dref representing it) is then at the same time the satisfier of the NP and therefore also of the NP's head noun. (This is in essence the way that quantified DPs were dealt with in our presentation of MG. In a DP like *every farmer* the quantificationally bound variable x in the logical form $\lambda Q.(\forall x)(\text{farmer'}(x) \rightarrow Q(x))$ of this DP is introduced by the Determiner, not by the NP or its head noun.) On this view NPs function as 'mere' predicates, which get all their arguments from elsewhere. Such an approach to the syntax-semantics interface of nouns and their projections militates against the assumption that nouns bring their arguments with them and so has nothing to suggest that something like this might be the case for verbs.

Once you start looking for support of the view that verbs have referential arguments (ant that these are always eventualities) you find it in several places. A particularly strong kind of evidence has to do with certain kinds of *deverbal nouns*. Many languages have ways of forming nouns out of verbs and among the different categories of deverbal nouns there are some that 'refer to' eventualities. In English many such nouns are *gerunds*, derived from their underlying verbs by suffixing *-ing* to the verb's stem. as in (3.31.a). But there are also many cases where the noun consists just of the verb stem, as in (3.31.b).

(3.31)

a. That morning he cleaned the bathroom. The cleaning was as thorough as any he could remember.

b. Suzie and Lara climbed Mt. Fuji. The climb took them a day.

In both (3.31.a) and (3.31.b) the deverbal noun occurs in the second sentence and the verb from which it is derived in the first. Both examples are instances of a frequently occurring pattern that for our present purposes is particularly telling. The DPs containing the deverbal nouns – *the cleaning* and *the climb* – seem to be 'anaphorically' related to what is said in the preceding sentences, by using the verbs that are their verbal bases. The cleaning referred to by the phrase *the cleaning* in the second sentence of (3.31.a) is the very cleaning spoken of in the first sentence. The same relation holds between the referent of *the climb* in the second sentence of (3.31.b) and the event described in the first sentence of that example. We often resort to deverbal nouns when we want to say more about an eventuality introduced in a previous sentence; a description formed from the name makes the eventuality available as referent for further predications by other verbs, prepositions and so on.

We have a direct and natural way of accounting for the 'anaphoric' relationships between the DPs *the cleaning* and *the climb* and the events described in the first sentences of (3.31.a,b), when we assume that the verbs in the first sentences have the same referential arguments as the nouns in the second sentences. In that case the semantic relations between the sentence pairs can be constructed along the lines of our reconstruction of anaphoric relations in our review of the top down treatment in DRT of donkey pronouns: The Universe of the DRS of the first sentence contains a dref e representing the event described by that sentence, and interpretation of the description in the second sentence then takes the form of setting the dref representing the description equal to e.[12]

[12] The details will of course have to be different from the treatment of anaphoric pronouns that we have been looking at, and this for two reasons. First, as noted, anaphora needs a very different kind of handling in a bottom up architecture, in which anaphoric noun phrases are treated as presupposition triggers; and second, the treatment of anaphoric definite descriptions isn't the same as that of pronouns. A crucial part of anaphora resolution for definite descriptions is that the dref chosen as antecedent can be shown to represent something that satisfies the descriptive content of the description (that is, its NP).

The distinction between referential and non-referential arguments is not only important for the interpretation of lexical predicates, but also at supra-lexical levels. One reason for this is that the nodes representing higher projections of verbs, such as VP or T', are interpreted as eventuality descriptions, and it is often important to know exactly which dref in the store of the representation attached to such a node represents the eventuality that is being described. That eventuality will always be the 'referential argument' of the eventuality description and so it is that eventuality which plays a special part in the execution of operators like the one triggered by the feature past (see (3.22)). In order that such an operator can apply correctly to a given input, it must be able to 'recognize' which dref in the store of the input representation is its referential argument. As a way to avoid possible indeterminacies in the application of operators like past it is therefore useful to single out the referential argument of the input description by giving it the subscript $_{ref}$. But we will use $_{ref}$ sparingly. We will use it to mark occurrences of referential arguments in stores, but won't attach the subscript to occurrences of the same dref elsewhere in the representation (i.e. in argument positions of DRS conditions in the DRS to the right of the store). We will also omit the subscript from the referential arguments of the Semantic Representations of lexical entries, since there the referential argument is always straightforwardly identifiable (viz. as the only dref in the store).

$_{ref}$-management isn't a completely trivial matter. One reason for this is that certain operators shift the role of referential argument from one dref to another. (We haven't seen an example of this yet – it does not arise in connection with the operator triggered by past – but the progressive operator discussed in the next section is such an example.) A further reason why a clear distinction between referential and non-referential arguments is needed is that verbs are not the only predicates with referential arguments. We have already seen that nouns have referential arguments too. (In fact, that was our starting point for the proposals we have made in this section.) Like verbs, nouns can be the lexical heads of complex phrases, and in the compositional construction of the semantic representations of the higher nominal projections shifts from one referential argument to another occasionally occur as well. Here too it is essential that the operator can identify the referential argument of the input representation. However, in what we are going to do in the remainder of these notes, these complications in the interpretation of complex nominal constituents won't play any role. So complications with $_{ref}$-management in the nominal domain won't be part of what follows.

Another lexical category that involves the distinction between referential and non-referential arguments is the category of prepositions. Prepositions are like relational nouns and intransitive verbs in that they have a referential argument and one non-referential argument. We will adapt our treatment of prepositions in DRS constructions accordingly.[13]

Many of the details of how $_{ref}$-bookkeeping works can't be shown at this early point. But some of them will become visible when we go through the DRS-construction for the sentence 'Frieda smiled' once more. We start with the representation for this sentence at the point where lexical insertion for *smile* and *Frieda* has just occurred. In order to mimic the selection of the correct $\overset{!}{\vee}$-disjunct from the entry for past the fact that the referential argument of *smile* is an event is represented once more in the 'canonical form of an eventuality dref ev together with the condition 'Event(ev)'.

[13]There are many cases where prepositions make common cause with some other predicate. An example is the presupposition *of* when it governs an adjunct to a relational noun, as in *friend of Mary.* Here *of* does not contribute a relation of its own, but rather is part of a complex morpho-syntactic construction that identifies the DP *Mary* as the non-referential argument of *friend.* The role played by *of* in this construction is comparable to that of a case marker, in that it helps to identify which argument slot of which predicate is the one for which the DP formally governed by the preposition is the filler. In fact, English, which has virtually no case marking left, still has the 'Saxon Genitive', expressed by *'s*, and this remnant case marker interacts with *of* on the one hand as competitor and on the other as a kind of work mate: *Mary's friend* is an alternative way of expressing the same thing as *(a/the) friend of Mary* – here the constructions involving *'s* and *of* seem to be competitors – but we can also combine *of* and Saxon genitive, as in *friend of Mary's*, and here they cooperate in identifying *Mary* as the non-referential argument of *friend.* (Quite a few English speakers seem to prefer *friend of Mary's* over *friend of Mary.* I do not know what to make of this fact – if indeed it is one.)

(3.32)

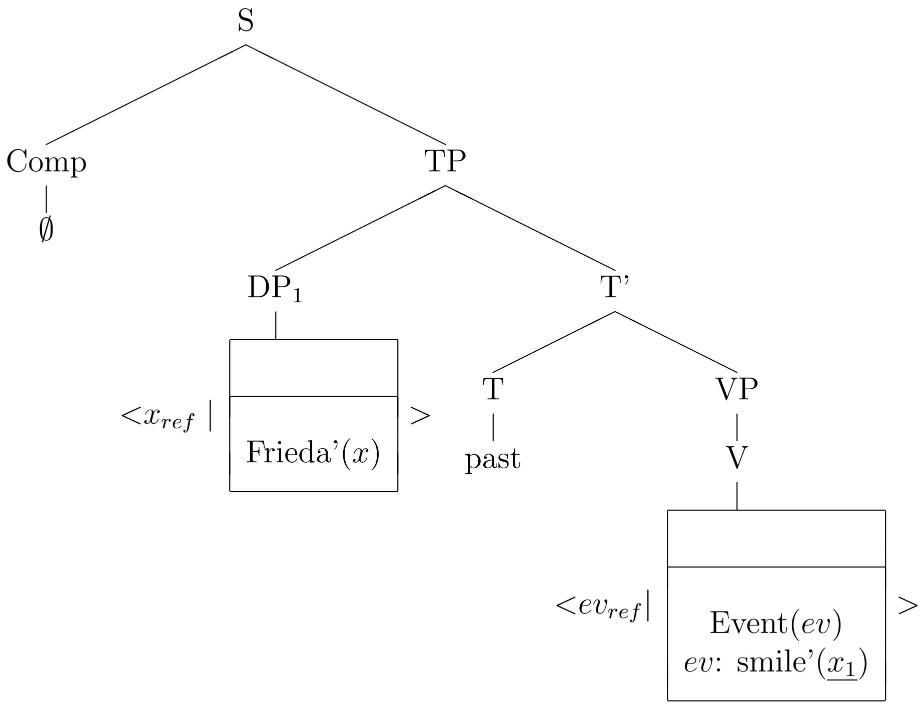

Combining this VP representation with the semantics of past leads in the first instance to the representation in (3.33) and after elimination of the inconsistent $\overset{!}{\vee}$-disjunct to the one in (3.34).

(3.33)

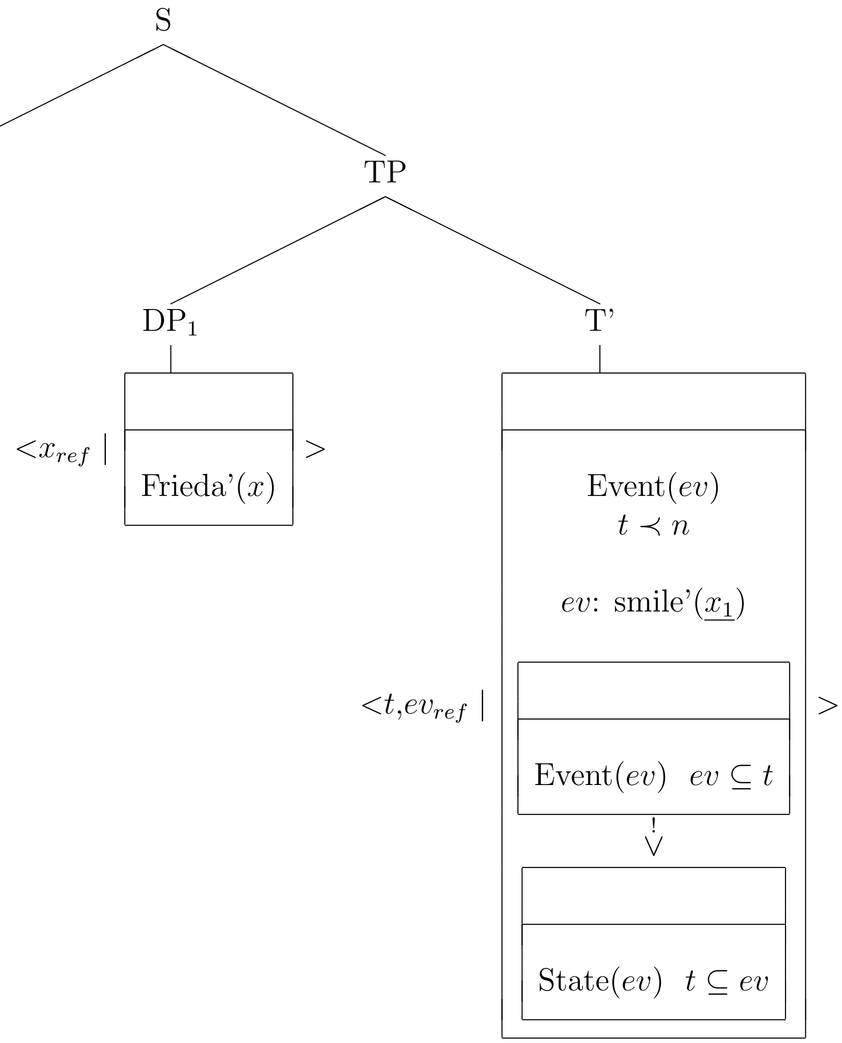

Note that once again the VP representation in (3.32) (which is identical to its V-represen-tation) is of the exact form specified in the entry for past. Second, after application of the past operator, leading to (3.33) and the simplification of (3.33) shown in (3.34), there is no longer any need for the subscript $_{ev}$. So the subscript could in principle be dropped in this particular DRS construction. However, in other constructions there still is a need for the subscript after this stage. So for reasons of uniformity we retain it also in (3.34); and it is also retained at the next stage, shown in (3.35), at which the T' representation has been combined with the subject representation. ((3.35) also shows a return to the more compact notation in which ev has been replaced by e and the condition 'Event(ev)' dropped.)

(3.34)

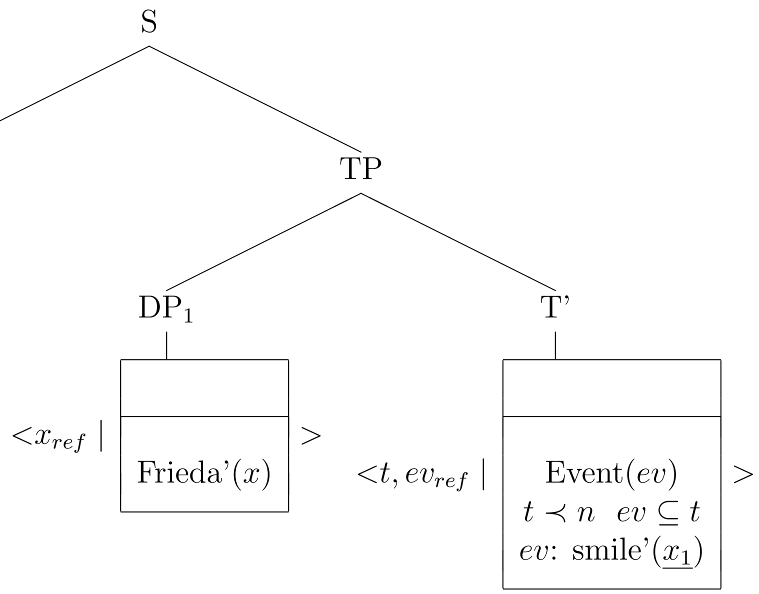

(3.35)

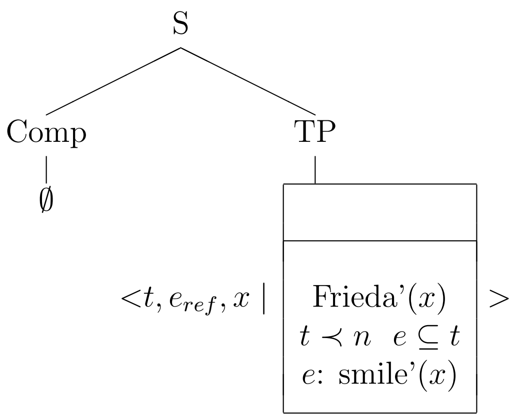

The transition from (3.34) to (3.35) concerns the combination of the representation of the T' node with that of the subject DP. Part of this transition is that after argument insertion (of the referential argument x of the DP *Frieda* into the argument slot $\underline{x}$ of *smile'*) the two representations are merged. This means forming a new store in which the drefs of the two stores are put together in a single list (as well as merging the two DRSs). In the assembling of the new store the referential argument x of the DP *Frieda* loses its referential argument status. The referential argument of the new representation – the dref which represents the bearer of the predicate represented by its DRS – is the referential argument e of the T' representation, not x. So as a member of the new store x is no longer $_{ref}$-marked.

In the last step of the DRS construction, in which the drefs that are still in store are transferred to the main Universe of the DRS to its right, the distinction between referential and non-referential argument status has become irrelevant. So any remaining $_{ref}$ subscripts are dropped. (3.36) shows the final DRS.

$$(3.36)\quad \begin{array}{|c|} \hline t \;\; e \;\; x \\ \hline \\ \text{Frieda'}(x) \\ t \prec n \;\; e \subseteq t \\ e\text{: smile'}(x) \\ \hline \end{array}$$

We conclude this section on the distinction between referential and non-referential arguments by returning to a point hinted at earlier. It has been argued in this section that the ways in which referential and non-referential arguments are realized in the grammar of English are strikingly different from what we find in the notation of predicate logic and other artificial languages of symbolic logic. And we surmised that all human languages are like English in this regard. In other words, this is an aspect of the syntax-semantics interface that has the status of a *linguistic universal.*

If the regime of predicate argument management that has been described in this section is a universal feature of human language, then it is one to which the literature appears to have paid little attention if any. The main reason may have been that for the most part linguistic typology has (at least until recently) not been high on the agenda of linguists with a formal logical background. It may need in depth 'brain washing' through an extensive symbiosis with formalisms like PC for someone to be struck by the ways in which referential and non-referential arguments are managed by natural languages along the lines this section has described.

With this sweeping hypothesis comes a plea. One way in which this entire section can and should be read is as a plea for information. Anyone who reads this and who is familiar with any of the (all - 4) languages that are beyond my personal horizon[14] is begged to come forward and make known to me wether the languages within their purview also differ from the architecture of PC in the way that this section has claimed for English.

[14]The 4 languages: Dutch, English, German and (with a question mark) French

3.5 Incorporating Adverbs and Aspect

3.5.1 Adding Adverbs

We have seen in Subsection 3.3 how the entry for past in (3.22) leads to the intuitively correct DRS for (3.10). One would hope that the same construction will also lead to the right result for the sentences in (3.20).

(3.20) a. At 18.00 Frieda closed the shop.
b. At 18.00 Frieda was closing the shop.

But as we have seen, there are two additional matters that we have to deal with in constructing semantic representations for these sentences, their temporal adverbs and the difference between the tense forms of the verb (Simple Past in (3.20.a) vs. Past Progressive in (3.20.b)).

We begin with a look at (3.20.a), whose only significant difference from (3.10) is that it contains a temporal adverb.

There are two types of questions that must be answered before we can extend our account of (3.10) to (3.20.a). The first type has to do with the syntax: How is the adverbial phrase *at 18.00* syntactically connected with the rest of the sentence?[15] The possible forms and points of adverb attachment are a notorious problem in English syntax; and this is true in particular for temporal adverbials. Temporal adverbials can occur at various points in well-formed sentences, and their surface positions impose constraints on where they may appear in the syntactic structures that we assume as LFs. But surface constraints do not fix the attachment points for adverbs completely. So the question where in our LFs – our starting points for the computation of semantic representations – they should belong still remains to be answered. For now we take some of the sting out of this problem by limiting our attention to sentence-initial occurrences of temporal adverbs, and of these we will assume that they are adjuncts to TP. (That assumption shouldn't be taken for granted either. For all we can tell on the basis of word order a sentence-initial adverb could just as easily be an adjunct to S; and there are further possibilities as well. A thorough investigation of all the possible options is possible only when other surface positions of temporal adverbs within the sentence string are taken into account, in particular those at the end of the

[15] N.B. I am treating prepositional phrases as a type of adverbials. Some may see this as requiring justification. But I do not believe that the matter is really controversial and so will not bother with this.

sentence. But adverb attachment is a topic beyond our current resources and we do best to limit ourselves to sentence-initial adverbs and make do with the assumption that they are adjuncts to TP.)

The second question concerns the syntactic structure of the adverbial itself. This question can be answered without further ado. The adverbial has the form of a Prepositional Phrase, consisting of the preposition *at* and the DP it governs; and the DP consists of just the expression *18.00*, which we will treat as a proper name.

When these assumptions about the form and attachment of the adverb are combined with the syntactic assumptions to which we are already committed, we are led to the LF for (3.20.a) that is displayed in (3.37).

(3.37)

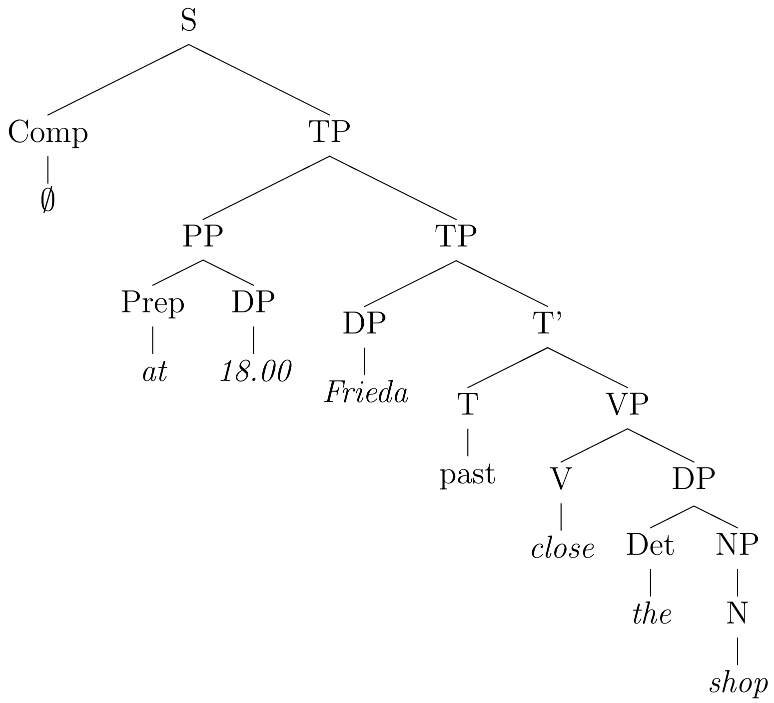

A further set of questions that need answering before we can carry out the DRS construction for this tree concerns the semantic contribution of the adverb. Now that we have decided on the place and the mode of syntactic attachment of sentence-initial temporal adverbs, we can bring these questions into clearer focus. Giving the definite description *the shop* an abridged treatment (consisting of a dref z and the surrogate condition "*the shop*(z)"), treating the verb *close* as a transitive event verb with the lexical entry in

(3.38) and otherwise proceeding as we did when constructing the representation for ‘Frieda smiled’, we get for the lower TP node of (3.37) the representation shown in (3.39).[16]

(3.38)

	close (V, trans.)	nom	acc
	e	$\underline{x}$	$\underline{y}$
Sel. Restr:	event	animate	
Sem.Repr:	$<e \mid \begin{array}{\|l\|}\hline \\ \hline e\text{: close'}(\underline{x},\underline{y}) \\ \hline \end{array} >$		

[16] One question that might have occurred to the reader is whether the dref z representing the shop should be put into the store or placed directly into the Universe of the outer DRS of the semantic representation of the VP node. This is a non-trivial matter, but it is one that will have to wait until we return to definite and indefinite noun phrases in Section 4. For now we will adopt the following policy: the drefs for proper names and definite descriptions are put into the store, with the eventual effect that they end up in the Universe of the main DRS (see below). Indefinite DPs present the complications that we encountered in our review of Top Down DRT in PART I. There is no reason for abandoning the position adopted there that indefinites can be analyzed either as quantifiers or as indefinite terms. When they are analyzed as quantifiers, then, as with all other quantifying DPs, their referential arguments get bound as part of the introduction of the duplex conditions to which quantifying DPs give rise. When they are analyzed as terms, then their referential arguments are first put in store, with the possibility that they will eventually be transferred to the Universe of the DRS following the store and sometimes to an even higher Universe. For more details see Section 3.9 and Section 4.

(3.39)

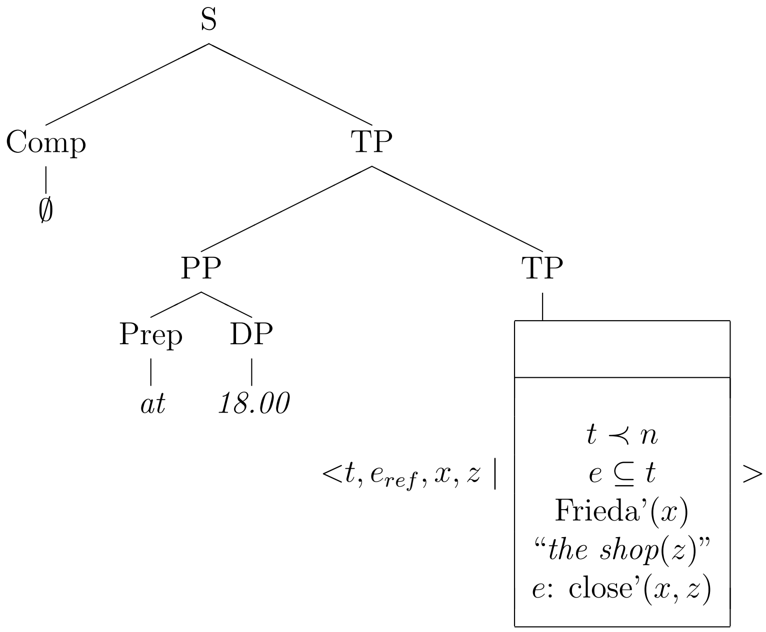

In order to determine what contribution is made by the adverb in (3.39) to the semantic representation of the upper TP-node we need to deal with the following problems. First, there is the problem we already noted about the denotation of the term *18.00*. If we make the prima facie plausible assumption that *18.00* denotes an instant of physical time, then it is impossible for the event of Frieda closing the shop to be temporally included within it. So we must assume that when used in a sentence like (3.20.a) *18.00* denotes some interval around this instant. This problem is related to an issue with the semantics of the temporal use of the preposition *at* (the use it has in (3.22)) and we will address it when we turn to the semantics of temporal *at* a couple of pages from here. A this point we just note for the record that when used in ordinary discourse clock terms like *18.00*, *six o'clock*, *18.01*and even *six o'clock sharp* allow for some denotational latitude: they aren't taken to denote instants of physical time, but 'blobs' of time, large enough for certain events of non-zero duration to fit within them.[17]

[17]We call terms like *18.00 calendar terms*. The semantics of calendar terms is based on our calendar, the system of conventions that we use to keep track of time by dividing it up into portions of various sizes. Calendar terms refer to such portions of time and they succeed in doing that because of their calendar-related syntactic and semantic structure. In fact, calendar terms form a kind of linguistic system, with its own syntax and semantics. In having its own syntax and semantics the system of our calendar terms is reminiscent of the system we use to name the natural numbers. Our conventional names for the natural

There is another problem with a DP like *18.00*: It will denote the clock time of six o'clock (am or pm). But which of the many times that fit this description will be denoted by a particular occurrence of *18.00*? Each day has one such time, so there are innumerably many such times (infinitely many, if we abstract from the physical limits imposed on the existence of our solar system). Evidently the way in which we tell which of those many times is denoted by a given occurrence of *18.00* depends on the context in which *18.00* is being used. Such context dependence affects most calendar terms. Only DPs like *at 18.00 on Tuesday February 3, 2015* escape this predicament; they refer to one and only one time. But the use of such fully explicit calendar DPs is cumbersome, so we try to make as much use of shorter phrases as we can, leaving it to context to settle what the phrase itself leaves open. Context-dependence of temporal adverbs is a topic on which we will have more to say later in Section 3.10. For now let us just assume that the DP picks out a unique time and that the information that determines which time that is is somehow packed into an 'ad hoc' DRS condition, which we will write as "*18.00*(t')".

Our next problem has to do with the semantic contribution of the preposition *at*. First a general remark about prepositions. As we noted in Section 3.4, prepositions are predicate words with a referential and one non-referential argument position. The non-referential position is filled by the prepositional complement DP of the Prepositional Phrase of which the preposition is the head – in (3.20.a) this is the DP *18.00* of the PP *at 18.00*. In keeping with our assumptions about nouns and verbs we will assume that the referential argument of a preposition is brought along by the preposition itself. Furthermore, when a PP containing the preposition is adjoined to some other phrase, its referential argument will be identified with the referential argument of that other phrase. We will return to this point below.

In addition there is a question about the particular preposition *at*, and more specifically about its use as a temporal preposition. When followed by a

numbers have semantically significant structure too: the form of a number name indicates where the number it refers to is positioned within the natural number sequence and for the most part it does so because of the general compositional principles that govern the form of these names. (This is true both of our decimal notation, but also of the words we use when we read decimal notation out; e.g. the internal structure of 'three hundred and twenty seven' mimics the internal structure of '327'.) Such naming systems are subsystems of the grammar of the language as a whole (in our case English) and a proper description of the syntax-semantics interface for the language should treat them as such. In these notes we do not develop a syntax-semantics interface for calendar terms. We will have more to say about such terms, however, in Section 3.10.

DP that denotes a time *at* expresses a temporal relation between its two arguments. And conversely, when it is used with its temporal sense *at* must govern a DP that denotes a time. Furthermore *at*'s referential argument must also be an entity of the kind that can stand in a temporal relation to a time. Thus *at* is subject to a combination of two selection restrictions, one for its non-referential and one for its referential argument.

I take it as beyond controversy that the only entities that can stand to each other in temporal relation to each other – relation such as temporal precedence, temporal inclusion or temporal overlap – are times and eventualities. So both the referential and the non-referential argument of temporal *at* must be entities belonging to on or the other of these two sorts. But in fat the selectional restrictions are stricter than this. First, as we already mor or lsee noted, the non-referential argument must be a time (and not an eventuality). second, the referential argument must be an eventuality. (No time can be said to be 'at' some other time.). And finally there is an additional restriction on the non-referential argument. Not only must this argument be a time; in fact it must be, in some appropriate sense, a temporal *point*, and not a temporal *interval*. (This is the respect in which it differs from the temporal prepositions *on* and *in*, which require that their non-referential argument *not* be a temporal point.) The sense in which the non-referential argument of *at* must be a 'point' is not so easy to articulate. Clearly the sense cannot be that of an indivisible, duration-less instant of physical time (a real number when, standardly, physical time is identifies with the on the real number line). For no ordinary event could temporally fit within a point of this sort, and certainly not a complex event like that of Frieda closing her shop. A point in time that can temporally contain such events must have substantial temporal extension. I do not know of a fully satisfactory story of what such 'extensive points' can be. But the only plausible partial story that I am familiar with derives the availability of such points from what is known as the conceptual *granularity* of time: When thinking and speaking about time we typically impose grids on it of varying coarseness or fineness – grids of varying 'granularity', as the official terminology has it. Any given granularity grid partitions time into minimal ('atomic') portions of that granularity; these minimal portions will be bigger or smaller depending on what granularity is chosen. Once a granularity has been fixed, it is these 'atomic' portions that play the part of 'temporal points'.

That granularity is somehow determined or influenced by context seems intuitively clear. But how it is determined by context is not as yet very well understood. There is one factor, however, that evidently plays an important part. The denotations of the non-referential argument phrases of temporal *at*, we saw, must be points, in the sense we are trying to identify. Often these phrases are calendar terms, and that is so in particular for the DP governed by *at* in (3.20.a). And when the DP chosen is a calendar term, then the particular choice of calendar term can be a decisive determinant of what granularity is assumed. The way on which calendar terms influence granularity is sometimes referred to as the 'round numbers effect'. *18.00* and *six o'clock* are 'round number terms' in a sense in which *18.01* and *one minute past six o'clock* are not. And *one minute and thirty seven seconds past six o'clock* is even 'less round'. Round number terms imply a coarser granularity than other calendar terms. For instance we may be willing to accept (3.20.a) as true even if Frieda started with the things she had to do to close the shop three minutes before six and completed the procedure only three minutes after six. Those six minutes might still count as a point according to the granularity that the use of *at 18.00* suggests. But to describe this event as 'At 18.01' frieda closed the shop' seems much less felicitous (and perhaps outright false) because the granularity suggested by *at 18.01* is significantly finer: the denotation of *at 18.01* won't extend to three minutes past six even though 18.03 is closer to 18.01 than it is to 18.00 when we think of these terms as denoting physical instants. For the 'round numbers effect and other questions relating to granularity see for instance (Lasersohn 2003)and (Krifka 2002).

Granularity is a topic that belongs to the theory of vagueness and imprecision. Earlier in these notes the decision was made to set those topics aside. That decision applies to granularity as well. No more will be said about it from here on.

The final question about *at* is how exactly we should understand the 'simultaneity' that it is supposed to express. Here we once more encounter the problem of perfective vs. imperfective aspect. When the representation of the adjunction site of an *at*-PP is the description of an event, then the temporal relation should be that of the described event being temporally contained within the non-referential argument of *at*; when it is the description of a state, then the temporal inclusion should be the other way round. In our formalization of the contribution that temporal *at*-PPs make to the semantics of the sentences in which they occur, we will make use of the lexical entry for temporal *at* shown in (3.40), in which the considerations and decisions of these last few pages are reflected.[18]

[18]We are making no effort here to relate the use of *at* that is identified in (3.40) to

(3.40)

	at (Prep) ev	acc $\underline{t}$
Sel. Restr:	eventuality	temporal point

Sem.Repr:

$$<ev_{ref} \mid \begin{array}{|c|}\hline \\ \hline \begin{array}{|c|}\hline \\ \hline \text{Event}(ev) \quad ev \subseteq \underline{t} \\ \hline \end{array} \\ \Downarrow \\ \begin{array}{|c|}\hline \\ \hline \text{State}(ev) \quad \underline{t} \subseteq ev \\ \hline \end{array} \\ \hline \end{array} >$$

We are now ready to combine the semantics of *at* with the semantic representation of *18.00*. This combination follows the by now familiar pattern of argument insertion and is no different from the way in which the DP *the shop* is combined with the transitive verb *close*, or the subject DP *Frieda* with the semantic representation of the T'-node in (3.39). There is one minor (and merely notational) difference with the cases considered so far. It has to do with the argument position into which the dref for the DP is to be inserted. For dref insertions into the non-referential argument positions of verbs we continue to rely on coindexations provided by the syntactic parser that our construction algorithm presupposes. We could adopt a similar coindexation strategy for the non-referential arguments of prepositions. But in this case

non-temporal uses of *at*. It is plain that there are non-accidental connections between the temporal use of *at* and its spatial use (as in: 'He can be found at 528 Elm Street'). (Spatial and temporal uses that are more or less systematically related to each other can be observed for a number of prepositions.) However, to capture these connections in a way that is conceptually insightful as well as extensionally correct has proved to be difficult. This is a challenge for a general theory of spatial and temporal prepositions. No such theory will be developed in these notes.

there is no real need for coindexation, since the non-referential argument slot of a preposition is *always* filled by the referential argument for the DP that is directly governed by the preposition.

The result of combining the semantics of *at* with that of *18.00* is given in (3.41).

(3.41)

$<t', ev'_{ref} \mid$

"18.00(t')"

Event(ev') $ev' \subseteq t'$

$\Downarrow$

State(ev') $t' \subseteq ev'$

$>$

Insertion of (3.41) for the PP constituent of (3.39) gives us (3.42).

(3.42)

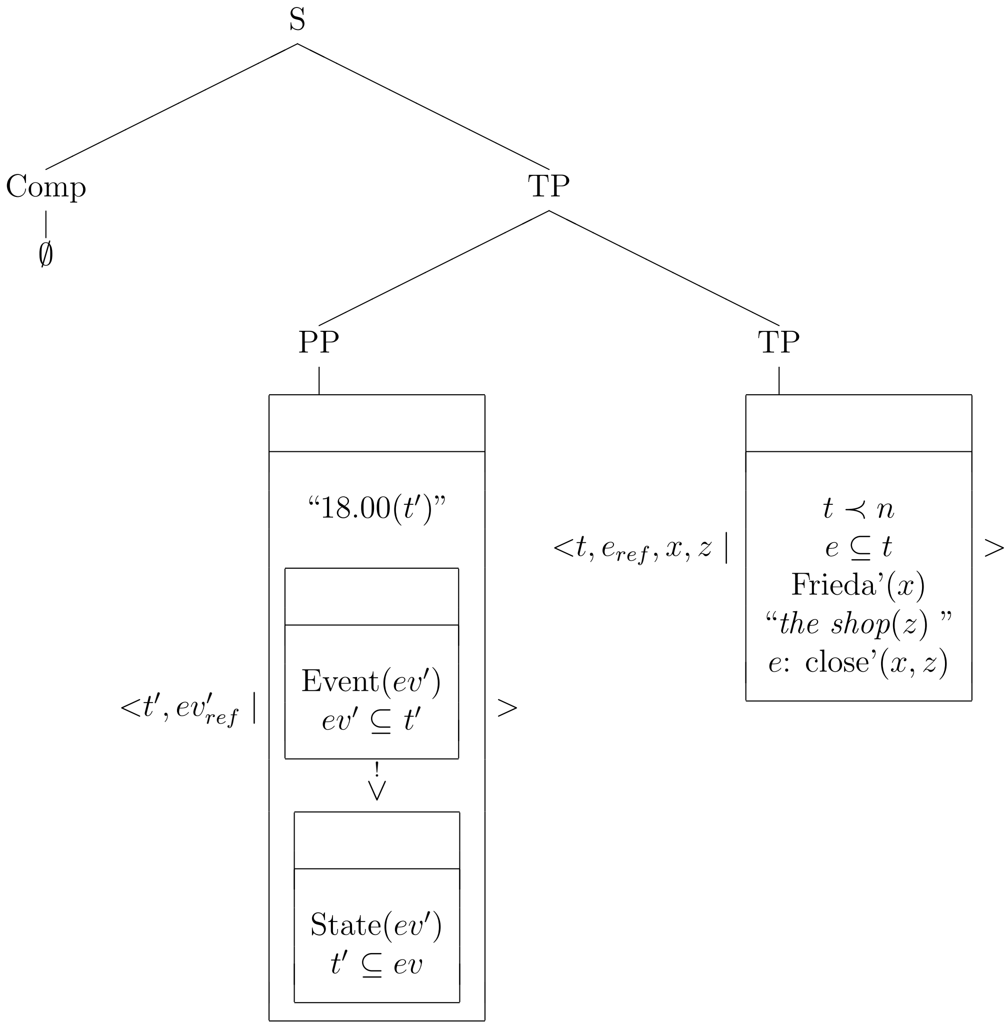

The construction step to be performed on (3.42) is an instance of executing the semantics of adjunction. Given the commitments that our present syntax-semantics interface has already made, the semantics of adjunction emerges as a form of *unification*: the identification of two drefs that each come with their own predications. As a result of this operation the predications involving the dref that plays the part of referential argument of the adjunct become predications involving the referential argument of the adjunction site.

In order for the unification operation to be well-defined, it must be clear which drefs from the respective stores are to be chosen as the two unificanda. This is one point where the algorithm relies on referential argument marking: the two drefs that are to be unified are the ones singled out (by virtue of bearing the subscript $_{ref}$, as referential arguments of their respective representations).

The execution of unification that we adopt here is in the spirit of the procedure we adopted in PART I for dealing with pronominal anaphora: We set the two referential arguments equal by adding an equation, in which the two drefs are connected by '=', to the Condition Set of the lowest DRS from which both drefs are accessible. After that the two representations are merged, in the by now familiar way: the stores are combined into a single list (in which the adjunct store appended to the end of the store of the adjunction site), and the two DRSs are merged into a new DRS.

In the case we are considering the unification triggers a further step, in which one of the $\overset{!}{\vee}$-disjuncts is selected. Here we find ourselves in the same situation as we were when combining the VP representation with the semantics of the feature past: one of the disjuncts is compatible with the condition 'Event(ev'_{ref})' and the other one is not; the compatible one gets selected.

In practice we will often simplify the representations that result from adjunct unification by substituting the referential argument of the adjunction site for all the argument occurrences of the referential argument of the adjunct and then eliminating all remaining occurrences of the referential argument of the adjunct, as well as the equality condition (which at this point has become a tautology of the form 'x = x', and thus redundant).

(3.43.a) gives the initial result of dealing with the adjunction in (3.42) and (3.43.b) the result of $\overset{!}{\vee}$-elimination. Eliminating the referential argument of the adjunct in favor of the referential argument of the adjunction site then leads from (3.43.b) to (3.43c).

(3.43)

a.

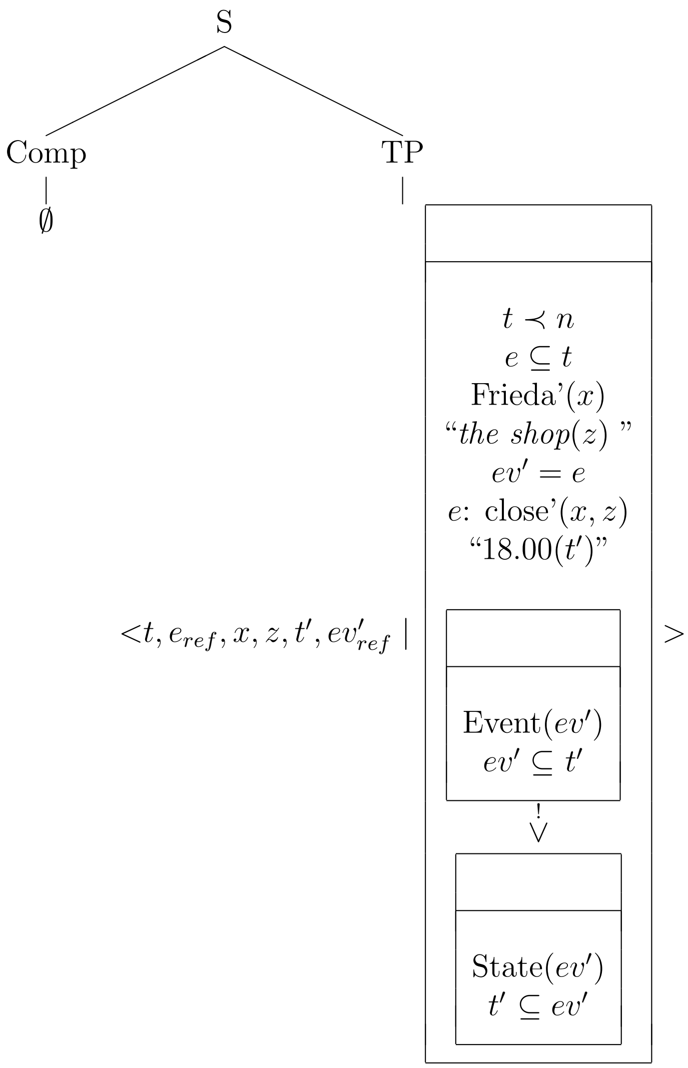
S
Comp
TP
$\emptyset$
$t \prec n$
$e \subseteq t$
Frieda'(x)
"*the shop*(z) "
$ev' = e$
e: close'(x, z)
"18.00(t')"
$<t, e_{ref}, x, z, t', ev'_{ref}$ |
Event(ev')
$ev' \subseteq t'$
State(ev')
$t' \subseteq ev'$
$>$

b.

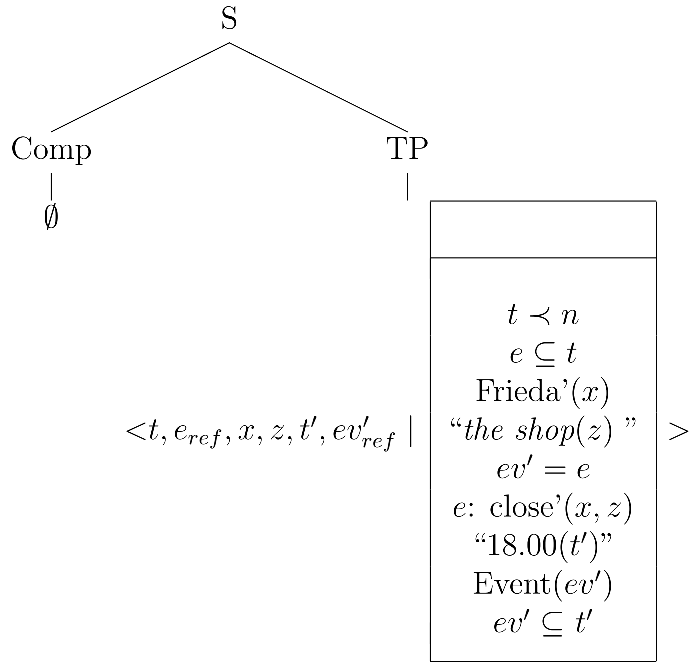

c.

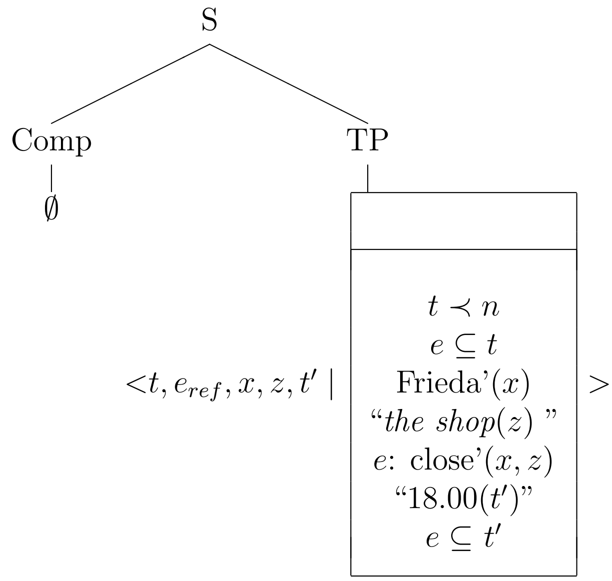

The final DRS can then be obtained from (3.43.c) by transferring the drefs from the store into the Universe of the DRS.

3.5.2 Adding Aspect

In the last Section we dealt with one of the problems presented by the sentence pair in (3.20) (repeated here for easier reference), viz the temporal adjunct that is part of each of them. In this section we address the respect in which (3.20.a) and (3.20.b) differ: their respective aspects.

(3.20) a. At 18.00 Frieda closed the shop.
b. At 18.00 Frieda was closing the shop.

In the DRS construction for (3.20.a) we focused on the treatment of the advertbial adjunct*18.00*. In the course of that we also dealt, implicitly, with the perfective aspect that distinguishes (3.20.a) from (3.20.b). (The event described by (3.20.a) was represented as temporally included within the time t' representing the denotation of *18.00*.) In the present section the primary focus will be on (3.20.b). But before we can start with the DRS construction for this sentence we first have to address the question how the LF for (3.20.b) differs from that of (3.20.a).

Our informal discussion of the sentences in (3.20) showed that the aspectual distinction between them was crucial for the interpretation of how their eventualities are temporally related to the time denoted by *18.00*. In (3.20.a) the event e of the representation of the lower TP node (i.e. the adjunction site for the temporal PP) is temporally included within the time denoted by *18.00*; in (3.20.b) the time denoted by *18.00* is included within the duration of the state described by the lower TP representation.

If this is the way in which the difference in meaning between the two sentences is to be captured, and if we also assume that the representation for (3.20.b) is to be constructed on the basis of the lexical entry for *close* given in (3.38), according to which *close* is an event verb, then there must be a switch from event description to state description somewhere along the path from lexical insertion of the verb to the point where the adverb is adjoined to TP. In fact, if we want to stick to our commitment that the eventuality described by the sister to T interacts with the tense feature specified at T in the same way that it interacts with the time introduced by the temporal adverb, then the transition should take place below the T' node.

The question what the precise point is at which the transition from event description to state description occurs is connected with what syntactic structure we should assume as input to the semantic representation for (3.20.b). This is just a special case of the more general question what syntax we should assume for sentences containing progressive morphology.

In principle there are two ways in which one could approach the difference between progressive and non-progressive forms. The first is to treat progressive forms as distinct tenses. On this approach there would be a 'past progressive tense', in addition to the Simple Past, a future progressive tense in addition to the Simple Future, and so on. And each of these progressive tense forms should be given its own lexical semantics.

There is no question that this could be made to work; but there is much that speaks against it. The most obvious objection comes from morphology. The progressive tense forms of English stand in a 1-to-1 correspondence to the non-progressive forms – for each non-progressive form there is a corresponding progressive one. This strongly suggests that the progressive form and the tenses are distinct and independent modifiers of verb stems. Furthermore it is easily seen that this morphological independence of the progressive from the different tenses is confirmed by the semantics. The semantic differences that can be observed between the simple past and the past progressive are replicated by the difference between simple future and future progressive, and likewise for other tense forms. All this suggests that the progressive should be treated as an operator that is distinct from, and independent from, the operators denoted by the tenses. And that means that the formation of a sentence like (3.20.b) involves two operations, one that transforms the verb into a progressive form and one that adds the tense.

The next question is which of these operations comes first, but here too morphology more or less dictates the answer. Forming the progressive of a verb consists in forming its gerundive and making that the complement of the auxiliary verb *be*, and then it is *be* which gets modified by finite tense. Clearly, this is the order in which a verb with a 'progressive tense' gets dressed up in its morphology. For a generative syntax like the one we are using the implication is that the progressive is introduced at a lower level than the tense.

The level at which the progressive is introduced is often referred to as the projection level of 'Aspect', and that is the term we will use too. We will represent this level as involving two nodes: (i) a functional head, with label 'Asp', which carries the information that is conveyed by the presence or absence of progressive morphology, and (ii) a node that is the mother to this node (as well as to the VP, her second daughter) and which has the label 'AspP' (for 'Aspect Phrase').

Thus the syntactic structure for (3.20.b) is as in (3.44). The information that the verb of the sentence has a progressive form is encoded in a feature with the name '+prog'.

(3.44)

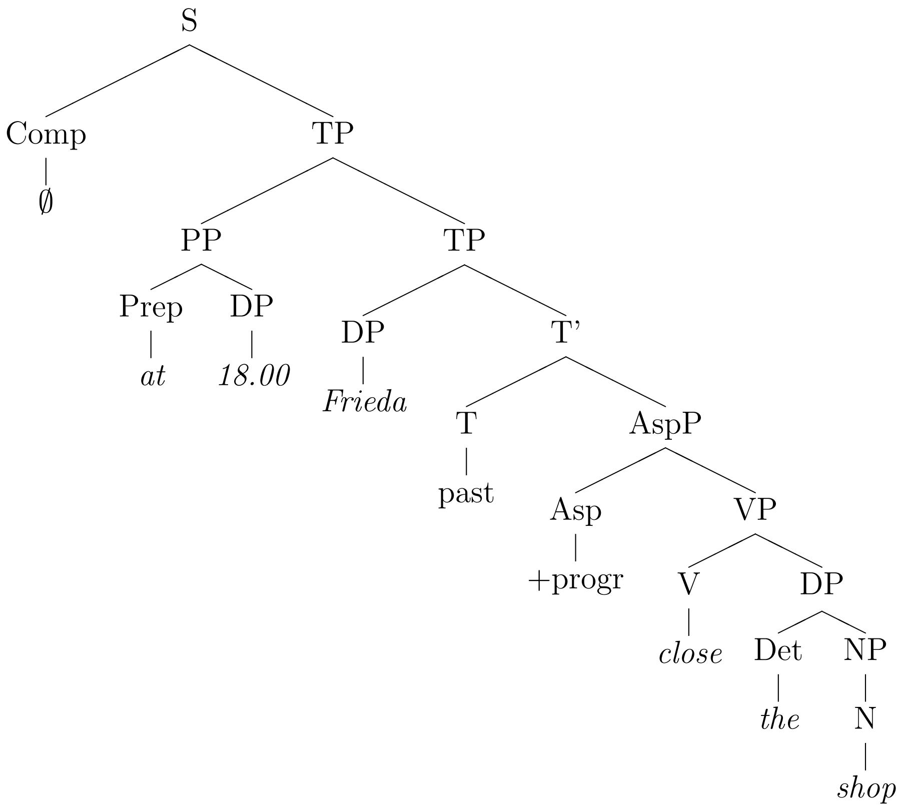

Now that we have introduced the Aspect level as a separate projection level, this also raises a question about the syntactic structure for non-progressive sentences like (3.20.a). Symmetry would seem to suggest that the syntactic structure of such sentences also contains the new level. After all, the information that the verb of the sentence is not in the progressive form is information no less than the information that its form is progressive. From this point of view the syntactic structure of (3.20.a) should be as in (3.45), with '-prog' the obvious name for the feature which conveys that there is no progressive.

(3.45)

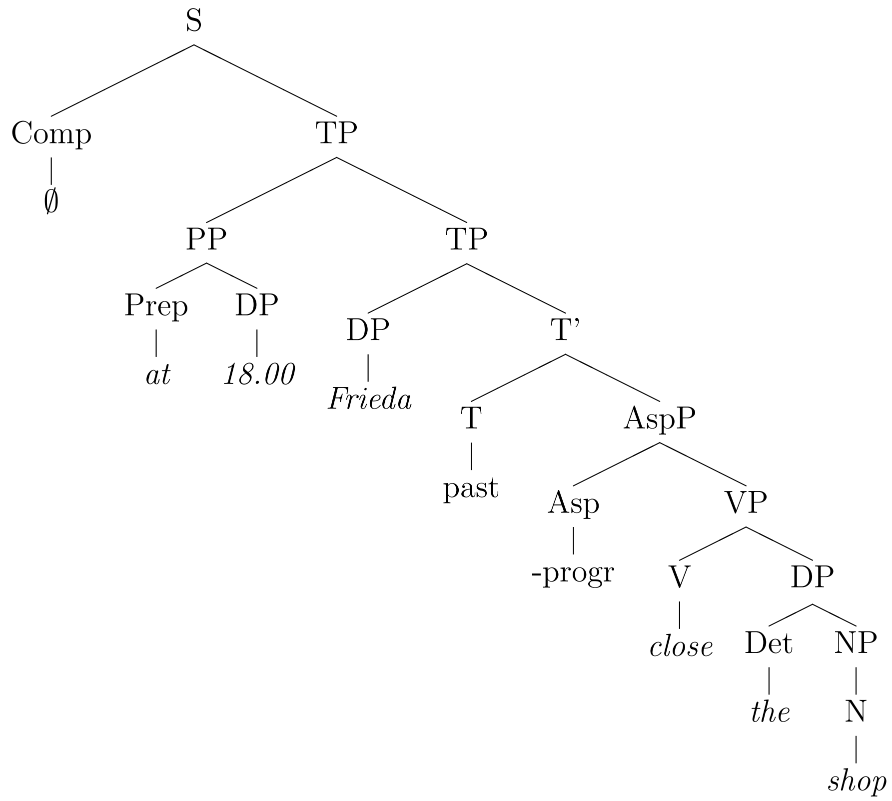

From a semantic point of view, however, this symmetry is artificial. Putting a verb into its progressive form has semantic consequences – as we will argue below, the meanings of the progressive forms of verbs are not the same as those of their non-progressive forms (an observation that should be obvious in any case, just from relying on your speaker's intuitions when comparing the progressive and non-progressive forms of run of the mill verbs). But *not* putting the verb into the progressive obviously is just a way of *not* changing its meaning. So for a sentence like (3.20.a) the Asp projection level is semantically redundant; for such sentences the semantic representation of the sister to Asp, the VP node, is passed up unaltered to its mother node AspP. Because of this the computation of the semantics will lead to exactly the same sentence representation whether the Asp projection level is included in the LF or whether it is left out: for sentences in which the Asp feature is -prog the Aspect projection level is otiose; it can be omitted without any change to the semantics. And so we will normally leave it out, unless there is a special reason for displaying it explicitly.

<u>Exercise</u> Construct the DRS for (3.20.a) on the basis of (3.45) as LF.

Now that the difference between (3.20.b) and (3.20.a) has been reduced to that between the features +prog and -prog, it seems clear where the decisive transition from event description to state description should take place. The element responsible for this change must be the feature +prog and the transition should occur as part of the transition from the semantic representation of the complement to Asp (i.e. the VP node) to the representation of the AspP node. But if +prog in (3.44) triggers the transition from an event description to a state description, how exactly are these two representations related? Or, to ask this question in a less theoretically loaded way, how does the meaning of the progressive, e.g. that of *be closing the shop*, relate to the meaning of the corresponding non-progressive *close the shop*?

Here is a first shot at a characterization of this relation: For an event verb like *close* the progressive form *be closing* describes a process that ideally should lead to the completion of an event of the kind described by the non-progressivized *close*. But what exactly should we mean by 'ideally should lead to'? That is a hard question, which has preoccupied semanticists since the early days of Formal Semantics (and probably before that). A glimpse of why it is hard might have been caught from our informal description of the difference between (3.20.a) and (3.20.b): (3.20.b) can be true even in situations where the closing isn't completed, for instance because the owner of the shop was called away in the middle of the process and then was unable to return to her shop until the next day. In such a situation there is no event that could have been described by the non-progressive 'x closed the shop', so we cannot characterize the process described by *was closing* simply as one part of an (actual) event of complete shop closing; at best the process could be characterized as part of an incipient shop closing, a closing that was 'on the cards'; but merely incipient events and events 'on the cards' aren't real events.

Much effort has been expended on trying to find a viable formal characterization of this intuitive idea: that the processes described by progressive forms of event verbs are parts of what might have become events describable by the non-progressive forms of those event verbs – provided things had turned out otherwise than they did. (E.g. if the shop keeper hadn't been called away while she was in the process of closing up). But this has proved to be surprisingly difficult, perhaps because too often those who tried had their hands tied by the formal framework within which they were working. (See in particular the attempts of Dowty (Dowty 1979) and Landman (Landman 1989*b*) to account for the semantics of progressives in terms of 'inertia worlds'. For another approach, involving non-monotonic logic, see (Asher 1992).)

We will not try to solve this problem here. (But see the subsection headed 'The Meaning of PROG' following diagram (3.43).) Our only concern will be to set things up in such a way that a more detailed account of the semantic relations between progressive and non-progressive forms can be fitted in without major structural changes. The crucial point here is one of which we will find other instances as we proceed: While it isn't possible to characterize the individual processes described by progressive forms in terms of *individual* events described by the non-progressive forms, it should be possible to describe the eventuality *predicate* expressed by the progressive form of an event verb as a function of the event predicate expressed by the non-progressive form of the verb. Or, putting the point in more overtly semantic terms, the *property* expressed by the progressive predicate should be determined by the *property* expressed by the non-progressive form.

But what in general is the property expressed by a verbal predicate (or, for that matter, the property expressed by any other predicate, from English or any other natural language)? The answer to this question that we adopt is representative of the *intensional* stance that has been standard within formal semantics since the work of Carnap, Kripke, Montague and others in the forties, fifties and sixties: the property expressed by a predicate P is a function that is defined on the set of possible worlds and that assigns to each possible world w the *extension* of P in w. (The *extension* of P in w is the set of all entities that satisfy P in w.) We denote this function as $^\wedge x.P(x)$. $^\wedge$ is the intensional variable binder familiar from Montague Grammar: '$^\wedge x.P(x)$' is to be read as 'the property of being an x that satisfies P'.

In the present context our concern is with properties of eventualities – properties whose extensions in the different possible worlds are always sets that have eventualities as their members. A term denoting such a property will take the form $^\wedge ev.P(ev)$, where ev is any eventuality dref (or, more specifically, $^\wedge e.P(e)$ when P is a predicate of events or $^\wedge s.P(s)$ when P is a predicate of states).

Although there is little we will have to say about what the progressive operator actually does to a property input of the general form $^\wedge ev.P(ev)$, we still need a name for it. We will use the name 'PROG'. In our representational framework 'PROG' will occur only as part of conditions in which it is applied to representations of eventuality properties and these will always be obtained by applying the operator ^ to representations of VP nodes. Exactly what that comes to is best shown by means of an example. But before we do so, there is one further point to be mentioned first. The progressive is an 'eventuality transformer'; both its input and its output are eventuality descriptions. However, they are descriptions with opposite aspectual properties. The natural inputs to the progressive are eventuality descriptions with perfective aspect – more about that below – and the outputs are eventuality descriptions with imperfective aspect. For us, who are identifying perfective eventuality descriptions with event descriptions and imperfective descriptions with state descriptions, this means that the inputs to the operators should be event descriptions and the outputs state descriptions. The form this will take is that the operation triggered by the feature +prog binds the event dref that is the referential argument of the input description and introduces a new state dref as referential argument of the output description. Binding the event argument involves taking it out of the store and putting it into the Universe of the main DRS that follows the store[19]; and introducing the new state dref involves adding it to the store, where it is appended to its front and marked as the referential argument of the output representation through the subscript $_{ref}$.

To see what this comes to let us have a look at the transition step in the DRS construction of (3.20.b) from the stage just before application of the operation triggered by +prog to the stage that results from this operation. The two stages are shown in (3.46) and (3.47).

[19] Insertion of the event dref, which is now bound by ^, into the Universe of the following DRS creates the same kind of situation that we also have in duplex conditions, in which the dref bound by the quantifier occurs both as argument of the quantifier (within the central diamond of the duplex condition to which the quantifier gives rise) and in the Universe of the restrictor DRS. Just as there, it looks as if the dref is bound twice over, once as 'bound variable' of the operator ^ and once through its occurrence in a DRS Universe. However, as in the case of duplex conditions, insertion of the dref into the DRS Universe makes no difference to the resulting truth conditions. In both cases – duplex conditions and the output representations of +prog – the semantically relevant binding is performed by the operator (i.e. the quantifier or ^). Adding the dref also to the DRS Universe in question makes no difference to the truth conditions of the DRS. Its only effect is that it somewhat simplifies the definition of dref accessibility, and it makes the accessibility of the given dref visible in an intuitive and easily surveyable form.

(3.46)

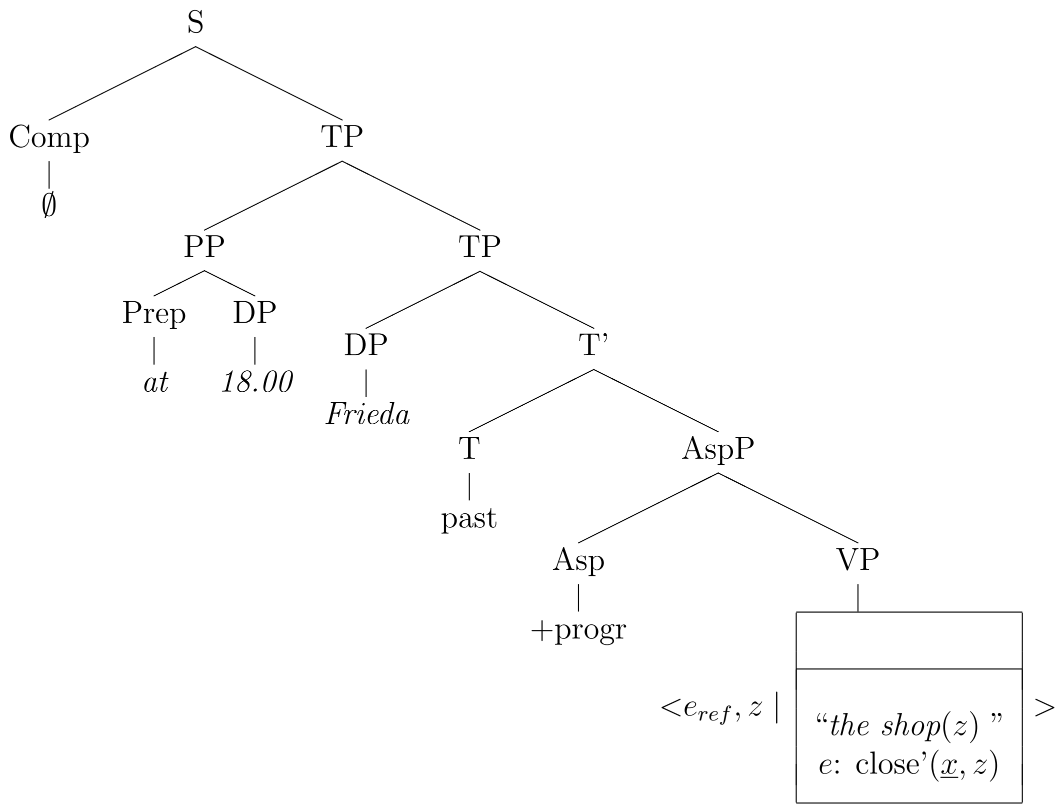
S
Comp
∅
TP
PP
Prep
at
DP
18.00
TP
DP
Frieda
T'
T
past
AspP
Asp
+progr
VP
<e_ref, z |
"the shop(z)"
e: close'(x, z)
>

(3.47)

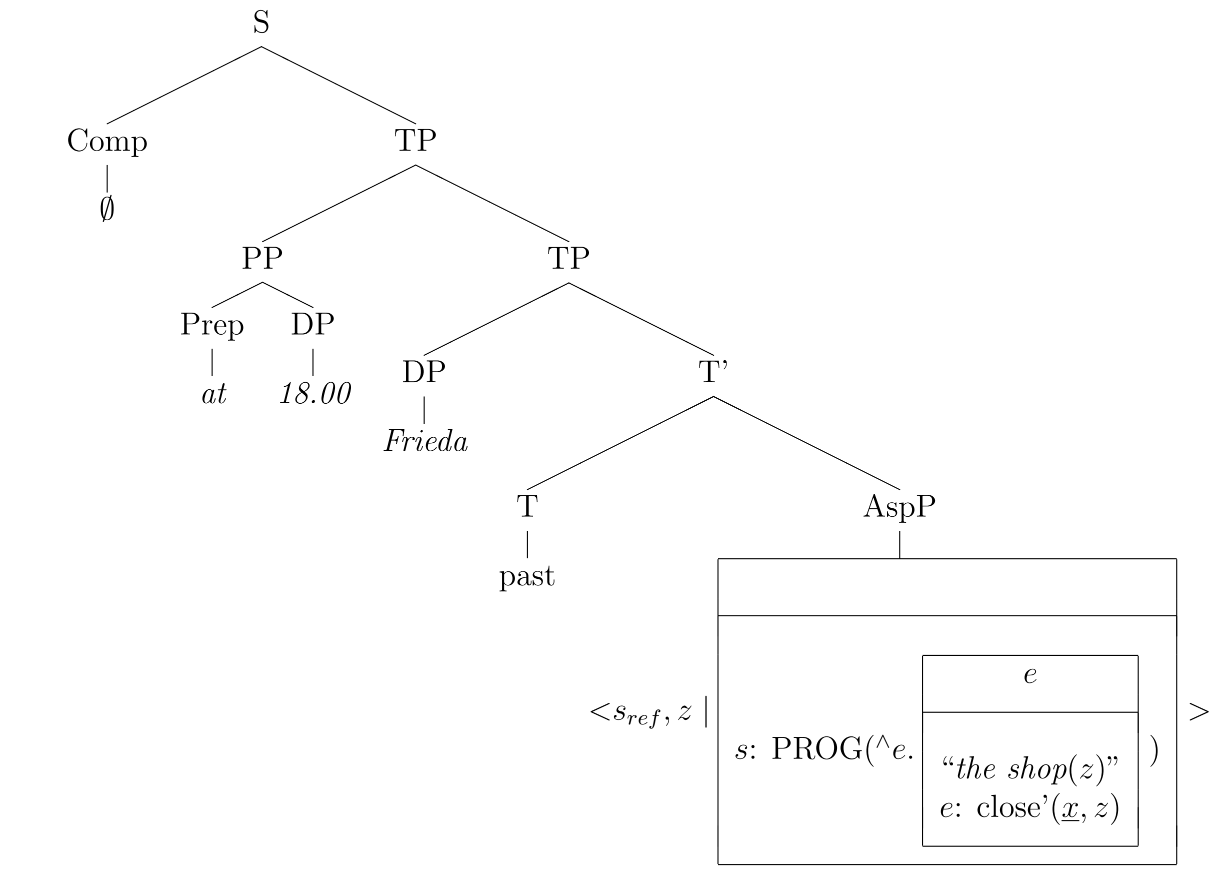

(N.B. Note the shift of the subscript $_{ref}$ from e to s.)

There is much that still needs to be said about the semantics of the operator PROG. Little of that will be said in these notes. There are some things that can and will said right here. But let us first complete the construction of which (3.48) is an intermediate stage.

The next construction step combines the AspP representation with the lexical semantics for the feature past. In the present case application of the semantics for past to the input representation, which is now a state representation, leads to selection of the second disjunct of the $\overset{!}{\vee}$ disjunction in the lexical entry for past. So the result of the next operation is the structure in (3.48).

(3.48)

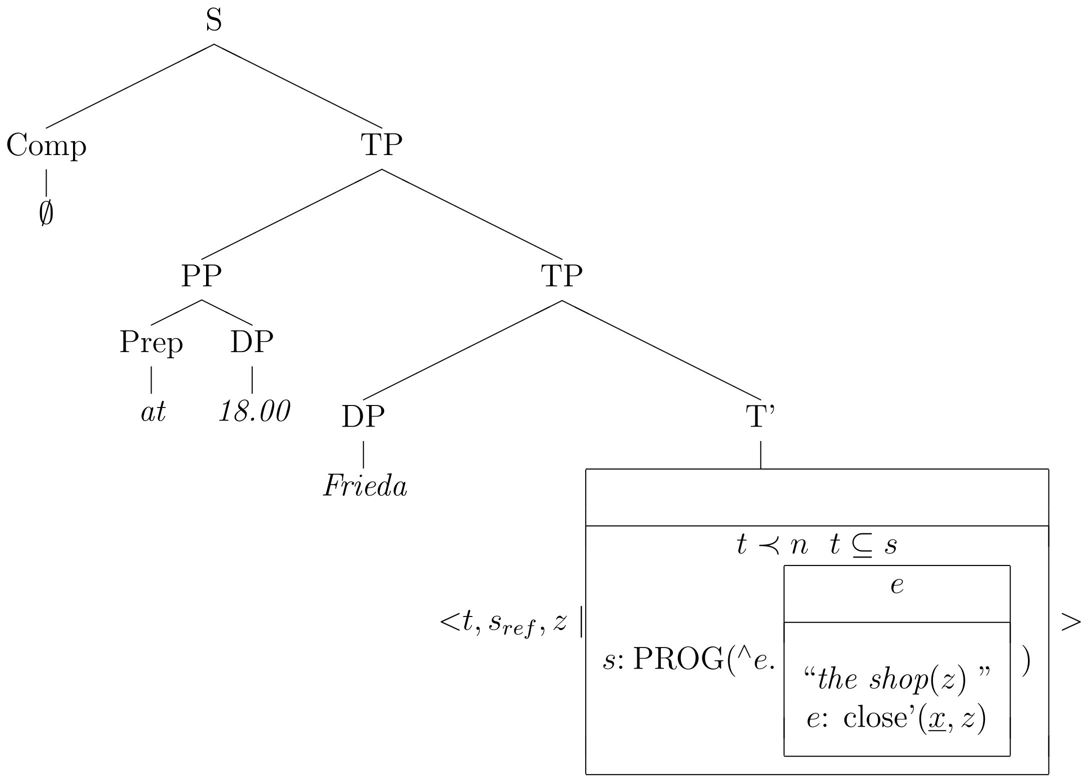

The remaining construction steps are the same as in the DRS construction for (3.20.b) (see (3.43.b,c)), except that when the semantic representation of the PP is applied to the representation of the lower TP, then it is once more the state disjunct of the $\overset{!}{\vee}$ disjunction that gets selected. This time we show, for a change, the final result of the construction, the DRS in (3.49).

(3.49)

t' t s x z

Frieda(x) $\grave{1}$8.00'(t')

$t \prec n$ $t \subseteq s$ $t' \subseteq s$

s: PROG(^e. [e | "*the shop*(z) " e: close'(x,z)])

The meaning of PROG

We now turn to the question that we saved for after completion of the DRS for (3.20.b). The representation in (3.47) that results from applying the progressive operator to the VP representation shown in (3.46) tells us next to nothing about how the second representation is related to the first. A semantic account of the progressive ought to have more to say about this. But what?

We already drew attention to this question when we noted that (3.20.b) can be said truly by a passer-by who didn't wait to see whether the closing of the shop was completed; and his statement would be no less true if it turned out that in the end Frieda didn't finish closing the shop. In such a case the statement with the progressive tense form would be true while (3.20.b) would be false. This phenomenon, that a statement involving the progressive form can be true while the corresponding non-progressive form is false, is part of what is known as the *Imperfective Paradox*. When this phenomenon was first noted in the literature, it was described as a paradox because the discrepancy – progressive form true, non-progressive form false – is attested for some VPs but not for others. Compare for instance the sentence pair we have been discussing, repeated once more below as the pair (3.50.c,d), with the pair (3.50.a,b). It was noted that the sentence (3.50.a) entails (3.50.b). But as we have just seen, such an entailment relation does not hold between (3.50.c) and (3.50.d), as it is possible for (3.50.c) to be true and at the same time for (3.50.d) to be false.

(3.50)a. Frieda was singing.

b. Frieda sang.

c. At 18.00 Frieda was closing the shop.

d. At 18.00 Frieda closed the shop.

The difference between the case presented by (3.50.a) and (3.50.b) and that presented by (3.50.c) and (3.50.d) has to do with the aspectual properties of their verbs and verb phrases. The difference is that between *telic* and *non-telic* verbs and verb phrases. Telic verbs/verb phrases describe events with an intrinsic end point, its *culmination point*. Every event satisfying the description must terminate in such a culmination point and comes to an end when and only when that point is reached. Events described by non-telic verbs and verb phrases do not have a culmination point. Take for instance the verb *sing* of (3.50.a,b). Suppose e is an event of Frieda singing. This event will come to an end at some point – because Frieda thinks that enough is enough, or because someone tells her to stop, or because she is knocked down by an irate member of the audience or whatever. Hadn't it been for her decision to stop, or for the interfering cause, she could have carried on singing. In such a possible continuation of the actual world the singing event would have lasted longer; but that doesn't make the shorter event that actually occurred any less of a singing event – one that one could describe later on as 'Frieda sang.' In the same vein, if Frieda sings the same song twice in a row and without interruption, we can distinguish at a minimum three singing events: the event of Frieda's first singing of the song, the event of her second singing of the song and the event of her singing the two songs one after the other. In these respects, Bach noted already in the early eighties (Bach 1986), the events described by non-telic verbs and verb phrases are much like the bits of matter described by mass nouns and mass noun phrases. (For instance, when you serve yourself two spoonfuls of apple sauce, then there is the first spoonful, the second spoonful and the two spoonfuls put together: the splash of apple sauce on your plate after you have put the first spoonful there and then the second spoonful on top of it. Each of these can be described as 'some (bit of) apple sauce'.)

Because of the non-telic character of *sing* pretty much every bit of singing can qualify as a singing event. And for that reason every time that there is singing by Frieda going on – every time that we could describe by 'Frieda was singing' – there must have been an event of her singing going on at that time, which could have been described by 'Frieda sang'. That is why we perceive the inference from (3.50.a) to (3.50.b) as valid. But no such entailment relation exists between (3.50.a) and (3.50.b).

Does this account of the difference between (3.50.a,b) on the one hand and (3.50.c,d) on the other get us any closer to an articulation of the relation between a state description like that attached to the AspP node in (3.47) and an event description like that attached to the VP node in (3.46)? Not really. But it does point us in a certain direction: the progressives of both telic and non-telic event descriptions describe processes which qualify as parts of events described by the corresponding non-progressives. The difference between the telic and the non-telic descriptions is that any process described by the progressive of a non-telic verb or verb phrase is ipso facto an event that can also be described by the non-progressive phrase. But for telic event descriptions this is not so. A proper initial segment of a shop closing event, for instance, is itself not an event of closing the shop; for that would require that the was reached of the shop being closed, and that state will only reached at the end of shop closing.

What exactly does it mean to be closing the shop?

What does being in the process of closing the shop precisely amount to? What are the things that you must be doing in order to qualify as being involved in a process that can be described with the progressive 'be closing the shop'? It seems quite impossible to give any generally valid specification of these things, even if it is only a partial one. What needs to be done to close a shop at the end of the day depends on all sorts of particularities pertaining to the shop in question: whether some of its wares are on display outside; whether the shop window has roller blinds that have to come down at night; what kind of lighting the shop has that has to be turned off or dimmed; whether the shop has an alarm system and how that has to be activated; and so on. An exhaustive description of all these different possibilities seems quite out of the question. Even spelling out what had to be done to close some particular shop on some particular evening would be a non-trivial task and one cannot help feeling that it wouldn't get us much closer to understanding the relationship between progressive and non-progressive uses of the verb *close*, let alone the semantic relations between progressive and non-progressive forms in general.

When you try to think about the truth conditions of the progressive sentence (3.50.a) it will dawn on you soon enough that what makes it so hard to state what they are is mirrored by the same uncertainties about what is or can be involved in events that are described by the non-progressive (3.50.b). But the truth conditions of (3.50.b) are not affected by these uncertainties. The truth of this sentence depends on whether the event culminates in the result state of the shop being closed; *how* that state is reached matters little. For the progressive sentence the situation is different. This sentence can be true even when the result state – that of the shop being closed – isn't reached. But not just anything will do in such cases. In order that the progressive description be true, what is being described must be like what is going on in completed closing events. The question is: How can we capture this notion of 'being like what goes in completed closing events' without drowning in a quagmire of obscure and obscuring detail?

The best-known accounts from the formal semantics literature have tried to avoid the quagmire by analyzing the meaning of the progressive entirely in terms of the concept of a possible (that is: actual or non-actual) completion of a process ((Dowty 1979), (Landman 1989*b*)). Suppose, these proposals argue, that we have a case where a progressive description is true while the corresponding non-progressive description is not. Then what makes the former true is that if it hadn't been for some interfering factor, which prevented the world from developing in the way it would otherwise have, an event satisfying the non-progressive form *would* have emerged. Or, stated somewhat differently, a process instantiates the progressive description if it would have culminated in the result state associated with events satisfying the non-progressive description, had it not been for some 'interference with the normal course of events'.

The merit of these proposals is that they avoid the seemingly impossible task of describing in concrete terms what a process must be like so that it can be truthfully described by the progressive form when there is no actual culmination. But they run into their own problems when it comes to saying more precisely what 'interference with the normal course of events' precisely comes to. Dowty tried to deal with this question by introducing a new concept into possible wold semantics.[20] This is the notion of an *inertia world*. Inertia

[20]At this point it is necessary to refer to aspects of semantics that we haven't discussed so far. Possible worlds play a central role in formal semantics. But I have been trying to stay away from them in this course, because quite a bit can be done without bringing them into play (i.e. without distinguishing between *different* possible worlds), including all that has been discussed in this course so far; and one simply cannot do everything within the limits imposed by a single semester. In first approximation different possible worlds correspond to different models (for the Predicate Calculus or for the λ-Calculus); but the models of possible world semantics are models which contain whole collections of possible worlds, and thus are more like bundles of the models that we have been working with. There is also a further complication that has to do with our present focus on temporal reference: our earlier models can be seen as describing the world as it is at one particular time. The models that are adequate for the interpretation of the representations of tensed sentences and discourses that we are constructing now must be models that describe worlds as developing through time. Such models can be thought of as temporal successions of models of the earlier kind, each one of which gives the extensions of the predicates of the language *at its point in time*. For instance, the predicate'cat' may now have different extensions at different times of one and the same model: the extension of cat' at time t need not be the same as the extension at some later time t', as some cats may have died between t and t' and others may have been born. So to each of the times of a model of the new kind there will be a corresponding model of the sort considered up to now.
The upshot of all this is that the models needed in a formal semantics of the kind that semanticists like Dowty and Landman assume will contain multiple worlds and each of those worlds will extend from the beginning of time to the end of time. It is of the possible worlds that make up such models that you should think when reading the remarks that

worlds are worlds where, intuitively speaking, the normal course of events is not interfered with. More particularly, we can for any possible world w and any time t on the time line of w consider the *inertial continuations of w after t*: these are those worlds w' that coincide with w up to t, but may diverge from it from that point onwards and that are inertial worlds in Dowty's sense. In the following discussion it is this second term – that of w' being *an inertial continuation of w after t*– that I will be using.

When w is itself among its own inertial continuations after t, then cases of the kind that triggered this discussion – where a progressive form is true but the corresponding non-progressive form false – will not arise in w at t. These cases arise only when w is not among the inertial continuations of w after t. And in those cases, Dowty's account maintains, the non-progressive form may be false (because what was going on in w at t did not develop into an event of the kind described by the non-progressive form) and yet the progressive form may be true because what was going on in w at t would have developed into an event kind described by the non-progressive form in each of the inertial continuations of w after t.

An analysis along these lines is appealing since it finesses the multitude of obscuring details I alluded to earlier. But our original problem threatens to return in a new guise when we ask ourselves what more in detail can be said about the notion of an inertial continuation.

Among the cases where the progressive form may be a true description while the non-progressive description is false there are some where it is intuitively clear what inertial continuation should amount to. These are clear cases of intervention, as in (3.51): The glass was rolling across the table top and applying the laws of physics to the system constituted by the glass and the table (but not including the speaker of (3.51)) one would have predicted the glass to go over the edge of the table and fall to the ground. However, the speaker interrupted this predictable course of events before the edge of the table was reached.

(3.51)The glass was rolling off the table when I caught it.

Let w be the actual world and t the time when the glass was rolling towards the edge of the table. Then the inertial continuations after t in w are those

follow. In particular, two worlds in such a model may be the same up to some time t, but then diverge. (More about models that extend over time can be found in the next subsection, Section 3.6. For more on multiple, temporally extended possible worlds see also Section 3.7.1 on the future tense.)

in which the speaker does not interfere with the physical course of events and in which there is no other interference with this course of events either. In all those inertial continuations the event of the glass rolling off the table will be completed, and so according to the inertia world analysis (3.48) will be true.

The scenario for (3.51) is about the interruption of a physical process with a predictable outcome. Many other scenarios in which progressive and non-progressive descriptions come apart involve the interruptions of intentional actions. We often describe such actions in terms of the goal that the agent means to accomplish by their performance, which makes them prototypical cases of*telic* event descriptions.[21] An example is the description 'cross the road', a phrase that describes those actions whose goals to each the other side of the road from where you are. Perhaps the most famous example from the literature on the Imperfective Paradox is the one in (3.52).

(3.52)The old lady was crossing the road when she was hit by a truck.

The scenario for (3.52) is like that for (3.51) in that here too there is an interruption of a system that would in the normal course of events have led to an instance of the non-progressive description if no interference had occurred. (In this case the 'system' is the one consisting of the old lady and the road, but not including the truck.) Progressives of telic action descriptions are particularly common instances of the Imperfective Paradox. The reason, no doubt, is that finding ourselves frustrated by outside interference when we are trying to realize a goal is an experience we all know only too well, sand which registers with us when we have it.

Progressives of non-telic descriptions

We saw that the problem we have been discussing – a non-progressive sentence is false and yet the corresponding progressive sentence true – arises for some event descriptions but not for all. Non-progressive VPs like *sing*, which describe processes without culminations, do not give rise to this problem. Since the events that instantiate the non-progressive description do not have intrinsic culmination points, proper parts of such events qualify as instantiations of the non-progressive description too. And since the converse entailment, from the non-progressive to the progressive form, holds as well, (3.50.a) and (3.50.b) are truth conditionally equivalent. (The entailment in

[21] *telos* is the Greek sword for 'goal'.

this second direction, from non-progressive to progressive past tense descriptions, holds in general, for telic descriptions like (3.50.c,d) no less than for non-telic ones like (3.50.a,b).)

This equivalence means that for 'non-telic' VPs like *sing* we do not need to worry about the problems related to the Imperfective Paradox. However, non-telic verbs confront us with a different question: If the non-progressive and progressive forms of non-telic VPs are equivalent, why should the progressive form be used at all for such verbs? Why don't we just make do with the morphologically simpler non-progressive form? I am not sure what a fully satisfactory answer to this question could be like. But part of the answer will have to be that the use of English progressive forms of non-telic verbs like *sing* is to some extent a matter of grammaticalization. The choice between progressive and non-progressive forms is determined partly by whether the verb is an *event verb* or a *state verb*. And this distinction is a grammatical one insofar as some verbs that behave like event verbs according to English Grammar describe eventualities that seem more like states than like events. Striking examples are English position verbs, among them *stand*, *sit*, *lie*. From a semantic point of view these verbs seem to describe states – a state of something standing or sitting or lying somewhere. But formally they behave like event verbs; this is what makes their progressive forms possible, and in certain contexts makes them obligatory.

The need for progressive forms of position verbs shows dramatically when the verb is in the present tense. This has to do in part with properties of the English Tense-and-Aspect system and the role of the present tense within it, about which we will have more to say in Section 3.7.2. (Section 3.7.2 is devoted to the present tense.) A proper discussion of the phenomena of which we will discuss some examples here will therefore have to wait until then. But since the examples I will provide below are such striking illustrations of the event-like character of position verbs, I want to show them here.

Consider the sentences in (3.53). (The first two of these were given at the start of the present section as (3.50.a,b).)

(3.53)a. Frieda was singing.

b. Frieda sang.

c. Frieda is singing.

d. Frieda sings.

(3.53.a) and (3.53.b) illustrate what we have just been saying about the progressive and non-progressive of non-telic event descriptions: Since the two forms are equivalent, the one can be used to make a true assertion when the other one can (even if one of them may be preferred for other than merely truth-conditional reasons). But between the present tense sentences (3.53.c) and (3.53.d) there is no such equivalence. When you want to give a straightforward description of what Frieda is doing right now, the correct form is (3.53.c). For instance, when I ask you about the noise that is coming from the attic, and you reply (3.53.a), then that is a straight answer, which describes Frieda's current activity as the source of the noise that has caught my attention. (3.53.b) cannot be understood in this way. What it conveys is something like a disposition of Frieda's or something that she is in the habit of doing: Frieda is a singer, or would like to think of herself as a singer, and she often engages in the sort of thing that singers do. By making this general statement about Frieda you are inviting me to draw the inference that what we are hearing right now is a manifestation of this habit or disposition.

Why present tense non-progressive event descriptions are subject to the kind of constraint that leads to a dispositional interpretation for (3.53.d) is a matter that will be taken up in Section 3.7.2. For now, just note that state descriptions are not subject to this constraint. When you want to say that Mary is at this moment in a state of anger, the straightforward way to to say this is 'Mary is angry'. The progressive counterpart of this, 'Mary is being angry', is highly marked, as if one wanted to convey that Mary's anger was a put on. As a matter of generality progressives of state descriptions are marked if in fact they aren't plainly impossible, and that independently of tense. (More about this below under the heading 'progressives of state descriptions'.)

Armed with these observations let us now go back to the English position verbs, of which we said that they behave like event verbs, although they could easily be thought of as describing states. Consider the present tense sentences in (3.54). These sentences form four pairs, each consisting of a progressive sentence and its non-progressive counterpart. The two sentences of each pair clearly differ from each other in meaning, much as we saw for vent verbs, although they could easily be thought of as describing states. Consider the present tense sentences in (3.49.c) and (3.49.d).

(3.54)a. John is sitting in the chair to the left.

b. John sits in the chair to the left.

c. The statue is standing in the middle of the square.

d. The statue stands in the middle of the square.
e. John is standing in the corner of the living room.
f. John stands in the corner of the living room.
g. The book is lying on the table.
h. The book lies on the table.

(3.54.a) and (3.54.b) differ in a way that is reminiscent of the difference between (3.49.c) and (3.49.d). If we want to convey which seat John is in right now, we say 'John is sitting in the chair to the left' as in (3.54.a). 'John sits in the chair to the left', as in (3.54.b), would not be right, but would rather suggest that the chair on the left is John's regular place. This sentence can be true even if John isn't sitting in the chair right now when it is being said.

The contrast between (3.54.c) and (3.54.d) is different and not quite as dramatic. Both can be used to indicate the statue's location. But note that with entities that are more easily movable, the difference becomes more like that between (3.54.a) and (3.54.b) . Especially when the subject is a person, as in (3.54.e) and (3.54.f), the simple present version seems rather funny, as if John was a piece of furniture and (3.54.f) an instruction to the movers about where to put him. Finally, the second member (3.54.h) of the last pair doesn't seem to have any coherent interpretation at all, not even as a statement of where the book normally is or ought to be.

At this point it should be clear that the contrasts we can observe between the sentences in each of the four pairs have to do with the fact that they are present tense sentences. In fact, when their present tenses are replaced by past tenses, the contrasts disappear, and we are back to equivalences like the one between (3.50.a) and (3.50.b) (except that the simple past form *lay* of the verb *to lie* seems to have become somewhat obsolete, which I take to be the reason why 'The book lay on the table' sounds rather stilted).

To repeat, the point of this discussion of the progressives and non-progressives of position verbs has been to show the role of grammaticalization in the use of the English progressive. When the progressive can or must be used, and what meanings progressives carry when it is all right to use them, varies between verbs. And the crucial divide, between event verbs and state verbs, is to some extent a conventionalized one: some verbs function grammatically as event verbs although their lexical meanings would seem to qualify them as state verbs. For these verbs the progressive functions much like it does for other non-telic event verbs. Its behavior relative to the true state verbs of English is quite different, as we will see below.

Progressives of state descriptions

By and large, progressives can be formed only out of event VPs. Progressives of state verbs are often ill-formed, as shown by the ungrammatical sentences in (3.55).

(3.55)

a. * Mary is knowing the answer to this problem.

b. * John is being six feet tall.

Can progressives of state verbs ever be grammatical? The examples in (3.56) seem to indicate that they can be.

(3.56)

a. Bennie is just being obnoxious.

b. Stella is being her usual innocent self again.

c. Carla is loving her new job.

d. As long as they are believing you are speaking the truth, there isn't too much you have to worry about.

The input descriptions to the progressives in (3.56) are all state descriptions. The first two involve the copular verb *be* – the VPs consisting of copular *be* and a complement (AP, PP or NP) – form an open class of state VPs – and the verbs of the last two examples, *love* and *believe* are standard examples of state verbs. But even so, the sentences in (3.56) do not show that the progressive accepts state descriptions as well as event descriptions. When you look more closely at the contents of these sentences, you see that in these sentences the progressive adds an element of gentility that is absent from their non-progressive counterparts. In (3.56.a) the predicate *is being obnoxious* seems to attribute a certain kind of willful and controllable behavior to the subject; it seems to imply that Bennie is involved in behavior that he could stop if he only made the effort. That meaning is not present in the non-progressive sentence 'Bennie is obnoxious'. in the other sentences the progressive also adds an agent-related dimension. (3.56.b) suggests that Stella is putting on her 'little innocent' act again, (3.56.c) that Carla's attitude to her new job has a dimension of active feeling and perhaps also that it is a condition which holds now but may not last; and in (3.56.d) a similar element of contingency seems to be implied about the false and potentially unstable beliefs that your lies have managed to install in others.

These examples point towards a general assessment: State descriptions are acceptable inputs to progressivization only when they can be (re-)interpreted as descriptions of something like activities and thus as event descriptions. The progressive can be applied only after such a reinterpretation from state description to event description has taken place. Reinterpretation mechanisms of this sort, which adjust representations to the requirements imposed by operators to which they are adjoined as inputs, are known as *coercion* mechanisms. In the present case the trigger of coercion is the operator triggered by +prog. The operator wants event descriptions as inputs. When its input is a state description, then coercion is needed to turn it into an event description first, via the kind of reinterpretation of which (3.56) offers a few examples.

But note well: Coercion of state descriptions into event descriptions isn't always possible, witness the sentences in (3.55). A precise description of +prog-triggered coercion should not only say what the semantic effects of coercion are when it happens, but also when it can happen at all.

Implicit in the claim that state descriptions have to be coerced into event descriptions so that the operator triggered by +prog can be applied to them is the principle that this operator can never be applied to a state description. The lexical entry for +prog that will be given in the next subsection endorses this principle. But no attempt will be made, either in this entry or elsewhere in these notes, to state the semantic effects of +prog-triggered coercion or the limitations to which saucy coercion is subject.

A lexical entry for the progressive

The feature +prog resembles the tense feature past in that both trigger *operators*, which transform input representations into output representations. To specify the lexical semantics of +prog we can therefore make use of the same format that we used in our formulation of the lexical entry for past. (A more elaborate entry would also specify when sand how this restriction can be overcome through coercion.)

(3.57) (lexical entry for the aspect feature '+prog')

+prog (aspect feature)

Sel. Restr: event description

Sem.Repr: $<e_{ref}, ... \mid K> \rightsquigarrow$

$<s_{ref}, ... \mid \text{s}:PROG(^\wedge e.K \cup \begin{array}{|c|}\hline e \\ \hline \\ \hline\end{array}) >$

Strictly speaking this section of the progressive should conclude with a statement of the verification conditions for DRS Conditions of the form 's : PROG($^\wedge e.K \cup \begin{array}{|c|}\hline e \\ \hline \\ \hline\end{array}$)'. But as was explained above, to do this would require much of the machinery of intensional model theory, and developing that is beyond of the tasks we have set ourselves in this course.

Formal statements of the entailment relations between progressive and non-progressive forms

We have seen that verbs and VPs vary with regard to the entailment relations between their progressive and non-progressive forms. Episodic uses of the non-progressive form generally license the corresponding progressive forms. But the converse entailment, from the progressive to the non-progressive form, holds only for non-telic verbs and verb phrases, but not for the telic ones. A semantic account of the progressive must capture this distinction in some way. One way in which this can be done within a formal theory of the kind we are developing is in the form of Meaning Postulates (MPs). Meaning Postulates go back to Montague's work on natural language semantics (and before that to the work of Carnap). They have been an important tool in Montague Grammar ever since Montague introduced them, and also in other approaches, including the logical form approach we are pursuing here. Within the framework we have adopted, MPs are best seen as statements of general properties of and relations between the various sorts of entities that can be described in the Logical Form Formalism that the given theory makes use of. (For us this Logical Form Formalism is the DRS language that we are in the process of developing in this part of the course.) The formal effect of a MP is to constrain the class of admissible models for the LFF to those members of the initially given model class $\mathcal{M}$ in which the MP is satisfied. Thus the effect of the MP given in (3.58) is to restrict $\mathcal{M}$ to those models M in which the interpretation of PROG, as a function from event properties to state properties, has the feature that for every event with the former property there is a multiplicity of states that have the latter.

(3.58)
$$\begin{array}{|c|}\hline e\ \ t \\ \hline \mathrm{K}(e) \\ t \subseteq e \\ \hline\end{array} \Rightarrow \begin{array}{|c|}\hline s \\ \hline s\colon \mathrm{PROG}(^\wedge e.\ \mathrm{K}(e)) \\ \mathrm{dur}(s) = t \\ \hline\end{array}$$

The intuitive content of this Meaning Postulate should be clear: When a non-progressive verb or verb phrase is instantiated by some event e, then for any interval of time t included in the duration of e there is a state described by the progressive form of the verb or VP with duration t.

Another Meaning Postulate, marginally less obvious than (3.58), is to the effect that if a verb or verb phrase is non-telic, then the corresponding progressive verb or verb phrase is instantiated *only* when the non-progressive verb or verb phrase is instantiated as well. For instance, (3.58.1) is the instance of this general principle for the non-telic verb *sing*. (It is this MP which formally accounts for the entailment from (3.50.a) to (3.50.b).)

(3.59)
$$\begin{array}{|c|}\hline s\quad x \\ \hline s\colon \mathrm{PROG}(^\wedge e.\ \begin{array}{|c|}\hline e \\ \hline e\colon \mathrm{sing'}(x) \\ \hline\end{array}\ ) \\ \hline\end{array} \Rightarrow \begin{array}{|c|}\hline e' \\ \hline \mathrm{dur}(e') = \mathrm{dur}(s) \\ e'\colon \mathrm{sing'}(x) \\ \hline\end{array}$$

(3.58) and (3.59) differ in that (3.58) is a general principle that applies to arbitrary event descriptions, whereas (3.59) is about a single verb, *sing*. Since there are many other non-telic event verbs, we will need more MPs like (3.59) in order to cover those other verbs. We can do this by introducing a separate MP for each such verb. But isn't there a more economical way of capturing the content of all those MPs? There would be if we had a general characterization of the non-telic event verbs, or of substantial subclasses of them. An MP could then refer to those characterizations and cover all verbs that satisfy it. Unfortunately, I do not know of any such characterizations, which would obviate the need to specify that a verb is a non-telic event verb in its lexical entry. Once this information is in the lexical entries of these verbs, then it is of course possible to formulate and MP that would state for arbitrary non-telic event verbs what (3.59) states for *sing*. But the gains that would bring us would be largely cosmetic. For applying the MP to any particular verb would require the consultation of its lexical entry.

3.6 Truth Conditions, Models, Ontology

A DRS is true in a model M if there is a verifying embedding of its Universe into the Universe of M. That is the central principle which governs the model-theoretic semantics of DRSs, the 'formulas' of our DRS languages and the logical forms via which a DRT-based semantics assigns truth conditions to the sentences and discourses of its object language. It is a principle that has been there from the start, irrespective of how DRSs are constructed and no matter what the precise form and expressive power is of the DRS languages adopted. This principle is to be upheld irrespective of the form and power of the DRS languages used. It imposes different constraints on the model classes that are appropriate to different DRS-languages: Each DRS language imposes its own requirements on the models in which its DRSs can be evaluated in an intuitively plausible way. This is true in particular for the DRSs we have been building in Part II of these notes. From now on we will assume that these DRSs, as well as all those that will play a part in what follows, belong to a DRS language $L_{DRS,t}$ (where 't' stands for 'time').

In our review of the top down approach to DRS construction we went through a brief account of the model theory for the DRS language we considered there. Let's refer to that language as 'L_{DRS}'. For L_{DRS} it was easy to determine a suitable model class. The natural choice were the models for first order Predicate Logic – more precisely, the models for a language of Predicate Logic whose non-logical constants are the same as those as those of L_{DRS}.

For $L_{DRS,t}$, however, simple first order PC models won't do. There are several reasons for this. The first and main reason is that $L_{DRS,t}$ has drefs standing for special sorts of entities: (i) times and (ii) eventualities, with (ii.a) events and (ii.b) states as its two mutually exclusive subcategories. Since evaluation of DRSs in models is by way of embedding functions that map DRS

Universes into model Universes, this means that the Universes of the new models must contain entities of these sorts. That is, each Universe U_M must contain (i) a set T_M of times and (ii) a set EV_M of eventualities, with a set E_M of events and a set S_M of states as subsets.

But that is not all. Our DRSs also make use of special symbols to express relations between times and events, such as $<$, $\subseteq$, $\supset\subset$. These symbols can be thought of as expressing part of a 'logic of time', with laws such as that if $t < t'$ and $t' < t''$ then $t < t''$, if $t < t'$ then not $t' < t$, if $t \subseteq t'$ and $t' \subseteq t''$ then $t \subseteq t''$, if $t \subseteq t'$ and $t' < t$ then $t = t'$ and so on. Intuitively such principles have the status of 'temporal laws'. And if they have the status if laws, then they ought to be *semantically valid* – that is, they should come out as true in every model M, no matter what times from T_M the drefs t, t', t'' are mapped to. This of course requires that T_M have a certain structure that 'validates' those laws. As for the laws listed above, it is plain that they are all valid on the assumption that T_M has the structure of a *linear order*.

The general assumption we will make about our models is that each M comes with a linearly ordered structure of temporal instants. We represent this structure as $T_M = <\mathrm{T}, \prec>$, where T is a set (of the instants of T) and $\prec$ a linear order of T.

Stronger assumptions about the structure of time would in principle be possible. For instance, one might consider whether time shouldn't be assumed to have the structure of the real line (i.e. whether we should assume it to be order-isomorphic to the real numbers). This after all is the assumption about time that is made in all of classical physics. I do not think, however, that it is either necessary or right to insist that time must satisfy this stronger assumption, and we won't make it. All we want to assume is that each model M contains a time structure $T_M = <\mathrm{T}_M, \prec>_M$ that is some kind of linear order.[22]

[22]In DRT there are special reasons for not wanting to make the stronger assumption in general (see (Kamp 2017)). Here we will only assume that $<\mathrm{T}_M, \prec>_M$ is a linear ordering, and impose no further general constraints. It has on the other hand also been contended that there is no justification for the assumption that the order of time is linear; in particular, time can be branching in the direction of the future. We will take this matter up in the next section, which discusses the future tense. As will be seen there, even if we allow for the possibility that at many times (and perhaps at all times) in the history of a world there is more than one way in which it can continue, this is no compelling reason for abandoning the assumption that time is linear in the direction of the future as well as in that of the past.

As mapping targets for the temporal drefs t, t', etc it is not the instants of T_M that we need, but rather its temporal intervals. To be formally explicit, we will identify the 'intervals of T' with the non-empty *convex subsets* of T, i.e. with all those non-empty subsets T' of T which have the property that whenever t and $t' \in$ T' and $t \prec t'' \prec t'$, then $t'' \in$ T'. We will denote this set as 'INT(T)'. We assume that the Universe of a model contains both the instants of T and the members of the set INT(T).[23]

It is not only the times of each model M that come with a temporal structure. The eventualities of M have temporal structure as well. However, given the temporal structure T_M we can account for the temporal structure of EV_M by assuming that each ev in EV_M has a *duration* in M, an interval $\mathrm{Dur}_M(\mathrm{ev})$ in INT(T) which covers all and only the instants during which ev is going on. That is, we assume that there is a function Dur_M which maps each eventuality ev in EV_M to its *duration* $\mathrm{Dur}_M(f(ev))$. We can then evaluate DRS Conditions like '$ev < ev'$' or '$ev \subseteq t$' in M via Dur_M. For instance, the embedding function f will verify the condition '$ev \subseteq t$' in M iff $\mathrm{Dur}_M(f(ev))$ is temporally included in $f(t)$. (In this last statement 'temporally included' means of course 'temporally included in the sense of T_M'.) In short: the

[23] The standard way of defining the set of 'intervals of T' is not fully satisfactory for our purposes, because of the distinction that it makes between open and closed intervals in those cases where T is a non-discrete ordering (as for instance when T cis like the real or the rational numbers. The distinction between closed and open intervals in the models for $L_{DRS.t}$ appears to be irrelevant to the semantics of tenses, aspect operators and other devices for referring to time that are found in natural languages. (It is maintained in quite a few accounts of tense and aspect that the difference between perfective and imperfective descriptions of eventualities can be characterized semantically in terms of whether the interval occupied by the described eventuality is open or closed on the right (i.e. towards the future): for the eventualities of perfective descriptions the occupied interval is closed on the right, for the eventualities of imperfective descriptions it is open. When this is a claim about intervals of real time that are occupied by the real eventualities that imperfective and perfective clauses describe, it is almost certainly either meaningless or false.) For non-discrete time structures the most elegant way to do justice to the position that natural language semantics is insensitive to the open-closed distinction is to 'divide' the structure by an equivalence relation which identifies, for any t and t' such that $t \prec t'$, the intervals (t,t'), $(t,t']$, $[t,t')$ and $[t,t']$. Let $\approx$ hold between two intervals I and I' iff there are t, t' such that $I \in \{(t,t'), (t,t'], [t,t'), [t,t']\}$ and $I' \in \{(t,t'), (t,t'], [t,t'), [t,t']\}$. We define INT($T$) to be the set of equivalence classes under $\approx$. Using the fact that for every convex subset T' of a lineally ordered set T there are $t, t' \in$ T such that T' $\in \{(t,t'), (t,t'], [t,t'), [t,t']\}$, we define for I, I' in INT(T), $I \prec I'$ iff there are t, t', t'', t''' such that $t \prec t' \prec t'' \prec t'''$, $I \in \{(t,t'), (t,t'], [t,t'), [t,t']\}$ and $I' \in \{(t'',t'''), (t'',t'''], [t'',t'''), [t'',t''']\}$; and $I \supset\subset I'$ iff there are t, t', t'' such that $t \prec t' \prec t''$, $I \in \{(t,t'), (t,t'], [t,t'), [t,t']\}$ and $I' \in \{(t',t''), (t',t''], [t',t''), [t',t'']\}$; finally, I is *temporally included in* I' iff there are t, t', t'', t''' such that $t'' \preceq t$, $t' \preceq t'''$, $I \in \{(t,t'), (t,t'], [t,t'), [t,t']\}$ and $I' \in \{(t'',t'''), [t'',t'''), (t'',t'''], [t'',t''']\}$.

function Dur_M allows us to transfer temporal relations between intervals of T_M to the eventualities of EV_M.

The set of eventualities of a model does not only have a temporal structure (imposed by Dur), but also a *mereological* structure, given by the part-whole relation $\preceq_m$. This relation is not to be confused with temporal inclusion. It entails temporal inclusion, but is much stronger. For instance, e might be an event of John doing the dishes, and e' an event of Mary coming into the kitchen while he is doing that. Then e' is temporally included within e. But it won't be the case that e' stands in the mereological part-whole relation to e. Similar observations about the difference between mereological and merely temporal relations obtain for states, and more generally for states and events (i.e. for eventualities generally). Note also that the mereologcal part-whole relation may obtain between eventualities even when one is an event and the other a state. One of the assumptions we made in our discussion of the progressive was that if e is an event described by a non-progressive eventuality description then there will be states described by the corresponding progressive eventuality descriptions whose durations are included in that of e. We assume that these states do not only stand in a relation of temporal inclusion to e, but that they are also parts of e in the sense of mereology.[24]

The part-whole relation between eventualities can be used as the basis for a notion of *mereological fusion*. There is more than one way in which this notion can be specified. A conservative specification is that according to which the mereological fusion $\oplus$ is defined only for eventualities whose durations abut or overlap. That is, $ev \oplus ev'$ is defined if and only if either there is a $t \in \mathrm{T}_M$ such that $t \in \mathrm{Dur}(ev) \cap \mathrm{Dur}(ev')$ or $ev \supset\subset ev'$. Moreover, when $ev \oplus ev'$ is defined, then $ev \preceq_m ev \supset\subset ev'$, $ev' \preceq_m ev \supset\subset ev'$, for every ev'' such that $ev \preceq_m ev''$ and $ev' \preceq_m ev''$ it is the case that $ev \oplus ev' \preceq_m ev''$ ($ev \oplus ev'$ is the mereologically smallest entity that mereologically contains both ev and ev'), and finally $\mathrm{Dur}(ev \oplus ev') = \mathrm{Dur}(ev) \cup \mathrm{Dur}(ev')$. More liberal definitions are possible as well, but most of them at the price that $ev \oplus ev'$ isn't guaranteed to cover a convex set of times (viz. when ev and ev' are separated by some non-empty interval of time).[25]

[24]The mereology spoken of in this section is in essence the same as that discussed in PART I, Section 2.7. The difference is that there we spoke about the mereological structure of the domain of individuals, whereas in the present section we are primarily concerned with mereological relations between times and eventualities.

[25]There is an unfortunate clash of notations here, between $\prec_t$ as the symbol for temporal precedence and the symbol $e \preceq_m$ for the mereological part-whole relation. The subscripts prevent actual confusion, but are a bit of a nuisance. We will leave them out whenever

Our treatment of the progressive presupposes a 'plenitude' of states, given in (3.60).

(3.60) (Plenitude of States)

Let ev be any eventuality of EV_M and t a time (i.e. an interval) of T_M such that $t \subseteq \mathrm{Dur}_M(ev)$. Then there is a state $s \in \mathrm{S}_M$ such that $\mathrm{Dur}_M(s) = t$ and $s \preceq_{m,M} ev$.

It should be noted that (3.60) combines two principles which from a conceptual point of view do not fully coincide. The first is a principle that applies when the eventuality ev is a state. In that case (3.60) asserts what is often referred to as the 'homogeneity' of states: For any state s and time t included in the duration of s there is a state $s_{s,t}$ that is 'of the same kind as s' but which covers only the temporal part t of s's duration. (In particular, when s is given by a state description SD, then s' will be 'of the same kind as s in that it also satisfies SD.)

When ev is an event e, then (3.60) can be seen as the expression of a different principle, which says that for each t temporally included in e's duration there is a state that is a mereological part of e and whose duration is t. Note that this second principle is closely akin to the Meaning Postulate (3.58) that we adopted as one of our constraints on the operator PROG. In fact, it is natural to take (3.60) and (3.58) as jointly entailing that when a state s satisfies the description obtained by applying PROG to the description of an event e, then s stands in the mereological relation $\preceq_m$ to e. We could make this relation between (3.58) and (3.60) explicit by adopting a further Meaning Postulate which relates PROG in the manner just described to the relation $\preceq_m$. But the constraint imposed by (3.58) and (3.60) will hold wether we do this or not.

This concludes the discussion of the general constraints that we will impose on the times and eventualities of the models for $L_{DRS,t}$. The next question we must address concerns the extensions of the predicate constants of our language. Here we must make a distinction between the constants that translate English verbs and those that translate words of other grammatical categories. For verbs the question is comparatively simple. This is because we treat verbs as descriptors of eventualities. For an example take the verb *sing*. In our approach *sing* is treated as a 2-place predicate, with an event as referential and an individual as non-referential argument. So its extension in

the risk of confusion is negligible.

a model M should be a set of pairs $<e, a>$, where e is a member of E_M and a a member of U_M. Note that each member of this extension (each 'case of singing') will be located somewhere along the time line of M, because Dur_M will map its first component e to some temporal interval $t = Dur_M(e)$. So t is the time when this particular 'case of singing' occurs.

Let us adopt this stipulation for all constants that translate verbs. That is, if V' is a constant that translates an (n+1)-place verb V, with an eventuality argument as referential argument and n non-referential arguments, then the extension of V' in M is a set of (n+1)-tuples $<ev, a_1, ..a_n>$, where ev is a member of EV_M and $a_1, ..a_n$ are members of U_M.

So much for $L_{DRS,t}$ constants that translate verbs; but what about expressions that represent words of other categories? In the DRSs we have been constructing so far the implications of this question haven't surfaced, because the only constants that have played a part in the DRSs constructed so far were proper names ('Frieda') or expressions that we treated as proper names because no better treatment was available yet (like 'the shop'). But consider the following two sentences.

(3.61)

a. Frieda bought a donkey.

b. Frieda hired a student.

There is an important difference between the two common nouns occurring in these sentences, *donkey* and *student*. *donkey* is a *sortal noun*. Sortal nouns are nouns used to classify individuals, as being of one sort rather than another. Such nouns are true of an individual in an absolute sense: something either satisfies them or it doesn't. Or, putting the matter differently, if something satisfies a sortal noun at one time then it does so at all times – once a donkey, always a donkey. Model-theoretically the constants translating sortal nouns can therefore be treated in the way we have been doing until now: their semantics can be captured by assigning them in each model an extension, consisting of all and only those things that satisfy them. So this is what we will assume for the constants of sortal nouns like *donkey*: In each model M for $L_{DRS,t}$ a constant of $L_{DRS,t}$ that translates a sortal noun will be assigned an extension of the familiar kind – a set consisting of all and only the entities that belong to the sort denoted by that noun.

The noun *student* is different. Being a student is not a property that you have once and for all. Some people become students at some time during their lives, then are students for some time and eventually stop being students (and some may then later return to being students again). So for each person a the predicate *student* has a *temporal profile* – a function that tells us for each time t whether a is a student at t. Note that this second case, represented by *student* as opposed to *donkey*, is if anything the more common one. It applies to most nouns, and moreover it applies to more or less all adjectives and all prepositions. (You can be angry at time t and not angry at time t', you can be in the garden at t without being in the garden at t' and so on.)

In spite of the fact that the temporal variability of non-verbal predicates is an extremely widespread phenomenon. (It is plain that it plays a part in the semantics of most sentences as soon as you look at them a little more closely.) But for a long time formal semantics ignored it (in much the way that we have been ignoring it so far in this course). The principal reason for this neglect is no doubt that until not so long ago time and tense weren't much of a concern in formal semantics generally. But even after it became widely accepted that verbs describe eventualities, the temporal variability of non-verbal predicates was mostly banned to the side lines. However, once proper attention has been drawn to the phenomenon, there is no way back to the innocent days of time-invariant extensions. In particular, now that we have brought it up, we will have to deal with temporal variability at some point. We will do so in Section 4.

For the moment, however, I want to set this problem aside and assume that in any model M for $L_{DRS,t}$ each nonverbal predicate has an extension of the same kind as it has in models for the Predicate Calculus. Thus *student* will have for its extension in M as subset of U_M, the extension of *in* in M will be a set of ordered pairs of members of U_M and so on. For these and most other non-verbal predicates this is of course wrong and just a stop-gap. We will return to the issue below in the subsection 'Temporal variability'.

Before we leave the subject of non-verbal predicates, however, there is one further observation to make. Even the extension of a sortal predicate like *donkey* has a temporal dimension in that its members have a limited time span. Donkeys come and go. Each donkey is born at some time and then at some later time it dies. While it is alive it belongs to the extension of *donkey* uninterruptedly, But that is nevertheless only a small part of the time line as a whole. For certain semantic purposes (e.g. in connection with the semantics of the present perfect) the question when an individual exists is important. So this is an aspect of the world that our models should be able to capture. One way to secure this is to adopt a device that was originally developed within quantified modal and temporal logic. We introduce a special predicate Ex which relates individuals to times and which holds between an individual a and a time t iff t is a time at which a exists. So the extension of Ex in a model M is a set of ordered pairs $<a, t>$ where $a \in \mathrm{U}_M$ and $t \in \mathrm{T}_M$. With Ex in place we can think of the extension in M of a sortal noun like *donkey* as in some sense varying with time too. In a derivative sense the extension of *donkey* in M at t is the set of individuals a that (i) belong to the time-invariant extension of *donkey* in M and (ii) are such that $<a, t>$ belongs to the extension of Ex.

There is one non-verbal category for which it can be argued that temporal variability doesn't apply to it. This is the category of proper names. It is of course true that an individual receives a name only at some point during its life time and that that point may be long after the individual came into existence. And of course the name could not have been used to refer to the individual at times before that point. Thus there is a change of some sort that takes place when the name is given to the individual. But the change need not be seen as a change in the individual. It is better thought of as a change in the language, which now has a new individual constant, consisting of the name *as* name of the individual to which it has just been given. As soon as this labeling event (or 'baptism', as philosophers have been calling such events since Kripke's *Naming and Necessity*) has occurred, the name can be used as name for the individual as if it had always been part of the language, and aa a name for the individual *simpliciter* (and not some later time slice of it that started only when the name seas introduced).

We are now in a position to give a first definition of the models for $L_{DRS,t}$. The definition doesn't deal with the temporal variability of non-verbal predicates and for this reason is preliminary.

(3.62) (Definition of models for the new DRS-language (to be revised in Section 4))

A *model for* $L_{DRS,t}$ is a structure $M = <T, EV, \preceq_m, U, F>$, where:

(i) T is a time structure $<\mathrm{T}, \prec_t>$, with $\prec_t$ a linear order of T;

(ii) EV is an event structure <EV,Dur,Event,State>, with Dur a function from EV into the set of intervals of T, and Event and State subsets of EV that are mutually exclusive and jointly exhaustive.

(iii) U is a set which includes EV, T and INT(T) (where INT(T) is defined from T along the lines indicated above).

(iv) $\preceq_m$ is a weak partial ordering (i.e. $\preceq_m$ is reflexive and transitive) with the additional property that if $ev \preceq_m ev'$ and $ev' \preceq_m ev$, then $ev = ev'$, and with a partial sum $\oplus$ as described above.

(v) F is a function that assigns each predicate constant of $L_{DRS,t}$ an appropriate extension. More specifically:

(a) for the temporal predicate '$\prec_t$' of $L_{DRS,t}$ F($\prec_t$) = $\prec_{t,M}$ (where $(\prec)_{t,M}$ is the second component of T_M). Note that this stipulation also fixes the interpretation of $\subseteq$ and $\supset\subset$, which are definable in terms of $\prec_t$). Also F(dur) = Dur.

(b) for each constant V' that translates an event/state verb V with n non-referential arguments (where n $\geq$ 0) F(V') is a set of n+1-tuples $<e/s, a_1, .., a_n>$, where e/s is an event $\in$ E_M/state $\in$ S_M and $a_1, .., a_n$ are members of U_M.

(c) For each DRS constant P' that translates a non-verbal n-place predicate P (n $\geq$ 1) F(P') is a set of n-tuples $<a_1, .., a_n>$, where $a_1, .., a_n$ are members of U_M.

(d) For each constant c' of $L_{DRS,t}$ that functions as a proper name F(c') is a member of U_M.

One important new feature of our DRSs is the presence of the indexical dref n. Occurrences of n in a DRS K represent the 'utterance time of the sentence or discourse represented by K'. But what exactly does that tell us about the verification of K in some model M? There are two ways we can try to make sense of this question. One is to assume that the model M contains information about which utterances are made at which times. Suppose that t is one of the times at which an utterance is made of a discourse D for which K is the semantic representation. Then the question whether this utterance of D is true in M is settled by the existence of a verifying embedding f of K in M such that $f(n) = t$.

This is one way to deal with the semantics of n, but as has been convincingly argued by Kaplan (one of the fathers of our current understanding of the semantics of indexicality, see (Kaplan 1989)), it is too restrictive. The alternative account that Kaplan favors considers it immaterial for the truth evaluation of a sentence or discourse D in a model M at a time t whether D

was actually uttered in M at t. Even in case there wasn't any such utterance at t, it would still be possible to ask whether D, as represented by K, *would* have been true in M at t *had it been uttered* in (the world represented by) M at t. In other words, on this alternative view DRSs are to be thought of as semantic representations of *possible* utterances; and the question whether a DRS K is true in M at t makes sense in principle for any time t of M as an assessment of whether an utterance at t of the discourse represented by K *was or would have been* true in M. We follow Kaplan in adopting the second answer to the question. In other words, we will consider K to be true in M at t iff there is an embedding function f such that $f(n) = t$ and f verifies K in M according to the familiar definitions for DRS-verification (as we discusses as part of our review of the top-down approach to DRT). Note that for the purpose of applying the standard approach to DRS verification we can think of the models for $L_{DRS,t}$ as models for a language of first order Predicate Logic with predicates 'Time', 'Eventuality', 'Event', 'State', '$\prec_t$', '$\preceq_m$', 'due', etc. and Meaning Postulates restricting the extensions of those predicates. (For more on Meaning Postulates see the section on Ontology below.)

These considerations about the role of n relate to another new requirement for our models; in fact they presuppose it: Our DRSs can contain, all at once, information about what is going on at n, about eventualities preceding n and about eventualities that come after n.[26] In other words, a single DRS of this kind can talk about different parts of the history of the world, as seen from a given temporal perspective point, and not only about what the world is like at one single point in time, while ignoring all that was the case before that point and that will be the case after it.

Temporal Variability

We noted that the extensions of non-verbal predicates often vary as a function of time. Many of the things we say depend on this for their truth. An example is the following sentence.

(3.63) Today's toddlers will all be adults twenty years from now.

This sentence seems undeniably true. But its truth depends on which extensions are taken to be relevant for the interpretation of the occurrences

[26]We haven't talked about the present or future tense, which are used in a language like English to talk about the present and future. Discussion of these tenses will follow in the next two sections.

it contains of the nouns *toddler* and *adult*. The qualification *today's* makes it clear that the relevant extension of *toddler* is its extension at the time of utterance, whereas the relevant extension of *adult* is the one twenty years after the utterance time, the time indicated by the temporal adverb *twenty years from now*. Crucial is of course that the times of the two extensions are distinct, since one cannot be a toddler and an adult at the same time.

For other sentences there are presumably different mechanisms that determine the times of the relevant extensions of the nouns and other non-verbal predicates those sentences contain, or that impose certain restrictions on what these times can be. An example is given by the noun *wife* and the adjective *married* in the sentence (3.64).

(3.64) When my wife and I first met she was married to Nicholas Parker.

The intended extension of *married* in this sentence was clearly that at the time of the mentioned meeting. The time of the relevant extension of *wife* must have been a different one. A default guess would be that it must be the utterance time. but depending on the contextual background of a given utterance of (3.64) it could also be some other time, somewhere between the utterance time and the time of that first meeting – for instance when the woman that the speaker is referring to no longer is his wife, because she died, or because they got divorced, or both. How we determine the times of the intended extensions of non-verbal predicates as part of interpreting the sentences in which they occur is a non-trivial problem, and one to which there exists to my knowledge no fully satisfactory solution.

For non-verbal lexical categories other than nouns temporal variability is if anything even more prominent. For instance, many adjectives are used to attribute properties that are contingent and that may be true of the things they are true of for limited, and often only short periods of time. Examples are *tired*, *angry*, *hot*. The same is true of many uses of prepositions. We often have some rough idea of how long it may take before a prepositional relation ceases to hold of two entities between which it holds now, but that is only a measure of the complexity of the problem. (Example: On the roof opposite there are are (a) a bird, (b) a cat and (c) a brick. The biddy may be gone in the next few seconds; the cat may take longer to disappear (it is dosing in the sun and seems quite content); and the brick is likely to stay where it is until I decide to remove it.)

The capacity that individuals have for going in and out of the extensions of non-verbal predicates is illustrated dramatically by sentences with *change of* *state verbs*, like *become*, *turn* or *make*. Other telling examples are those involving directional forms of spatial prepositions like *into* and *onto*. Temporal presupposition triggers like *no longer*, *again* or *back* also are a copious source of sentences that assert changes in the extensions of one or more non-verbal predicates. Some illustrations can be found in (3.65).

(3.65)

a. When she said that, he became angry.

b. In the autumn the leaves turn green.

c. She made him furious.

d. She went into the house.

e. She is no longer with him.

f. She is back in London.

g. He is quiet again.

There are several ways in which temporal variation of the extensions of non-verbal predicates can be formalized. A way that is congenial to the DRS language we have adopted is to assume that non-verbal predicates are state describers, just as we have been assuming for state verbs, and that they give rise to DRS Conditions of the same general form. For instance, we can no longer represent the information that j is a student as 'student'(j)'; rather, the Condition expressing this information should now take the form 's: student'(j)'. Correspondingly we encode the extensions of 'student" in a model M by taking F_M(student') to be a set of pairs $<s, a>$ where $s \in \mathrm{S}_M$ and $a \in \mathrm{U}_M$. Being members of the Domain of the function Dur, the 'non-verbal states' s that occur as first members of these pairs are assigned a temporal location. So the member a of U_M belongs to the extension of student' in M at the time t of T_M iff there is a state s in S_M such that $<s, a> \in \mathrm{F}_M$(student') and $t \subseteq s$.

In Section 4 we will switch to the representation of non-verbal predications in this new form (writing 's: $\mathrm{P}(a_1, ..a_n)$' when P is an n-place non-verbal predicate instead of '$\mathrm{P}(a_1, ..a_n)$'. But for the remainder of Section 3 we stick to the familiar notation for representing non-verbal predications, ignoring temporal variation in such predicates. For this older notation, which ignores temporal variability, extensions for non-verbal predicates can be of the simple form stipulated in Definition (3.62). (See Section 4 for a revision of this definition.)

Ontology

The word 'Ontology' appeared in the title of this subsection but then it never showed up again. What does ontology have to do with what this section has been about? In philosophy 'Ontology' is used to denote that part of Metaphysics that is concerned with 'what there is', or, in the more exotic terminology of the philosophers, with the 'ultimate constituents' of the world. In more recent times the project of making sense of this question and trying to answer what one succeeds in making of it has been extended to include also the characteristic properties of the sorts of entities that are claimed to be among those constituents and the characteristic relations between them. It is in this wider sense that the term is understood within Artificial Intelligence (see for instance: (Baader, Calvanese, McGuinness, Nardi & Patel-Schneider 2003), (Staab & R.Studer 2003)). It is in this wider sense also that we understand the term 'ontology' here.

For us there are two levels at which ontological assumptions can manifest themselves. One level is that of the model theory for the given Logical Form Formalism (here L_DRS, t): One of the things a model theory must specify is the class of models. Part of that specification is a specification of the various sorts of entities that make up the models' Universes and the structural relations between those entities. (Recall Definition (3.62) which specifies both that times and eventualities are among the sorts that make up the Universes of the models and that these sorts are structured in certain ways.) The other level is that of the Logical Form Formalism itself. Its notation and vocabulary carry certain commitments about what sorts of things there must be. Here too our LFF L_DRS, t) offers a good illustration: its 'special purpose drefs $t, t'.., ev, ev', .., e, e', .., s, s', ..$ (or – what comes to the same thing – the sortal predicates 'Time', 'Eventuality', 'Event' and 'State') carry a commitment to the existence of times, events and states. In view of how truth-in-a-model is defined for DRT-based languages – a DRS K is true in a model M if there is a verifying embedding f from the Universe of K into the Universe of M – it is plain that ontological commitments at the level of the LFF entail like commitments at the level of the model theory: If such a truth definition is to deliver intuitively plausible results, then a 'legitimate' embedding function which testifies to the truth of a DRS K in a model M should map special drefs in the Universe of K – drefs that purport to stand for entities of some particular sort – to entities of those very sorts in the Universe of M. But that presupposes that the Universes of the models do contain entities of those sorts. Models for which this is not so should *ipso facto* disqualify as models for the given LFF.

On the other hand there is in general no entailment of ontological commitments in the opposite direction: Certain commitments made at the level of the model theory – for instance that all models must have entities of a certain sort or that all entities of a certain sort are really complexes of entities of one or more other sorts (and thus are 'reducible' to entities of those other sorts, as it is often put in Metaphysics) – need not be reflected in the LFF. (One example: an LFF might talk about physical objects without making any commitments about their underlying structure, but jot may nevertheless be the case that the model theory for this LFF only admits models in which every physical object is a complex composed of atoms.)

The distinction between ontological commitments at the level of the LFF and ontological commitments at the level of its model theory can be extended to the structural relations. So far we have spoken of such commitments only at the level of the model theory – for instance, when we discussed the question whether the temporal ordering relations in models should be linear orders, or whether they should be linear orders of a particular kind (such as that of the reals). But there is also a way of tying such commitments directly to the LFF itself, which makes use of Meaning Postulates.

We have already been using MPs for the purpose of imposing constraints on the class of admissible models. In view of this purpose it is natural to think of them as addenda to a central definition of the class of models (like Definition (3.62) of the model class for $L_{DRS,t}$) and thus as part of the model theory. But there is also a slightly different way of looking at MPs. We can also see them as *axioms*, and therewith as part of the syntax of the LFF. Adding axioms to the syntactic specification of a formal language amounts to turning the system into a *theory*, in the logical sense of the word 'theory': A system that consists of (a) a definition of the set of its syntactically well-formed expressions (and in particular of its well-formed formulas) and (b) a set of axioms, special formulas that have the status of fundamental assumptions that count as necessarily and invariably true. (Typically, syntactically defined formal theories also specify certain formal rules for deducing formulas from other formulas, so that additional formulas can be deduced from the axioms that are explicitly given, but for the present argument this is not essential.) Familiar examples of theories specified in this form are first order and second order theories of arithmetic and of other parts of mathematics, but also formalizations of parts of empirical science such as Newtonian Mechanics, Thermodynamics, Chemical Bond and so forth.

When we look at MPs in this second way, as part of the syntactic specification of a DRT-based language, then the structural constraints they express appear as commitments at the level of the language itself. These commitments translate directly into commitments at the level of model theory, since, trivially, each model that satisfies the MPs will exhibit the constraints they express. For instance, if we add to $L_{DRS,t}$ Meaning Postulates expressing that $\prec_t$ is a linear ordering, then in any model M that satisfies those MPs the relation $\prec_{t,M}$ will be a linear ordering. But here too there is no entailment of commitments in the opposite direction: The model theory may impose restrictions other than those expressed by the MPs that we treat as axioms. For instance, the model theory for $L_{DRS,t}$ could be formulated in such a way that for every model M that it admits $\prec_{t,M}$ is isomorphic to the reals, even though this is not required by the MPs we have added as axioms to $L_{DRS,t}$.

A final observation of this section is one that only concerns ontological commitments at the level of model theory. At this level ontological commitments can take two forms. The first of these is the one we have so far discussed: The Universes of the models that the model theory postulates must contain entities of certain sorts, usually standing to each other in certain structural relations. To illustrate the other way we consider once more the case of commitments to time. In the next section, which will be devoted to the future tense, we will touch briefly on the subject of *Tense Logic* (or *Temporal Logic* as it is now more commonly referred to). Tense logics are formal systems in which there is no explicit reference to time, but which have operators – so-called 'tense operators' – that refer to time 'implicitly' in that they articulate the semantic contributions these operators make to the formulas in which they occur in terms of what things are the case at which times. The models for Tense Logic make it possible to do this because they represent courses of events through time. In this respect they are like the models for the model theory we have described for $L_{DRS,t}$. But the difference between the two kinds of models is that in the models for Tense Logic 'time remains on the outside'. The models for Tense Logic are functions from times to sets of eventualities, but they do not have universes that contain instants or periods of time among their members. So the model theory for Tense Logic doesn't literally acknowledge times as entities, you might say, and in this regard it differs from the model theory for $L_{DRS,t}$. The model theory for $L_{DRS,t}$ is committed to time as part of what there is. The temporal commitment of the model theory for Tense Logic is only to a kind of temporal fabric to the structure of what there is.

These few remarks were not meant as a proper introduction to Ontology (which they obviously are not). Their central point is that in a semantic framework of the kind we are developing ontological commitments can arise at different levels. One reason why this difference is important is that the languages we speak, and which in many ways reflect the ways we think, may come with ontological commitments that are rooted in human cognition. The tracing of these commitments, for which Emmon Bach cast the term 'Natural Language Metaphysics' some four decades ago, must be sharply distinguished from what through the centuries has been the target of Metaphysics as pursued by philosophers: the enterprise of trying to break through the wall of cognitive prejudice that screens 'true' reality from those who, for better or worse, are equipped with that cognition.[27]

3.7 Other Tenses

3.7.1 The Simple Future

(3.66) gives some examples of sentences in the Simple Future tense.

(3.66)

a. Louise will go to Paris on Sunday.

b. Louise will love Paris.

c. Louise will be visiting Paris.

There is a very plain account of the semantics of the Future tense as it manifests itself in these sentences: the Simple Future is just the mirror image of the Simple Past. In the syntax-semantics interface we are developing this assessment of the Simple Future tense can be captured by adopting a lexical entry that is the mirror image of our entry for the Simple Past in (3.22). More precisely, this entry for the Simple Future is identical with (3.22) except that the condition '$t \prec n$' is replaced by '$n \prec t$'. (See (3.69) below.)

This is the semantics for the future tense that we will adopt and as far as that is concerned we could end this section right here. But in relation to the

[27] In Western philosophy this quest has been a dominating concern throughout its history, from Thales to the likes of Russell, Whitehead, Quine, Davidson and Lewis. In non-Western philosophies questions of what there is, as distinct from what there seems to be, have been no less important, though the ways in which those other philosophical cultures have tried to think about such matters may have been very different.

Future Tense the semantics that such a lexical entry provides is much more controversial than the entry in (3.22) has generally been thought to be as a way of capturing the semantics for the Simple Past. It is important to have some appreciation of why the simple account of the future tense described in the last paragraph has been thought problematic, and what alternatives have been proposed. In the next few paragraphs we will untangle a few of the strands of this controversy.

Doubts that the simple account of the future tense can be right are in one part of a conceptual nature and in another part they are linguistic; but there are evident connections between the conceptual and the linguistic worries. The conceptual doubts go back to antiquity, with Arsitotle's discussion of his 'sea battle example' as central focus. In Book 9 of *De Interpretatione* Aristotle discusses the statement (3.67).

(3.67)There will be a sea battle tomorrow.

This statement is true, his discussion implies, if it is already determined now, at the time when the statement is made, that a sea battle will take place tomorrow. And the statement will be false when it is already determined that there won't be a sea battle. But if the matter is still open – if the question whether there will be a sea battle depends on decisions that are still to be made or on eventualities that will affect the outcomes of those decisions – then the statement is neither true nor false.

Whether or not one accepts Aristotle's intuitions and arguments, they cannot be dismissed out of hand. And it seems clear that in this regard the future tense statement in (3.67) is different from its past tense 'mirror image' in (3.68).

(3.68)There was a sea battle yesterday.

The parallel argument – that (3.68) would be true if it was determined now that there was a sea battle yesterday, that (3.68) would be false if it was determined that there was no sea battle, but that the statement would be without a truth value if there would be no determination either way – seems absurd: how could this matter fail to be determined at the present moment of time, well 'after the fact'? Either there was a sea battle or there wasn't. *Tertium non datur.*

This intuition of a fundamental asymmetry between the future and the past – that the future is 'open', whereas the past is 'closed' – is a deeply anchored portent of our perception of the world and our role in it, as players that help shape the future, and who thus share responsibility for what it is going to be like. There is nothing that can be done about the past. Its only active role now is to set the limits within which shaping of the future is still possible. The notion of an open future, as forever emerging in one of a number of different possible ways from an already closed past, is inseparable from the conception we have of ourselves as free agents.

But what does that tell us about the semantics of the future tense? In the absence of further facts or assumptions: very little. Even if the future is open – if at this point in time there are many things of which it is not yet determined whether or not tomorrow (or next week, or next year) will bring them – it remains true that of all the futures that are still open at any given point there is exactly one that will have emerged after any given amount of time further down the line. And what could be wrong with saying that when we make a statement in the future tense, it is about that future that we are talking – about the one that will unfold in actual fact, so that what we are saying is true if and only if it is true of that one going-to-be-actual future? Of many such statements we may not know now, at the time when we are making them, with certainty that they are true; and our failure to know if they are really true need not reflect any shortcomings of our own, but simply be a consequence of the fact that there simply isn't anything to be known yet about the future matter that they describe. But even if that is so, that doesn't show that such statements are lacking in truth value at the time when they are made. It is a perfectly coherent position to claim that they too are either true or false, no less than statements about the past or statements about a future that the past and present have already fixed. The predicament with future tense statements is just that often we will simply have to wait until the world has developed far enough to fix their truthvalues.

The claim that Future Tense statements are about the actual future, as that future will emerge with time, seems to capture adequately how people understand future tense statements when they make bets. I assert that the democrats will win the next presidential election. (That is the exact form of words I use.) You, who think it is much more likely that the Democrats will lose, challenge me to a bet, giving me pretty good odds – one dollar of mine for three of yours. I accept and so we have a bet: I bet that what I said is true, you bet that what I said is false. If things turn out as I say they will – the democrats win – then that shows that my statement was true and I have won. If things turn out otherwise, then my statement was false and I have lost.

On this view, history draws, as it progresses, a unique timeline through the thicket of possibilities, and we can refer to those parts of the line that are still to be drawn just as we can refer to the parts that have been drawn already. But, to repeat, that still doesn't settle the question whether it is this emergent future that is referred to by any one of the various future tense forms that we find in human languages. Take for instance the form that concerns us here directly, the English future tense in which the auxiliary *will* is followed by a bare infinitive form of the verb. One possibility is that this form is invariably used to refer to the one future that will emerge as time passes (the assumption which as I suggested above best fits our use of the future tense in the context of betting). But it is also possible that this form is used to talk, sometimes or even always, about what will be the case no matter how the world will or would continue from what it is now: that it is true not simply when it proves to be true in the future that will actually emerge, but only when it is true in *all* the futures that *could* emerge from the given present. On such a view the verb form consisting of *will* + infinitive acts as a kind of necessity operator, referred to in the philosophical literature as *historical necessity*.

The logic of historical necessity is an interesting and important topic in its own right.[28] But such a logical investigation has to be distinguished from the *linguistic* question what the semantics is of some particular bit of natural language morphology such as the *will* + infinitive combination in English. A logic and semantics of possible futures with the actual future as one among them provides us with the formal environment in which the implications can be assessed of particular proposals for the semantics of such bits of morphology. But by itself such an environment cannot settle these linguistic questions. Settling those is possible only bycarful investigation of of how native speakers use and understand the given words or constructions.

Compelling arguments for either of the two possibilities mentioned in the one but last paragraph – (i) English *will* refers to the actual future that will

[28]The classical setting in which this logic has been studied is that in which a propositional operator of *historical necessity* is added to a Priorean system of classical propositional or predicate tense logic (For these classical systems of tense logic see (Prior 1967), (Cocchiarella 1965); and for the systems resulting from the addition of the historical necessity operator see in particular (Thomason 2002), (Zanardo 1985)). For a brief précis of the formal semantics of this approach see the end of this section.

emerge; (ii) *will* quantifies over all possible futures – aren't easy to come by. In fact, these two are not the only possibilities. There is the Aristotelean position we already discussed, which is similar to option (ii), but differs in that when the infinitival complement of *will* is true in some possible futures but false in others, then the statement is, at the time of its utterance, neither true nor false. It has also been surmised that English *will* is ambiguous between the first option and either one of the last two, and various suggestions have been made as to how individual uses of *will* can be disambiguated.

Since in my own perception it is clear that in betting contexts the semantics of future tense statements is the one given by option (i) and since arguments that these or other contexts its semantics is different strike me as inconclusive, I propose that (i) is the correct option generally, and that we adopt as lexical entry for the English Simple Future tense (expressed with the help of the future tense auxiliary *will*) the 'mirror image' of our entry for the Simple Past. The new entry is given in (3.69).

(3.69) (lexical entry for the tense feature 'fut')

fut (tense feature)

Sel. Restr: eventuality description

Sem.Repr: $<ev_{ref}, ... \mid K> \rightsquigarrow$

$$<t, ev_{ref}, ... \mid K \cup \begin{array}{|c|}\hline \\ \hline n \prec t \\ \boxed{\text{Event}(ev_{ref}) \quad ev_{ref} \subseteq t} \\ \overset{!}{\vee} \\ \boxed{\text{State}(ev_{ref}) \quad t \subseteq ev_{ref}} \\ \hline \end{array} >$$

English is like many other languages in having various forms for expressing futurity. Apart from *will* + infinitive the most common form is *be going to* + infinitive. It is not all that easy to make out a principled difference between this form and the *will* + infinitive construction. But there appears to be a tendency to prefer *be going to* + infinitive in cases where the described eventuality is in the more proximate future, and perhaps also when its occurrence is presented as the natural outflow or consequence of what is the case now, or of what has just been decided. (To the extent that this is true, it suggests that *be going to* is more like option (ii), or like Aristotle's conception of the future tense.) But the differences are subtle, and judgments appear unreliable. The data are complicated by dialectal variation and diachronic drift.[29]

<u>Exercise</u> Construct DRSs for the sentences in (3.66). (You may treat *go to* as if it were a transitive verb, here with *Louise* as subject and *Paris* as direct object.)

Possible continuations of a single present

We conclude this section on the future tense with a sketch of the formalization of the notion of 'possible alternative continuations of a world w from a given point t in time onwards'. In speaking formally and explicitly about ensembles of alternative possible worlds this subsection reaches beyond a boundary that for the most part these Notes do not cross. We have seen that certain natural language constructs are intensional in that they depend on the intensions of one or more of their parts. Our one example so far was the progressive operator PROG, discussed in Section 3.5.2, which turns an event descriptions into a state description, and where the existence of a state satisfying that description at a given time t depends not just on the actual continuation of the world after t but may also depend on what will be the case in its non-actual inertial continuations. But in Section 3.5.2 we recoiled from the effort to make the semantics of PROG formally precise. Moreover, world plurality was excluded from the models defined in Section 3.6, which each involve just one single world (passing through a multitude of times). Although we will stick with our one-world models at least for the time being, what follows now is an exception to this. The models discussed

[29]In this regard the comparison of the uses of these two forms in contemporary English may be an instance of a general tendency: When a number of linguistic forms are available for a small range of closely related semantic functions, they often end up dividing that space between them. But how the space is divided up will often shift in the course of time and may also vary between different sub-groups of the language community at any given moment.

below involve not only a multiplicity of times but also a multiplicity of worlds.

The focus of this interlude differs from most of the work on modal and intensional logic and semantics in that we will be concerned not just with pluralities of worlds, but with the interactions between worlds and times. The logical languages that have been used traditionally to study these interactions are combinations of classical systems of *modal logic* and *tense logic*. The most prominent of these combined systems result from merging a modal logic and a tense logic both of which are extensions of the same underlying system of classical logic, which can be either classical propositional logic or classical predicate logic. In this brief introduction we only consider the propositional case.

The system of modal logic we consider here is obtained by adding to classical propositional logic a single 1-place sentence operator $\Box$, read as 'it is necessary that'. The standard semantics for this system involves models each of which is based on some set W of possible worlds. Formulas of the system are evaluated for truth and falsity at worlds in models. That is, the basic semantic notion is that of a formula Φ of the system being true in a model M at a world w, where w belongs t the given world set W_M of M. We write this as '$[[\Phi]]_{M,w} = 1$'. Each model M specifies for each sentence letter q of the system a truth value $[[q]]_{M,w}$ (1 for 'true' or 0 for 'false'). The truth values of formulas whose main operator is a truth-functional connective are computed 'locally' from the values of the immediate subformulas from which they are formed – 'locally' in the sense that the value of the formula at w depends only on the values of its subformulas at that same world w. For instance, $[[\neg\Phi]]_{M,w} = 1$ iff $[[\Phi]]_{M,w} = 0$ and $[[\Phi \mathbin{\&} \Psi]]_{M,w} = 1$ iff $[[\Phi]]_{M,w} = 1$ and $[[\Psi]]_{M,w} = 1$. But the truth value of a formula of the form $\Box\Phi$ is not locally determined. Whether $\Box\Phi$ is true in M at w depends not, or not just, on the truth value of Φ in M at w, but on the truth values of Φ in M at other worlds than w. Which other worlds is a choice point in the model theory for systems of modal logic like this one. The simplest and oldest proposal, going back to Leibniz, is that

$$(\Box.1) \quad [[\Box\Phi]]_{M,w} = 1 \text{ iff } [[\Phi]]_{M,w'} = 1 \text{ for all } w' \in W_M.$$

But other truth clauses for $\Box\Phi$ are possible as well. In particular, since the work of Kripke and others from the fifties and sixties it has become standard to assume that the world sets W come with additional structure, and this structure is then used in defining the conditions under which $[[\,\Box\Phi]]_{M,w} = 1$. Specifically, one assumes that W comes with an 'accessibility relation'

R, which specifies for each world $w \in W$ which worlds $w' \in W$ count as 'possible alternatives' to w, or as *accessible from* w, as the terminology has it. In models whose world sets come with such a relation R the truth value of $\Box\Phi$ at w can be specified as determined by the truth values of Φ at the worlds w' that are accessible from w:

($\Box$.2) $\quad [[\Box\Phi]]_{M,w} = 1$ iff $[[\Phi]]_{M,w'} = 1$ for all $w' \in W_M$ such that wRw'.

When the truth values of formulas $\Box\Phi$ are given by ($\Box$.1), then that fixes the logic of the system: The model theory determined by this truth clause for $\Box\Phi$ together with the local clauses for the truth-functional operators fully determines which formulas come out as logically true (i.e. true in all models at all worlds) and likewise which formulas are logical consequences of which others. (Φ is a logical consequence of a set Γ of formulas if the following holds: at each world w of each model M Φ is true in M at w whenever all the formulas from Γ are true in M at w.) When the truth value clause for $\Box\Phi$ is taken to be ($\Box$.1), then the logic of the system is fully determined. (It is the modal logic known as 'S5'.) But when world sets are assumed to come with accessibility relations and the truth clause for $\Box\Phi$ is assumed to be given by ($\Box$.2), then the logic depends on the properties of the accessibility relations. For instance, the formula $\Box\Phi \rightarrow \Phi$ will be a logical truth iff it is assumed that all accessibility relations are reflexive.[30]

The system of Tense Logic that is the other component to the combined system to which this interlude is devoted is set up in a manner closely similar to the Modal Logic just discussed and in fact it could be seen (and was originally conceived) as a modal logic in temporal clothing. The system, originally due to Prior (Prior 1967), is obtained by adding to classical propositional logic the operators P, for 'it was the case at some past time that', and F, for 'it will be the case at some future time that'. The models for this system are equipped with a set T of temporal instants, which play the same formal role as the sets W of possible worlds in the models for Modal Logic; and these sets T always come with additional structure, in the form of a temporal ordering relation $\prec$: for t, t' from T, '$t \prec t'$' means that t is earlier than t'. Formulas are now evaluated for truth or falsity in models at instants of time. Again the truth clauses for the truth-functional connectives are local (e.g. $[[\Phi \ \& \ \Psi]]_{M,t}$

[30]Much is known about how the strength of the logic (i.e. which formulas are logically true and which are logical consequences of which) depends on what general properties the accessibility relations are assumed to have. Much of the work in the early days of modern modal logic – in the fifties, sixties and seventies – was concerned with this question. For a standard work on the topic see (Blackburn, de Rijke & Venema 2001).

$= 1$ iff $[[\Phi]]_{M,t} = 1$ and $[[\Psi]]_{M,t} = 1$), while the clauses for formulas of the forms $P\Phi$ and $F\Phi$ are given by:

($\underline{P}$.1) $[[P\Phi]]_{M,t} = 1$ iff $[[\Phi]]_{M,t'} = 1$ for some $t' \in T_M$ such that $t' \prec t$.
($\underline{P}$.1) $[[F\Phi]]_{M,t} = 1$ iff $[[\Phi]]_{M,t'} = 1$ for some $t' \in T_M$ such that $t \prec t'$.

As we noted for the modal system outlined above, the strength of the logic depends on the assumptions that are made about the properties of $\prec$. (Here is an example: the formulas $P\Phi \rightarrow PP\Phi$ and $F\Phi \rightarrow FF\Phi$ are logically true if it is assumed that the relation $\prec$ is always *dense* (i.e. that for any two elements t, t' of T such that $t \prec t'$ there is a t'' in T such that $t \prec t'' \prec t'$), but not without this assumption.) Like for the modal system described above, there have been extensive investigations of the relation between logical truth and consequence on the one hand and properties of $\prec$ on the other.

The combined system of modal and tense logic that is the focus of this interlude is obtained by adding to classical propositional logic both the modal operator $\Box$ and the tense operators P and F. The models for this system must cater on the one hand to the needs of the tense operators and on the other to those of the operator $\Box$. That is, they must provide a temporal structure as well as a possible world structure, and it is here that in the set-up we have chosen the interaction between worlds and times manifests itself. The guiding intuition in articulating this interaction is this: The formula $\Box\Phi$ should get a formal semantic analysis that justifies its paraphrase as 'it is necessary that Φ is true in virtue of the present and past'. That is, $\Box\Phi$ should be true at a time t in a world w iff Φ is true at t irrespective of how w may evolve from t on.

Here is a way to make this idea formally precise. We assume that each model now supplies a set W of possible worlds, each of which runs from the beginning of (its) time to the end. In principle it is compatible with this assumption that different worlds in W come with different time structures (as a reflection of the different ways in which their histories unfold). And for certain applications of the formalism the possibility of such variation is important. But models in which every world w in W has the same time structure are easier to handle and they are adequate in many contexts, including those we are concerned with here. So we make this simplifying assumption: Every model M comes with a time structure $\mathcal{T}_M = < T_M, \prec_M >$, which is to be regarded as the time structure of each of the worlds in M.

Each model M involves a set of worlds W_M, and each of the worlds w in this set is like a model for the tense logical system described above. So the central notion of our model theory is now that of a formula being true in a model *in* a world *at* a time. The truth-functional connectives are evaluated locally both in the sense of time and in that of possible worlds, thus $[[\Phi \mathbin{\&} \Psi]]_{M,w,t} = 1$ iff $[[\Phi]]_{M,w.t} = 1$ and $[[\Psi]]_{M,w,t} = 1$, for any M, any world $w \in W_M$ and any time $t \in T_M$. The evaluation of the tense operators P and F is local in the sense of worlds, but not (of course) in the sense of time. For instance we get for P:

$$(P.2) \quad [[P\Phi]]_{M,w,t} = 1 \text{ iff } [[\Phi]]_{M,w,t'} = 1 \text{ for some } t' \in T_M \text{ such that } t' \prec t.$$

In contrast to P and F, the evaluation of $\Box$ formulas is local in the sense of time but not local in the sense of worlds. To state the truth conditions of such formulas we need to fill in one piece that is still missing from the structure of our models for the (P,F,$\Box$)-system. It is here that we need a formalization of the notion of two worlds being indistinguishable at least up to t (but possibly diverging later on in time). To denote this relation we use the symbol $\approx$. $\approx$ is a 3-place relation, between a world, a world and a time. That the relation holds between w, w' and t is written thus: '$w \approx_t w'$'. (So '$w \approx_t w'$' means that w and w' coincide at least up to t.) In view of this intended interpretation, it should be clear that $\approx$ should have the following general properties: (i) if t, $t' \in T_M$, $t \prec t'$ and $w \approx_{t'} w'$, then $w \approx_t w'$; (ii) for fixed t the 2-place relation $\approx_t$ is an equivalence relation. (That is: (a) $\approx_t$ is reflexive, i.e. for all $w \in W_M$ $w \approx_t w$; (b) $\approx_t$ is symmetric, i.e. for all $w, w' \in W_M$, if $w \approx_t w'$, then $w' \approx_t w$; and (c) $\approx_t$ is transitive, i.e. for all $w, w', w'' \in W_M$, if $w \approx_t w'$ and $w' \approx_t w''$, then $w \approx_t w''$); (iii) when q is any atomic sentence and t' any time $\preceq t$ and w, w' any worlds from W_M such that $w \approx_t w'$, then q is true in M at w at t' iff q is true in M at w' at t'. (An implicit assumption behind condition (iii) is that the propositional letters of the system represent 'atomic statements', which have no internal temporal structure to them, but simply state what is the case at the time at which they are asserted – statements like 'It is raining' or 'Mary is in Paris'.) We assume that in every model the relation $\approx_M$ has the three properties (i), (ii) and (iii).

We are now in a position to state the truth conditions for formulas of the form $\Box\Phi$:

$$(\Box.3) \quad [[\,\Box\Phi]]_{M,w,t} = 1 \text{ iff } [[\Phi]]_{M,w',t} = 1 \text{ for all } w' \in W_M \text{ such that } w \approx_t w'.$$

The necessity operator characterized by ($\Box$.3) is known as 'historical necessity'.

To get a sense of the expressive capacities of this system, consider a model M, a world w_0 from W_M (which we think of as the actual world) and a time t_0 from T_M (which we think of as the 'present') and certain statements that are made in w_0 at t_0. Let q and q' be two atomic sentences. Then the statement Pq says that q was true at some time in the past of t_0 in w_0 and the statement Fq' that q' will be true at some time in the future of t_0. Assume that in M there is a time $t \prec_M t_0$ such that q is true in M in w_0 at t. Then it should be clear from what has been said about $\approx$ and $\Box$ that in M the formulas Pq and $\Box Pq$ are both true in w_0 at t_0. But on the other hand, if there is a time $t' \succ_M t_0$ such that q' is true in M in w_0 at t, then that is no reason why q' should be true at t' in other worlds w' such that $w_0 \approx_{t_0} w'$. In fact, let us suppose that W_M contains a world w'_0 of just this kind: (a) $w_0 \approx_{t_0} w'_0$, and (b) for no t' such that $t_0 \prec_M t'$ is q' true in M in w'_0 at t'. Then Fq' is true in M in w_0 at t_0, but $\Box Fq'$ is not.

What is illustrated by the behavior of these two formulas – the historical necessity of a statement about the past like Pq is entailed by its truth, but the truth of a statement about the future like Fq' is not – can be generalized to arbitrary statements about the past and arbitrary statements about the future (though it isn't always all that easy to recognize from the structure of a formula whether it is a statement about the past (and nothing but the past) or a statement that is about the future). The possibility of distinguishing between the historical necessity of statements about the past and the historical contingency of many statements about the future is what makes the (P,F,$\Box$)-system of interest in the context of the semantics of future tense expressions (such as the *will* + infinitive construction of English).

But to repeat an earlier observation, providing ourselves with a formal environment in which these distinctions can be formally described cannot be a substitute for a linguistic analysis of how the natural language constructions actually work in the languages we speak. For a concrete example, the logical system we have outlined gives us two potential formalizations of the future tense expressed by English *will*: as '$F\Phi$' and as '$\Box F\Phi$', where Φ expresses the content of the infinitival complement of *will*. These two formalizations correspond to (i) the semantics of the English *will*-future that we have adopted and (ii) the 'deterministic' semantics according to which a future tense statement is true now only when its truth is already fixed by what has been. In this way a formalization of the concept of historical necessity may help us in seeing more clearly what the logical implications are of different semantic proposals; but of course it cannot tell us which of them is right.

3.7.2 The Present Tense

We start with a selection of present tense sentences. Some of these are in the Simple Present and some in the Present Progressive. But on our approach to the semantics of the tense forms they all share the same tense feature, which we refer to as 'pres'.

(3.70)

a. Louise loves Paris.

b. Louise is visiting Paris.

c. Louise visits Paris.

d. Louise is writing a letter/two letters/several letters/some letters.

e. Louise writes a letter/two letters/several letters/some letters/the letter.

f. Louise writes letters.

g. Louise writes several letters a day.

h. Louise plays the violin.

i. Louise is eating an apple.

j. Louise eats an apple.

k. Louise eats an apple a day.

l. Occasionally Louise eats an apple.

m. I am hearing a nightingale.

n. I hear a nightingale.

o. I promise to submit the paper by Friday.

p. I am promising to submit the paper by Friday.

q. And now the moment has come that we have all been waiting for: The Queen steps forward and cuts the ribbon. The bridge is open for general use.

Not all of these sentences are equally ‘good’. Some of them may require a special kind of interpretation or a special kind of context and some may strike you as simply ungrammatical. And some of the sentences will produce different reactions in different readers. The sentences have been left without any ‘goodness’ markers ($\surd$, *, ?, ??), so that you can make up your own mind about which of them are fine, which awkward and which downright out without potential bias due to somebody else’s judgments. But of course I too have my own opinions about these sentences and – for better or worse – it is these that will inform much of what follows in this section.

Of particular importance for what is to be discussed here are the pairs of sentences in (3.70) that are identical except that one is in the Simple Present and the other in the Present Progressive, among them <(3.70.b),(3.70.c)>, (3.70.d),(3.70.e)> and <(3.70.h),(3.70.i)>. Of these three pairs the Present Progressive sentences seem to be straightforwardly acceptable and interpretable. Each of them can be understood as describing something that is going on at the time at which the sentence is uttered; or, according to our analysis of the progressive: the states expressed by the progressive forms in those sentences – the states that get introduced through the application of the +-prog operator – hold at the utterance time. This is something that all three sentences have in common with (3.70.a), which also describes a state, but does so by virtue of the fact that its verb *love* is a lexical state verb. The corresponding Simple Present Tense sentences cannot be interpreted in this same way. In fact, English textbooks for second language learners have had a tendency to mark such sentences as ‘ungrammatical’. That is surely not quite right. But what is true is that for the typical English speaker these sentences aren’t the right ways to say what is naturally expressed by their Present Progressive counterparts: that a certain activity that the sentence is used to describe is going on right now, at the time when the speaker produces her utterance. The Simple Present members of the pairs can have their own uses, however. For instance, (3.70.c) can be used to say that Louise goes to visit Paris ever so often, or on a regular basis; and (3.70.i) could occur as a stage instruction in a play (although this doesn’t feel completely natural; more easily understood in this way would be ‘Louise picks up an apple from the bowl on the table’). (3.70.e) is even harder to understand as a stage instruction, presumably because writing a letter is thought of as taking even more time than eating an apple; but even in this case such a use doesn’t seem completely out of the question..

There is also a striking difference between (3.70.e) (in each of its four alternative versions, with singular direct object or any one of the three plural ones) and (3.70.f). (3.70.f) is a natural way of describing a habit or disposition of Louise, something she does now and then, or that she usually or typically does under certain conditions, or towards which she has a natural propensity. In that respect (3.70.f) is similar to (3.70.c). But we need an explanation why the verb *write* needs a bare plural direct object to bring about this effect, whereas with *visit* a 'referential' direct object like *Paris* will do as well. (We won't try to come up with an explanation of this particular difference here.)

The contrast between (3.70.e) and (3.70.f) finds an echo in that between (3.70.i) and (3.70.j). While (3.70.i) invites some kind of stage-like setting, (3.70.j) requires no such interpretation. It can be understood as a plain statement to the effect that the eating of an apple is something that Louise does occasionally.

Before we look at the remaining examples in (3.70) let us try and see if we can recognize a principle that justifies the judgments about the ones considered thus far. First a declaration about how we will proceed. For the time being we will focus exclusively on uses of the Present Tense that serve to describe what is the case or what is going on at the time of utterance. We will refer to this use (somewhat tendentiously) as the 'Standard Use' for the Present Tense. (The English Present Tense also has certain other uses, in which it describes past or future eventualities. Some things will be said about these uses towards the end of this section. But for now we set those uses aside. We will turn to them at the end of the present Subsection.)

Internal vs. External Viewpoint Aspect

There is a principle that governs the Standard Use of the Present Tense and that is found in language after language. The principle reflects the special relationship in which speakers stand to the information about which they speak as they observe it. But it manifests itself with particular prominence in a language like English where the distinction between Perfective and Imperfective Aspect is marked for event verbs through the presence or absence of the progressive. In fact, we already saw the principle in action when we made the familiar observation that when you want to use an event verb to describe what is going on while you are speaking, you have to use the progressive form; what you should say is 'She is eating an apple', not 'She eats an apple'.

As this observation may suggest, one way to state the principle is that the Standard use of the Simple Present requires Imperfective Aspect:

(3.71) The Standard Use of the Present Tense is restricted to input representations with Imperfective Aspect.

Given our identification of Imperfective aspect with the description of states, this principle can also be formulated as the rule that state descriptions are appropriate as inputs to the Present Tense in its Standard Use, but event descriptions are not. This means that in its Standard Use the Present Tense accepts two kinds of eventuality descriptions: (i) eventuality descriptions that are state descriptions because their verb is a stative verb (as in the case of the verb *love*) or (ii) eventuality descriptions that are state descriptions because they involve an operator like the progressive, which turns event descriptions into state descriptions. Event descriptions as such are not compatible with the standard use, but they can be turned into state descriptions and thus become suitable input representations for the present tense. As some of the examples in (3.70) show, applying the progressive is not the only way to turn event descriptions into state descriptions that can serve as inputs for the Standard Use of the Present Tense. We will turn to such other operations presently.

Principle (3.71) can be seen as a special case of a more general principle which is connected with an aspectual opposition that we haves not yet encountered. This opposition, prominent in the work of Carlota Smith ((Smith 1991)), is that between *External Viewpoint Aspect* and *Internal Viewpoint Aspect*. Smith noted that there are two ways of looking at what was, is or will be going on at a time, or over some period of time: from the outside and from the inside. With these two modes of looking come two ways of describing what one is looking at, the internal and the external mode of describing eventualities. Smith uses the terms 'External Viewpoint' and 'Internal Viewpoint' to refer to these two perspectives and she uses the terms 'External Viewpoint Aspect' and 'Internal Viewpoint Aspect' to refer to the description modes that these two viewpoints invite and require. In some languages the distinction between External and Internal Viewpoint Aspect descriptions is overtly marked and the ways of marking them are the very same that are used to mark the distinction between what we have been referring to as that between Perfective and Imperfective Aspect. It is not entirely clear to me whether these two oppositions – that between Perfective and Imperfective Aspect and that between External and Internal Viewpoint Aspect – line up perfectly. But it seems quite plausible that Internal Viewpoint always requires Imperfective Aspect, and it is that assumption which entails Principle

(3.71). The connection is this: whereas for past and future eventualities there is in general a choice between describing from an internal or an external viewpoint, for eventualities that hold or are going on at the time of speech only the internal viewpoint is possible: the speaker is in the midst of what is going and so has no choice: her viewpoint can only be the internal one. Therefore, since Internal Viewpoint requires Imperfective Aspect, the only appropriate way to describe what is going on while one speaks is to use an imperfective eventuality description.

Many event verb sentences in the non-progressive present tense seem weird or simply ungrammatical. But some of the examples in (3.70) show that this is by no means true for all of them. Among the simple present tense sentences that have felicitous uses are (3.70.c), (3.70.f) and (3.70.h). But note well that they too can be seen as conforming to Principle (3.71), and that they do so because their interpretation involves a reinterpretation of the event descriptions that are their ostensible inputs to the Present Tense. In this regard they resemble present progressive sentences. But they also differ from sentences in the present progressive, and in fact they do so in two respects. First, the transformation of event descriptions into state descriptions that is effected by the progressive operator is overtly expressed (by progressive morphology), whereas the transformations involved in the interpretation of sentences like(3.70.c), (3.70.f) and (3.70.h) are not overtly marked. Second, the semantics of the state descriptions resulting from these latter transformations is quite different from the semantics of the progressive. The states described by the sentences in (3.70.c), (3.70.f) and (3.70.h) are *habitual* or *dispositional* states, states to the effect that events of the kind described by the untransformed event description occur 'habitually', or to the effect that there is a disposition towards the occurrence of such events.

Habitual and dispositional reinterpretations of event descriptions are typical for the Standard Use of the Simple Present Tense. Since this use of the Present Tense does not accept event descriptions as inputs, it isn't surprising that such sentences should involve state descriptions of some kind or other; for if it weren't possible to reinterpret event descriptions as state descriptions in some way, then such sentences wouldn't be interpretable at all. What we see here, in other words, is an instance of a phenomenon of which we encountered another manifestation in our discussion of the Progressive in Section 3.5.2, viz. coercion. In Section 3.5.2 it was argued that the Progressive does not accept state descriptions, but that sometimes state descriptions can be reinterpreted as event descriptions, which can then serve as proper inputs to the progressive operator. What we see illustrated by (3.70.c), (3.70.f) and

(3.70.h) is coercion too, though it is coercion 'in the opposite direction' – from event descriptions to state descriptions.[31]

For a formal treatment of sentences involving habitual and dispositional reinterpretations of event descriptions we need predicates in our DRS language that enable us to represent the results of those reinterpretations. To see what is needed let us begin by focusing on sentence (3.70.c). This sentence can be characterized informally as describing Louise's habit of visiting Paris – or, in the more technical vocabulary we have been using, as describing the state of affairs that consists in Louise having this habit. To clothe this still informal characterization in formal dress we include in our representation language an operator HAB which turns event descriptions into state descriptions, where the resulting state description is to the effect that events which fit the input description occur on a frequent or regular basis. Thus, (3.70.c) can be analyzed as involving the application of HAB to the event description '$^\wedge e.e$: $\underline{x}$ visit Paris'.

If reinterpretation of underlying event descriptions takes the form of subjecting them to the application of HAB, then there has to be some point in the construction of the semantics of the complete sentence at which this application takes place. Clearly this point must come before temporal location by tense, since it is the reinterpreted eventuality description that gets temporally located by the Present Tense. In other words, the syntactic level at which this operation takers place must be below the node T'. But exactly how far below T'? That is a non-trivial question, which has to do with details of the structure of VPs that we do not want to get into here. Presumably the HAB operator interacts with some components of more detailed and refined VP structure and so will have to be inserted into VP structures at some level that isn't represented in the structures we have been assuming in these notes and that we will continue to assume throughout. For the purposes of these Notes we will assume that HAB is triggered at the level of Asp, and, more specifically, that it is triggered by a new value of the ASP feature, to which

[31]Note well: to describe a certain interpretation as the 'result of coercion' doesn't account for why the eventuality descriptions that result from the coerced reinterpretation have the meanings they have. The Standard Use of the Present Tense may require a state description of some kind and thus 'coerce' the reinterpretation of what are event descriptions to begin with. But why should the results be descriptions of habitual or dispositional states? This is an aspect of the semantics of sentences like those in (3.70.c), (3.70.f) and (3.70.h) for which I do not have a satisfactory explanation; and perhaps no truly satisfactory explanation can be given. As things stand, all we can do is accept that coercion by the Simple present Tense takes these forms, in the possibly vain hope for an explanatory story that will be told by someone else.

we will refer as +hab. We will briefly return to this decision below. On this assumption the input representations of HAB are the representations of the sister nodes to the Asp nodes, in other words, of the (highest) VP nodes. This means that in the case of (3.70.c) the input representation to HAB is as in (3.72):

$$(3.72) < e_{ref}, p \mid \begin{array}{|c|}\hline \\ \hline Paris'(p) \\ e : visit'(\underline{x}, p) \\ \hline \end{array} >$$

HAB then turns this representation into the one in (3.73)

$$(3.73) < s_{ref}, p \mid \begin{array}{|c|}\hline \\ \hline s\text{: } HAB(\hat{}e. \begin{array}{|c|}\hline e \\ \hline Paris'(p) \\ e : visit'(\underline{x}, p) \\ \hline \end{array}) \\ \hline \end{array} >$$

The intuitive meaning of the DRS Condition beginning with 's:' in (3.73) is that s is a state to the effect that the represented habit can be said to exist over the course of the state's duration.

What makes a state S a state of the kind specified in (3.73) depends on the distribution of events of the sort specified by (3.72) over the duration of S. But what should that distribution be like? How many events of the specified sort must occur within the duration t of S, or with what kind of regularity – how widely or narrowly spaced in time – or on what sorts of occasions within that period? These questions cannot be answered in general terms. How many events of the specified kind must occur within t and how distributed will depends on the kind. Even when the distribution of such events over t is fully known, speakers may vacillate in their judgments or differs from each other. There appears to be one minimal condition, however, that any habitual state should satisfy: there must be within its duration at least some occurrences of events of the specified sort. We will adopt this condition, by adding a Meaning Postulate expressing it to our given stock of Meaning Postulates. is straightforward The formulation of such a Meaning Postulate is straightforward and left to the reader.

A habitual interpretation also seems the most natural option for (3.70.f). But here a somewhat different interpretation seems possible as well, according to which the sentence ascribes a certain *disposition* to Louise – e.g. the disposition to reply, when you contact her, by letter rather than by e-mail or telephone. (3.70.h) is similar to (3.70.f) in this regard. It can be understood as saying that Louise plays the violin on a regular basis, or simply that she *can* play the violin – the violin is her instrument – even if she hasn't actually played it for years. Dispositional interpretations differ from habitual interpretations in that they can be true even if there are no instances of the specified kind of event within the relevant period of time. A well-known example from the semantic literature con be found in (Carlson & Pelletier 1995*a*) (p. 72). They note that the sentence 'Mary handles the mail from Antarctica' may be true even if the period t to which it applies contains no occasions of Mary actually handling any mail from this virtually uninhabited part of the globe; it suffices that Mary was put in charge of the Antarctica Department throughout t. Other compelling illustrations that a dispositional claim can be true even though the can be found in the Philosophy of Science literature. Classical examples of dispositional properties that can be true over a period t without being instantiated during t are those that we describe with the verb *dissolve*. Consider for instance the event description 'dissolves in sulphuric acid'. If you point at a thing and say: 'This thing dissolves in sulphuric acid', the natural interpretation of what your statement is that the thing you refer to has the dispositional property of dissolving when it is put into sulphuric acid. If this statement is true, then the thing cannot have been involved in an actual dissolving event, for in that case it would no longer exist. So having this dispositional property – that of being involved in an event of being immersed into sulphuric acid and dissolving – *entails* that the event description hasn't been instantiated.

We can formally distinguish between habitual and dispositional interpretations by introducing besides the predicate HAB a dispositional predicate DISP, with the condition 's: $DISP(^\wedge ev.K)$' expressing that s is a state of the relevant individual having a disposition towards displaying instances of the eventuality property $^\wedge ev.K$. Again this proposal is no more than a representational shell, within which something of substance could be articulated at some later point about what it is to have a disposition of a certain kind, but which doesn't say any such thing in and of itself. In fact, our formal treatment isn't saying anything about the semantic differences between habitual and dispositional interpretations. The only difference it makes explicit consists in adopting 'instantiation Meaning Postulate' of the one but last paragraph for HAB but not for DISP.

In the syntax-semantics interface architecture we are developing, reinterpretation of event descriptions as descriptions of dispositional states should have a trigger just as habitual reinterpretations. We have already made the assumption that habitual reinterpretations are triggered by the feature value +hab of the feature ASP. If dispositional reinterpretations are to be triggered in the same way, then we need a further value for ASP, to which we will refer as +dis. This is the assumption we will adopt.[32]

We briefly return to the decision we made above about how and where the reinterpretation of event descriptions as descriptions of habitual or dispositional states is determined. We argued that this determination must occur below the node T', but then presented our actual decision – the reinterpretation is triggered by a feature value of the feature ASP – without further argument. The motivation for making the Asp node responsible for the transformation of event descriptions into habitual or dispositional state descriptions is the following consideration. The range of sentence interpretations that our construction algorithm is meant to be able to handle at this point includes sentences in which event descriptions are transformed into progressive state descriptions, habitual state descriptions or dispositional state descriptions as well as sentences that require no transformation, but where the event description is passed on to the next higher node unaltered. These are four mutually exclusive options. Their being mutually exclusive is most easily captured by setting things up in such a way that the choice is made at one single point in the course of representation construction, according with information that is available at that particular point. Given our earlier decision to make the Asp node the locus of the choice between progressive and non-progressive interpretations, the only natural move now that the repertoire of choices has been extended from two to four is to retain Asp as the point where this now four way choice is made. The only difference between the original two options and the ones that have just been added is that the former is explicitly marked by progressive morphology. The choice between the other three – habitual reinterpretation, dispositional reinterpretation and no change – is not marked.

[32]With this last assumption the value range of ASP has grown to a set of four elements: +prog, -prog, +hab, +dis. At this point the term '-prog' is somewhat misleading, since it no longer denotes the complement of the feature value +prog. As member of the extended, four-element value set for ASP -prog still triggers an interpretation which passes its input representations unchanged up to AspP. But this operation is now in opposition not only to +prog, but also to +hab and +dis.

So the choice that the parser has to make between the ASP values +hab, +dis and -prog is one it has to make in some other way. We cannot exclude the possibility that the parser will sometimes make the wrong choice between these three feature values, which will then lead to uninterpretability higher up in the semantics construction process. (For instance, when the parser has chosen -prog as value for ASP, but the semantic representation of the VP is an event description and the sentence is an instance of the Standard Use of the Simple Present Tense, then the representation construction will abort at the point where the tense information must be combined with the semantic representation of the sister node to T; for that semantic representation will be an event description, but it should have been a state description.)
It might be thought that the reinterpretation of event descriptions as habitual or dispositional state descriptions could be handled in a different way, which would avoid the possibility of semantic representation constructions that abort because the parser made the wrong choice when picking value for ASP. Could we not leave the decision to reinterpret to the point where the tense information is combined with the semantic representation of T's sister? The effect of that would be that if the tense information is to the effect that we are dealing with a Standard Use of the Simple Present and the sister representation is an event description, then a coercion step – reinterpretation of the event description as that if a habitual or dispositional state – will have to be performed (if possible) to the input representation before the tense information is combined with the result of this step. The possibility and form of such event-to-state coercions would then have to be incorporated into the lexical entry for the Present Tense in its Standard Use. But that is, one might want to argue, exactly where this information belongs, since the constraints involved are precisely what sets Standard Use of the Simple Present apart from other tenses (and from other uses of the Present Tense).

Unfortunately this alternative solution won't quite work without further specifications that rob it of much of its attractiveness. This is because other tenses, among them in particular the Simple Past, also allow for habitual or dispositional reinterpretations of event descriptions, although they do not require them. If coercion of such reinterpretations in Standardly Used Simple Present Tense sentences is incorporated into the lexical entry for this particular Tense Form, then room will still have to be made for the optional reinterpretations that often exist when the tense is different. In the light of this the better strategy would seem to be to allow the choice for such reinterpretations to be made at some suitable point independently of tense, in combination with a principle that the input representations to the Simple Present in its Standard Use must always be state descriptions.

In the interlude following below we have a fleeting look at habitual and dispositional interpretations in non-present tense sentences.

Before moving to this interlude, one further observation about habitual interpretations. The sentences (3.70.k) and (3.70.l) are also examples of sentences whose semantics involves the transformation of event descriptions into state descriptions, and where the resulting state descriptions have much the same flavor as the habitual interpretations we have discussed. But the difference is that the transformations into state descriptions in these last two sentences are the result of overt quantification (triggered by the adverb *occasionally* in (3.70.k) and by the adverbially used DP *a day* in (3.70.j)). We will leave such cases of 'temporal quantification', where quantification os over times or eventualities, for now, but will turn to them later, in Sections 3.11.4 and 3.11.5.

Interlude: Habituals, Dispositionals and Temporal Quantifications with non-present Tenses

Now that habitual, dispositional and quantificational readings have been mentioned, this is as good a place as any to dwell a little on the fact that the Simple Present Tense is not the only tense form that licenses such readings. Non-progressive past tense and future tense event descriptions can get habitual or dispositional interpretations too. But there is a difference: in past and future tense clauses episodic interpretations are possible, but in general such interpretations compete with 'episodic interpretations', in which the initial event description is not reinterpreted. In fact, other things being equal, there appears to be a preference for episodic interpretations, suggesting that that is the 'default option'.

For some event clauses in the Simple Past Tense only an episodic interpretation seems to be possible (at least in neutral contexts, in which there isn't a particularly heavy pressure in the direction of a non-episodic interpretation). As an illustration compare (3.74.a) with (3.74.c).

(3.74)a. Louise ate an apple.

b. Louise eats an apple. (= (3.70.j))

c. He made dinner. She did the dishes.

d. He makes dinner. She does the dishes.

(3.74.a) only seems to allow for an episodic reading: at some (not further specified) pats time there was an event of Louise eating an apple. But (3.74.c) can be interpreted both episodically – about how the two dealt with dinner on one particular occasion – and non-episodically: as a description of how the two handled the problem of dinner in general; he would make dinner and she would do the dishes. Note that this difference between (3.74.a) and (3.74.c) pairs up with the difference between the present tense sentences (3.74.b) and (3.74.d): (3.74.d) is naturally and unproblematically understood as describing the habitual: he makes dinner and she does the dishes, that's how they go about dinner generally. But (3.74.b) was our paradigm example in the list (3.70) of a sentence that somehow seems to defy a felicitous reading altogether. It is a plausible conjecture that this alignment – between (3.74.a) and (3.74.c) on the one hand and (3.74.b) and (3.74.d) on the other – holds generally: an event description in the simple past tense only has an episodic reading iff the corresponding present tense sentence has no reading at all. (Of course, even if this conjecture should be correct, that won't help us in defining the difference between those event descriptions that pattern with (3.74.a,b) and those that pattern with (3.74.c,d). Finding an interesting definition of this difference is to my knowledge still an open problem.

As we noted, one difference between present tense and non-present tense sentences whose VPs express event descriptions is that when a habitual or dispositional reinterpretation is possible at all, then the present tense sentence will get that and only that interpretation, whereas its non-present tense counterparts are ambiguous between this interpretation and an episodic one. However, for past tense sentences English has a couple of simple devices that resolve this ambiguity in favor of the non-episodic reading. One of them is the word *would* and the other the aspectual verb *used to*. (3.75) gives the results of using these as alternatives to (3.74.c).

(3.75)a. He would make dinner. She would do the dishes.

b. He used to make dinner. She used to do the dishes.

c. Louise would eat an apple.

d. Louise used to eat an apple.

Note also that applying *would* or *used to* to an event description that seems to allow only for an episodic interpretation when it is used in the simple past eliminates the option of an episodic reading, thus pushing the interpreter in the direction of a non-episodic one.[33] When *would* or *used to* are applied to an event description like *eat an apple* the result seems not optimally felicitous (though it might be in the right context). This too seems to correlate with the observations made in connection with (3.74).

When the Simple Present Tense locates events at utterance time

In what we have been calling the 'Standard Use' of the Present Tense, present tense sentences with non-progressivised event descriptions either get a habitual or dispositional interpretation or they get no interpretation at all. But there are Simple Present Tense uses of event descriptions which serve to describe events located at the utterance time and which are felicitous. One such use of the Simple Present is what is known as its *reportive use*. An example is (3.70.p). In reportive uses of the Simple Present the time of the described event is to be understood as coinciding with the utterance that reports it. Paradigmatic examples are on-site reports that are broadcast (in real time) over the radio. The reportive use differs from the Standard Use in that each new sentence utterance that is part of a discourse in which the present tense is used reportively defines its own utterance time (thereby giving the interpreter a feeling that he is witnessing the sequence of reported events as they evolve in the 'real time' defined by the reporter's successive utterances). The reportive use can be characterized as combining the External Viewpoint aspect – the events are presented as completed wholes, which the reporter can report in this form because they have just been completed as she speaks – with the basic function of the present tense, of locating the described eventuality at the utterance time. Part of this ploy is that each successive utterance defines its own utterance time – put in the terms of our DRT-based account of discourse interpretation, each determines its own n value. So it is the time defined by the utterance event that is to be understood as locating the event described in the utterance, and not the utterance time of the discourse as a whole that the individual utterances are part of. Note well that nothing that was said before the present subsection about how the Present Tense temporally locates the eventualities it describes entails that the Present Tense can be used in the way of the reportive present. But given that the Present Tense can be used in this way, it is not surprising that sentences in which it is so used employ the non-progressivized form to convey their message: the

[33] The *would* form also has a 'future of the past' interpretation, according to which the described event is in the future of some past reference point. This form can be discussed properly only after we will have introduced presuppositions into our representation language; it will be addressed in Section 4. The present discussion ignores this use of *would*.

speaker's mental registration of a complete event of the kind described at the time of the utterance she produces as a verbal reflection of that registration. The utterance doesn't represent the event as going on at the time when it (the utterance) happens to be made. Rather it defines its own time, which manifests itself to the listener via the speaker's registration of it sand the utterance that is the direct outflow of that registration.

The reportive use of the Simple Present Tense is not the only one that locates events at the time of utterance. There are two other types of examples of this that we briefly mention, each of which is associated with a particular use of a particular type of verb. (As far as I know there are no special names for these uses in the literature, but the naming question need not detain us.) The first type of verbs that allow for episodic event-describing uses of the present tense are so-called *performative verbs*; the second type consists of *perception verbs*. Illustrations of such uses of the two types of verbs can be found among the examples in (3.70), (3.70.o) for the performative use of performative verbs and (3.70.n) for the 'direct perception report' use of perception verbs.

Performative uses of speech act verbs, in which the verb occurs in the first person present tense, are speech acts in which the verb is used to describe the act that is performed by uttering the sentence of which it is part. For instance, an utterance of (3.70.o) is understood as the *making* of a promise, and the verb *promise* is used to describe that very act. In such cases the utterance is identical with the event described by its verb. As with the reportive uses just discussed, this rules out both past and the future tenses as tense markings of the describing verb – since the reported event coincides with the utterance, it follows that it can be neither in the past nor in the future of the time of the utterance. Furthermore, like in the reportive use, the utterance serves to identify the time of the event: in particular, if utterance and event are one and the same thing, then of course their times must be identical too. This entails – trivially – that the event time is temporally included in the utterance time.

The use of the Simple Present in perception reports like (3.70.n) invites an account along somewhat similar lines. Note that the events that such sentences describe are perceptions. A typical use of a sentence like (3.70.n) is to describe a perception at the very instant the speaker becomes aware of what she perceives, as if the utterance is a form in which the perceptual experience becomes public – the utterance counts as a direct manifestation of the perception. On this view utterance and perception coincide in much the same way as performative utterances and the acts that they describe, and thus will coincide in particular in time.

This special character of the utterances of sentences like (3.70.n) and (3.70.o) becomes more clearly visible when we compare these sentences with their present progressive variants, as in (3.70.m) and (3.70.p). The contrast between (3.70.o) and (3.70.p) is easy to appreciate. (3.70.p) is a rather odd sentence. (Here it is the present progressive version that seems strange, in contrast to the pairs <(3.70.b),(3.70.c)> and <(3.70.j),(3.70.k)>, where the non-progressive variant seems the odd one out.) If (3.70.p) has any use at all, it is one in which the sentence describes an act that is distinct from the one that is identified by its own utterance. An example would be something like this: the speaker is writing a letter in which she makes the promise to submit her paper by the mentioned date, and while she is doing this, she utters sentence (3.70.p) in reply to someone who has just come into the room and asks what she is doing.

In the case of perception verbs the situation is different. (3.70.m) is a perfectly natural sentence, no less than (3.70.n). But the uses of the two sentences are different. As I said above, a natural use for (3.70.n) is that where the speaker signals her becoming aware of the nightingale she hears. (3.70.m) on the other hand seems to be preferred in a situation where the speaker has been aware of the nightingale for some time already. The sentence has the status of a testimony to the speaker's awareness of her perception, rather than a direct manifestation of that perception. (For instance, the speaker might say (3.70.m) in response to a query about the rapt expression on her face, or why she appears to be paying no attention to what her interlocutor is telling her.) But the difference between (3.70.m) and (3.70.n) is subtle, and nothing like the easily graspable difference between (3.70.n) and (3.70.o).

When Present Tense locates Eventualities at other Times than n

So far we have looked at interpretations of present tense sentences in which the tense is understood as locating the described event or state as occurring or holding at the utterance time. But in addition the English Present Tense can also be used to locate the described eventuality at times other than that of the utterance (or of the current discourse of which the utterance is part). On the one hand there is the so-called 'historical present'. The historical present is a rhetorical device designed to impart liveliness to a report of past events: by using the present tense instead of a past tense the author confers upon the reader a sense of being present at the events she recounts, as a kind of eye witness. The classical example of this ploy is Julius Caesar's 'de Bello Gallico' the report that Caesar sent to the Roman Senate in order to inform it about his accomplishments as Governor of Gaul. (Since Caesar always wins in this report, in battle after battle and siege after siege, his story has a tendency to get rather monotonous. The author is to be commended for having seen this himself and for having realized that a little stylistic livening up of his narrative would not go amiss. (Unfortunately, the powers of the historical present are limited. In Caesar's extensive use of it the historical present soon turns into a tedium of its own making.)

The Historical Present, however, isn't restricted to literary works (assuming that *De Bello Gallico* qualifies as such). They can also be employed in spoken language and often are. The following report of a election canvasser might be a natural example.

(3.76) So I go up to this house and ring the bell. And at first nothing happens. And then I ring the bell again and the door flings open and a bullet whistles past my head. 'What do you want?' the guy says. 'I am from the Democratic Party' I say, we are trying to talk to people in this neighborhood.' 'Get off my property' the guy says, 'or you'll find out I am a better shot than you thought.' So I back off and it is only after I am past the gate that I dare to turn my back to him.

The English Present Tense can also be used to refer to the future. In English this possibility is rather restricted, much more so than in many other languages, including close cousins like German or Dutch. The restricted futurite use of the Present Tense in English is known as its 'time table use'. It can be employed when the future event it describes is already settled at utterance time: either by some law-like natural process or else by an official human arrangement or decree. Among the familiar examples are the sentences like those in (3.77.a,b).

(3.77) a. The sun rises at 7.42 tomorrow.

b. The train for Paris leaves at 12.44.

c. Chelsea plays Arsenal next Sunday.

The typical uses of (3.77.a) and (3.77.b) are those in which *7.42 tomorrow* and *12.44* refer to times after the utterance time, but where it is already determined that the described events – the rising of the sun and the departure of the train for Paris – will occur, and that they will occur at the times at which the sentences say they are going to happen.

In the absence of a natural process or public arrangement that settles the future occurrence of an eventuality at the mentioned time the Simple Present cannot be used to describe this occurrence. Note well that this is so even when there is a strong commitment to the occurrence of the eventuality. For instance, the subject of the sentences in (3.78) may be fully resolved to accept the position that has been offered to him but has decided to ask for the later starting date. Even under those conditions what is felicitously expressed in (3.78.a) cannot be stated in the form of (3.78.b).

(3.78)a. He will take the job, but he will ask them if he can start a little later.

b. He takes the job, but he asks them if he can start a little later.

It is not easy to say how the situations in which the time table use of the English Present Tense is possible differ from those in which it is not. One feature that the salient examples of permissible instances of the time table use that are known to me have in common is that the future occurrence of the eventuality as such is presupposed; what the sentence asserts is just at what particular time the eventuality will occur. The time table use is then felicitous iff at the moment when the utterance is made the time at which the event will occur is fixed as well as the mere fact that an event of described sort will occur. But I am unsure whether this presuppositional dimension is part of all legitimate time table uses. Further explorations will haves to decide if this hypothesis holds in general.

A Rump Lexical Entry for pres

The survey of different uses of the Present Tense in this section should have made it plain that a comprehensive formal treatment of this tense form within the syntax-semantics interface we are developing would be quite difficult. The survey has shown (i) that the Present tense can be used not only to locate the described eventuality at the utterance time but also at times in its past and in its future. And second (ii) there is a considerable range of different uses all of which locate the described eventuality as coincident (in some sense or other) with the utterance time. A comprehensive treatment of all these possibilities would not only have to provide formally correct accounts of each of them but would also have to address the question how the interpretation process can identify the instances of those different possible uses. A mere specification of the times to which occurrences of the different uses of the Present Tense can anchor the eventualities they can be understood to describe, but which has nothing to say about which Present Tense occurrences are instances of which of its uses, will make the lexical semantics for the Present Tense void of practical applications: when the entry provides no means for recognizing which occurrences of the present Tense instantiate which of its possible uses, any eventuality described in any Present Tense sentence could be located anywhere in time.

The solution for which we will opt errs in the opposite direction; it simply ignores many of the uses to which the Present Tense can be put. First, we will confine ourselves here with a simplified lexical entry for the Present Tense which ignores those uses that locate at time before or after the utterance time. But as we have seen, even when we restrict the semantics of the Present tense in this manner the difficulties are far from over. A distinction is to be made between those cases in which the input to the present tense feature is a state description and those in which it is not. The latter cases, we saw, include on the one hand uses of the Present Tense as reportive present, and on the other utterances involving special categories of verbs – perception verbs and speech act verbs – which stand in special, witness-like relations to the events they describe. A lexical entrance capable of dealing with all those uses of the Present Tense, which locate an event at the utterance time, should offer the means to recognize the different instances of this more restricted range of options. Since even this task is beyond our present capacities, we will reduce the scope of our lexical entry for the Present Tense further yet, by setting aside also all cases in which Present Tense sentences locate *events* at the utterance time.

What we are left with is an entry for the Present tense which comes with a selection restriction to the effect that the input representation must be a state description. That entry will have the effect that if the input representation is the description of an event, 'coercion' to a state description (habitual or dispositional) is needed. Because of the way in which we have chosen to implement the triggering mechanism for such 'coercions' (by the values +hab and +dis of the feature ASP), the needed 'coercions' are taken care of automatically if the parser delivers the appropriate choices of those feature values. So no information relating to possibilities and modes of coercion need to be added to such a restricted lexical entry. So see end up with a quite simple entry for the Present Tense, albeit at the price of severe truncation of its actual range.

The Present Tense entry that is the result of these simplifications is shown in (3.79). In keeping with our entries for other tense forms it is assumed that the Present Tense form is the morphological expression of a value of the feature Tense that is associated with the node T. the name of this feature value is 'pres'.

(3.79) (lexical entry for the tense feature 'pres')

pres (tense feature)

Sel. Restr: state description

Sem.Repr: $<ev_{ref}, ... \mid K> \rightsquigarrow$

$$<t, ev_{ref}, ... \mid K \cup \begin{array}{|c|}\hline \\ \hline t = n \\ t \subseteq ev \\ \hline\end{array} >$$

3.8 Perfects

In Section 3.5.2 we saw that the English Progressive is an aspect operator which is distinct from Tense and which appears below Tense in the syntactic trees for sentences with progressive forms. Arguments similar to those we gave in favor of the syntactic analysis of Progressive sentences also apply to the Perfect. Perfect tense forms come in a similar variety as progressive forms – there are the Past Perfect and the Future Perfect as well as the Present Perfect; and we will see below that there also is a common semantic element to all of these, which sets them apart from the corresponding non-perfect forms. These considerations strongly suggest that the Perfect should be treated, like the Progressive, as an operator that is distinct from the category of tense and that makes its own, independent semantic contribution. Furthermore, again as with the progressive, English morphology suggests that in the syntactic construction of a sentence or clause with a perfect tense form, the Perfect makes its entry before Tense. For perfect formation involves putting the verb in its past participle form and combining that with the auxiliary *have*; and it is then this auxiliary which becomes the morphological carrier of finite tense (just as finite tense is put on the auxiliary *be* in the formation of a sentence with a progressive tense form).

The final point we have to settle is the order of application between perfect and progressive in those sentences and clauses which have both. That they can co-occur is plain, and illustrated in (3.80.a). But it is also clear from these examples that the order in which progressive, perfect and tense must have been put on is that in which I am listing them right now: the progressive comes first, turning the bare infinite form of the verb into its present participle form and adding the auxiliary *be*; next comes the perfect which turns *be* into its past participle form *been* while adding the auxiliary *have*; and finally the tense is put on *have*. Applying perfect and progressive in the reverse order is not possible, as demonstrated in (3.80.b).

(3.80)a. Frieda has/had/will have been closing the shop.

b. * Frieda is/was/will be having closed the shop.

To deal formally with the conclusions from these observations we need to assume yet another projection level in the syntactic structure of finite clauses – that of *Perfect Phrases*, or *PerfP*'s – which is situated between the AspP level and the TP level and at which the perfect is introduced into those clauses that have it. We encode the information as to whether a perfect is introduced at this level in terms of a feature '+perf', indicating the introduction of the perfect, and -perf to indicate that no perfect is introduced. Thus, for instance, the syntactic form of the first of the three sentences compressed into in (3.80.a) is as in (3.81).

(3.81)

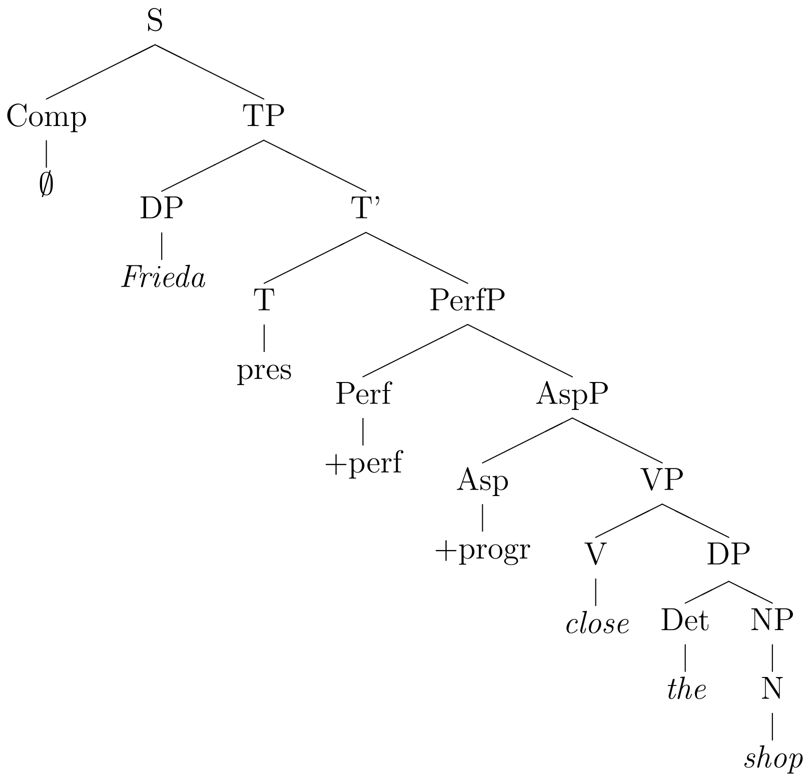

So much for the syntax. The first point we can make about the semantics of the perfect is that present tense uses are unproblematic and that that is as true for non-progressive as it is for progressive perfects. For instance, (3.82) is just as grammatical as the sentence represented in (3.81).

(3.82) Frieda has closed the shop.

This suggests that perfects are state descriptions. And there is a very simple explanation of why that should be so. Perfects describe *result states*, states that are the result of the occurrence of an event instantiating the underlying non-perfect eventuality description. It is the result states of such eventualities that perfect forms describe and that are then located by tense. The result state s of a given event e starts the very instant at which e comes to an end and then goes on for some time after that.

For sentences in the Present Perfect this implies that an eventuality of the kind described by the input to the perfect must have happened at some time before n. In this regard Present Perfect sentences are much like Simple Past sentences with the same verb and arguments. For instance, the truth conditions of (3.82) are similar to those of (3.83). Both are true, when uttered at a time t only if there was an event of Frieda closing the shop that preceded t.

(3.83) Frieda closed the shop.

But how similar are the truth conditions of these two sentences? If the result state s of an event e goes on forever once it has started (and, thus, once e has come to an end), then the truth conditions for (3.81) collapse into those that our account attributes to the Simple Past sentence in (3.83). But do result states always go on forever? That depends. For one thing, languages differ on this point. For instance, current French and German impose few if any constraints on result states. Here the truth of a Present Perfect sentence amounts to no more than the past occurrence of the described event, and the truth conditions are thus indistinguishable from those of the corresponding sentence in the Simple Past. In fact, in these languages there is a tendency to use the Present Perfect in situations where at an earlier development stage of the language one would have used the Simple Past. (This kind of diachronic drift, in which Simple Past forms are gradually pushed out by Present Perfect forms, has been documented for many languages; see for instance, (Dahl 1985), (Bybee & Dahl 1989).) But the English Present Perfect is different in this respect. In order that the result state of an event of a given kind obtains at some time t after that event has occurred more is required than the mere fact that there was a previous occurrence of e. For the result state to still hold at some time after the occurrence of e, some kind of causal aftereffects of e must continue to obtain. For instance – to give just one example – you can describe someone as ‘having left the house’ when she went out of the house but has not yet come back. But after the person has come back, to say she ‘has’ left seems infelicitous; compare (3.84.b) with (3.84.a). (Note that the Present Perfect in the second sentence of (3.84.a) and (3.84.b) is unproblematic so long as it is assumed that at the time of utterance Frieda hadn’t left the house for a second time.)[34]

(3.84)

[34]Recently, when presenting contrasting pairs like that in (3.84) to students I have found a striking (and disturbing) diversity of judgments about the comparative acceptability of the discourses making up such pairs. I do not quite know what to make of this. One possibility is that the sharp distinction between Present Perfect and Simple Past that is supposed to be part of the native speakers’ understanding of English according to received descriptions of English Grammar is no longer part of the Grammar of many current native speakers and that English is also following the path down which other European languages have slid before it.

a. Frieda left the house. But she has come back.

b. Frieda has left the house. But she has come back.

Such additional constraints – to the effect that the result state must hold at the relevant evaluation time – are found with the English Present Perfect but not with other perfect forms of English. For instance, the restriction indicated by the awkwardness of (3.84.b) doesn't apply to its past tense equivalent in (3.85.a), nor in the closely similar discourse in (3.85.b).
Such additional constraints on the existence of result states of events of certain kinds – constraints to the effect that the result state doesn't obtain any more when they are no longer satisfied – are found with the English Present Perfect but not with other perfect forms of English (the Past Perfect and, to the extent that I am able to tell, the Future Perfect). For instance, the restriction indicated by the awkwardness of (3.84.b) doesn't apply to its past tense equivalent in (3.85).

(3.85) Frieda had left the house. But she had come back.

There doesn't seem to be general agreement about the explanation of this difference between present perfects and other perfects in English. Part of the story may be that while there is some kind of competition between the Present Perfect and the Simple Past, there is only one form – viz. the Past Perfect – that is available to an English speaker when she shifts either of these forms in the direction of the past. That is, the right diagnosis may well be that the Past Perfect is structurally ambiguous between a past-shifted Present Perfect and a past-shifted Simple Past. If that is so, then we would expect a restriction that applies to one of Present Perfect and Simple Past but not to the other to be no longer detectable in relation to the common form that results from past shifting either of them; for it should then be possible in principle to interpret any instance of this common form as resulting from a shift of the form that is not subject to the restriction.

Whether constraints like that responsible for the oddity of (3.84.b) also apply to the Future Perfect is not easy to decide. For instance, is (3.86) acceptable or isn't it?

(3.86) At that point Frieda will have left the house. And she will have come back inn the meantime.

I suspect that one reason why our intuitions appear to be not as sharp in relation to such Future Perfect examples as they are with regard to similar examples invoking present and past perfects has to do with factors that affect our judgments about the truth conditions of future tense sentences (those factors that have to do with the indeterminacy of the future; see Section 3.6 about the Simple Future Tense) and that may interfere with judgement that pertain directly to the perfect. Since the judgments about future perfects lack clarity, I will set this matter aside and assume that they are like past perfects rather than present perfects.

One of the challenges posed by the English Present Perfect is to articulate the properties of the result states that their interpretations invoke – those properties that must hold at a time t in order that the Present Perfect sentence counts as true when uttered at t. Attempts to meet this challenge have given rise to an extensive literature, but a definitive statement of what kinds of event descriptions give rise to what kinds of constraints continue to be elusive. (For a comparatively recent example of work devoted to this problem see (Portner 2003).) Many different constraints have been identified; they vary with the event descriptions to which the Perfect is applied, but often also with uses of Perfects of the very same event descriptions in different contexts. No effort will be made in these notes to give an overview of these various constraints. But it will nevertheless be of some use to draw a distinction between two kinds of result state that are implicated in the interpretation of English Present Perfects. – the kind that is subject to the constraints alluded to and another kind that is not. (Parsons 1990) refers to result states of the second kind, which obtain merely in virtue of an earlier occurrence of the events whose result states they are, as *resultant states* and to result states which are tied in a more substantial way to the events that give rise to them – those which are subject to the additional constraints spoken of above – as *target states*. We adopt the second term, but use the term *formal result state* instead of *resultant state*.

One important difference between formal result states and target states is that the relation between events and their formal result states is a 2-place relation that holds simply between *individual* eventualities, i.e. between an event e and the formal result state s that it initiates. We represent this relationship as 'res(s,e)'. But the relation between events and target states is a more complex one. Or, more accurately, it is a more complex relation for those who hold that a single individual event is almost always specifiable as an instance of more than one event description. When in a cinema you get up from your seat and leave the building, then in some places this is tantamount to terminating the validity of the ticket you bought to get in, so that you cannot return to the seat where you had been sitting while watching whatever they were showing. On the conception of event identity that I favor (see (Davidson 1967), (Davidson 1970)), we are dealing in this case with a single individual event, which can be described alternatively as 'you leave the theatre' and 'you terminate the validity of your ticket'. This event gives rise to different target states depending on how we classify or describe it, and these target states may have different lifespans. When you reenter the building after having left it and having thereby invalidated your ticket, one of these target states, that of having left the building, comes to an end; but the other one, that of having invalidated your ticket, does not.

Since what is to be considered the target state produced by a given event depends on the way the event is being described, the target state relation cannot be treated as a 2-place relation between individual states and individual events; it has to be treated as a 3-place relation between (i) an individual state, (ii) an individual event and (iii) an event description of which the event is an instance. We represent this second relationship as 'Res(s,e,E)', where s is an individual state, e an individual event and E an event description of which e is an instance. (In the uses we will be making of Res the third slot will be typically filled by event property terms of the form '$^\wedge e.K$', which are obtained from event descriptions to which the perfect operator is applied.)

Note that although the predication 'Res(s,e',$^\wedge e.K$)' involves the event property term '$^\wedge e.K$', and in that respect resembles the relations PROG and DISP, there is nevertheless an important difference: Res cannot hold unless there is an actual event e' that instantiates the property denoted by the term in its third argument place and to which s is directly related as its result state. In the notation we have adopted this is made explicit by the assumption that Res has an argument slot for the event instantiating the property as well as for the property itself.[35]

The predicate Res expresses a 3-place relation between target states, the events of which they are the target sates and the description of those events

[35]An alternative way in which we could account for the difference, which would bring the treatment of the perfect more closely in line with the treatments proposed earlier for progressive, habitual and dispositional sentences, would be to adopt an operator RES which forms target state descriptions out of the eventuality descriptions that are fed to the perfect operator as inputs. We could then write 's: RES($^\wedge e.K$)' to express that s is a target state of an event that instantiates the description '$^\wedge e.K$', in analogy with conditions like 's: PROG($^\wedge e.K$)', 's: HAB($^\wedge e.K$)' and 's: DISP($^\wedge e.K$)'. The difference between RES on the one hand and PROG and DISP on the other would then be that RES is *veridical*: 's: RES($^\wedge e.K$)' only holds if there is an event e' that instantiates $^\wedge e.K$ and which initiates s (that is, $E \supseteq\subseteq s$).

under which the state is a target state of the event. In the notation introduced above these three arguments follow the predicate symbol 'Res' in the manner familiar from the notation for predications in predicate logic. however, as discussed in the last footnote, Res has much in common with the predicates PROG, HAB and DIS: all are used in state descriptions, in which the state described is related in a certain way to an underlying event description. To lend prominence to this analogy we will write 's: Res(e',$^\wedge.K$)' instead of 'Res(s, e',$^\wedge.K$)', with the dref that plays the part of referential argument before the semicolon (just as we write 's: PROG($^\wedge.K$)', for instance).

We will return to the special properties of the English Present Perfect towards the end of this section. But first a couple of DRS constructions for sentences with prefect tense forms and then a lexical entry for the feature +perf that represents perfect aspect in our LFs for sentences with such forms. The two sentences for which we are going to construct DRSs are the Present Perfect sentence (3.81), repeated as (3.87.a), and its Past Perfect counterpart in (3.87.b).

(3.87)a. Frieda has closed the shop.

b. Frieda had closed the shop.

We start with the DRS construction for (3.87.a). The LF we use as input to the DRS computation for (3.87.a) closely resembles the tree in (3.81):

(3.88)

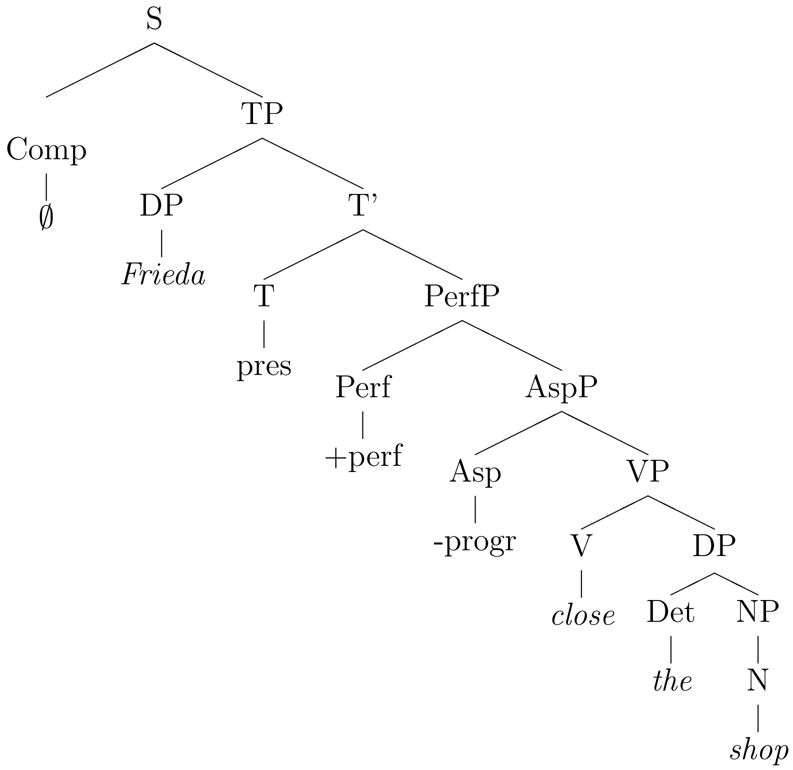

The first few steps, which lead to the semantic representation of the VP node, are as before (see Sections 3.4 and 3.5.1) and we do not repeat them. Since in the case before us nothing changes between VP and AspP (because of the feature value -prog), we start our construction at the point where the AspP node has received its representation.

(3.89)

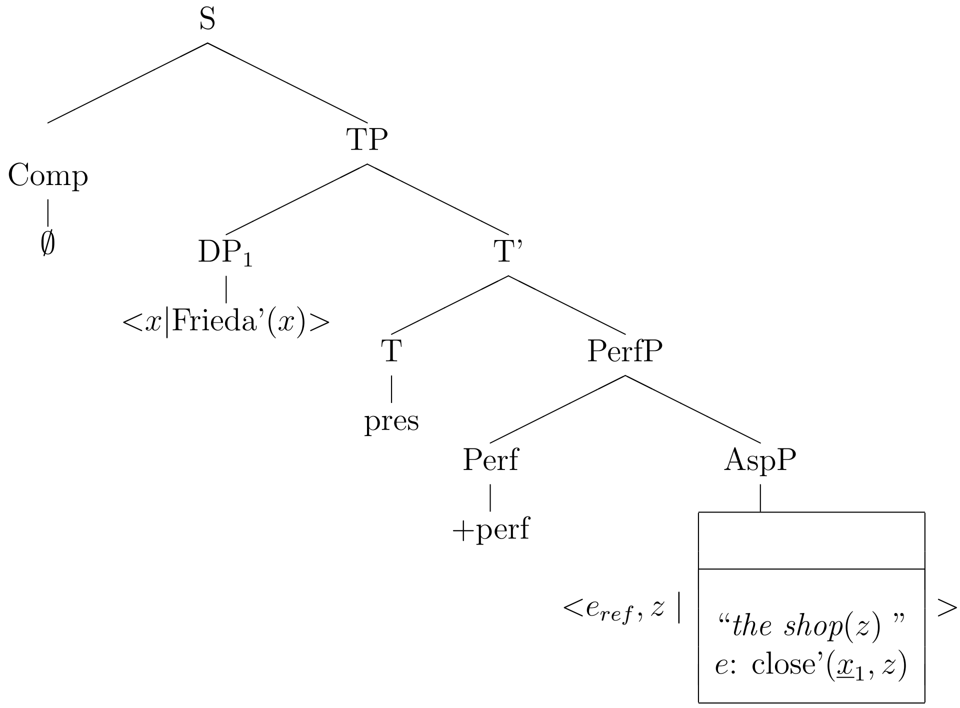

We must now specify the semantic import of the feature +perf. One of the operations involved in the execution of +perf is the introduction of a new state dref s. Furthermore, s gets inserted into the first argument slot of a DRS Condition whose predicate is the 3-place predicate Res. In addition, a new event dref e' is chosen which is inserted in the second argument position of Res; and the third argument is the property term '$^\wedge e.K'$', where K' is formed from the DRS K of the input representation in the way we describe presently. As indicated above, the notation we use to express the resulting predication is 's: RES(e',$^\wedge e.K'$)', consistently with what we have been assuming for DRS-conditions involving PROG, HAB or DISP. The store of the output representation is formed by removing the dref e from it and replacing it by s and e', which are placed at the front of the list of the input store. s takes over the role of referential argument from e and is accordingly subscripted with $_{ref}$. The DRS K' is formed from K by (i) adding to the Condition Set of K the conditions 's: Res(e',$^\wedge e.K'$)' and '$e' \supset\subset s$' and (ii) adding e to the Universe of K. (This last operation may seem odd since e is already bound by the property abstract operator $^\wedge$; but given the way the model theory of DRS languages is set up, adding e is harmless and it is convenient for the definition of accessibility, on which we do not dwell here, but which is a crucial part of the account of pronominal and other types of anaphoricity.)[36]

[36]Note that the condition '$e' \supset\subset s$' is strictly speaking redundant, since it is entailed by

When applied to (3.89) these operations lead to the structure in (3.90).

(3.90)

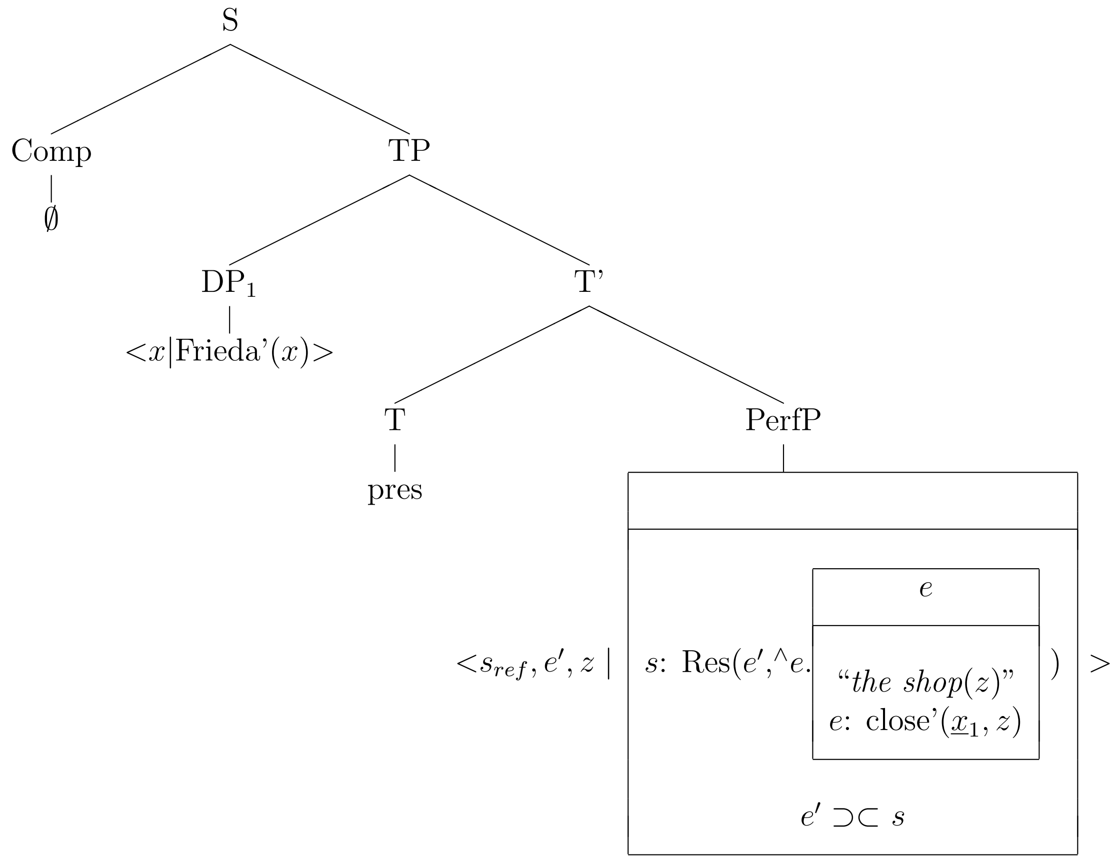

The next step is the contribution made by pres. Since we haven't yet presented an application of our entry for pres (given (3.79)), we display this step separately too, by showing the T' representation to which it leads. Note that since the PerfP representation is a state description, pres has the input that it wants.

---

the Res-Condition: when s and e' stand in the relation expressed by the Res-Condition, then they always stand in the temporal relation $\supset\subset$. We could make the redundancy explicit by stating this entailment in the form of a Meaning Postulate. Stating such an MP is left to the reader.

(3.91)

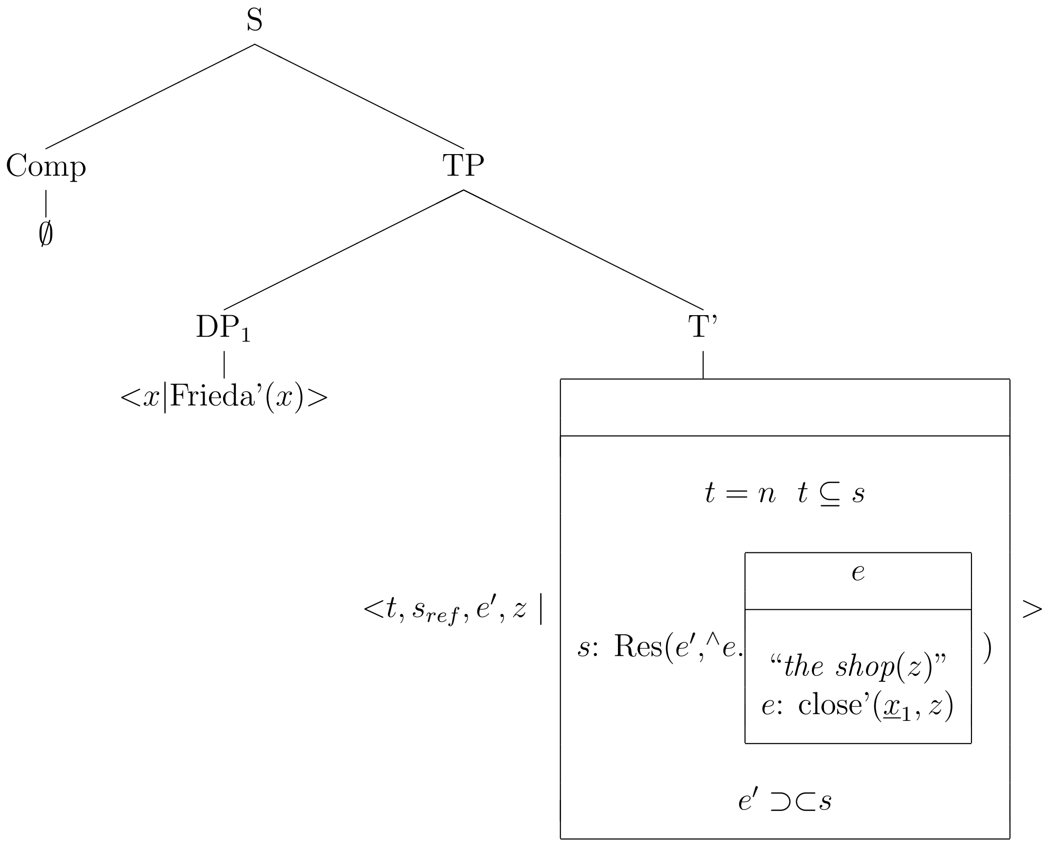

The remaining steps are familiar. They lead to the DRS in (3.92).

(3.92)

$$\begin{array}{|c|}\hline t \;\; s \;\; e' \;\; z \;\; x \\ \hline t = n \;\; t \subseteq s \\ \text{Frieda'}(x) \\ s\text{: Res}(e', {}^\wedge e. \begin{array}{|c|}\hline e \\ \hline \text{“}\textit{the shop}(z)\text{”} \\ e\text{: close'}(x, z) \\ e' \supset\subset s \\ \hline \end{array} \;) \\ \hline \end{array}$$

(N.B. An awkward feature of this DRS is that the condition "*the shop*(z)" has ended up in the DRS that identifies the description of the event e that is part of the result state description. This problem will be sorted out in Section 4, when we revise our treatment of definite noun phrases as part of our account of identification presuppositions in Sections 4.3 and 4.4.)

Now that we have presented the DRS construction for (3.87.a) in this much detail, we can deal with (3.87.b) without much additional effort. There are two differences between this example and the last one. The first difference is that the result state s introduced by +perf is now to be a formal result state, one that is related by the condition res($s.e$) to the event described by the input representation to +perf. This requires somewhat different operations than in the case of a target state. To obtain the store of the output representation, a new dref s must be introduced and added to the store of the input representation. But no $^\wedge$-abstraction over the referential argument e of the input representation is needed in this case. The DRS of the output representation is formed simply by adding the Conditions 'res(s,e)' and '$e \supset\subset s$' to the Condition Set of the input representation.[37]

The second difference is that the location time t introduced by past is declared to be in the past of n – that is, we get the condition '$t \prec n$' instead of '$t \subseteq n$'. (This now involves a choice between the two $\overset{!}{\vee}$-disjuncts of the $\overset{!}{\vee}$-disjunction contributed by past; but since the input to past is a state description, it will be the second of the two disjuncts that gets selected.) The resulting DRS is given in (3.93).

$$
(3.93)\quad
\begin{array}{|c|}
\hline
t \;\; s \;\; e' \;\; z \;\; x \\
\hline
t \prec n \quad t = s \quad \text{Frieda'}(x) \\
e'\text{: close'}(x, z) \\
\text{res}(s,e') \\
e' \supset\subset s \\
\hline
\end{array}
$$

The DRSs in (3.92) and (3.93) are the ones we want. But to get to them we had to do something that isn't quite legitimate as things stand. We simply assumed that to get the DRS for the Present Perfect sentence (3.87.a) we had to treat the feature +perf as contributing a target state (and thus a

[37]In this case too the condition '$e \supset\subset s$' is strictly speaking redundant and could be dispensed with as part of the lexical semantics of +perf so long as we adopt a Meaning Postulate that states the entailment of '$e \supset\subset s$' by '$res(s, e)$'.

condition involving the predicate Res), whereas the construction of the DRS for the Past Perfect sentence in (3.87.b) involved +perf operations that lead to the introduction of a formal result state. How do we know which of those two ways of executing +perf is the right one in each of these two cases? This question depends, we have seen, on two factors. One is whether the input to the Perf operator is an event description that comes with a target state specification. This in turn depends on two further factors, viz. (a) whether or not the verb is a target state verb and (b) what happens to target state specifications along the way from lexical insertion to the final step in the computation of the AspP representation that serves as input to the Perf operator. The second factor is the tense form of the sentence – whether it is a Present Perfect or some non-present perfect: Non-present perfects are analyzed as giving rise just to formal result states, irrespective of their input representations. This second factor involves information that isn't directly available at the Perf node, where the construction algorithm needs it. But it can be read off the tree by looking higher up for the tense feature value at T. (We could redefine the trees delivered by the parser in such a way that this information is explicitly encoded also at the lower Perf node.)

Whether the input representation to the perfect operator is a target state description will in general depend on a number of factors. One of these is the lexical verb. For instance, the choice to use Res in the execution of the perfect operation in the DRS construction for (3.82) finds its justification in the fact that *close* is a target state verb. However, in general, the mere fact that the main verb of a clause is a target state event verb is no guarantee that the input representation to the perfect operator will be a target state representation, as various modifications can occur between lexical insertion for the verb and arriving at the semantic representation of the AspP node that serves as direct input to the perfect operator. In particular, the operations triggered by +prog, +hab and +dis will change their input representations into state representations, which evidently are not target state event representations. In fact, for the sentences considered in the Notes the operations triggered by +prog, +hab and +dis are the only ones that can turn target state representations into non-target state representations.[38] So for these

[38] In richer fragments of English than are considered in these Notes there are also other modifications between lexical insertion of AspP representation that can turn target state event representations into non-target state representations. For such fragments the question when target state ergs lead to target state representations for AspP can be a quite complicated matter. I do not know how much of the various mechanisms that can be involved here and the interactions between them has been charted; but there is no need to worry about these complications here.

sentences the situation is quite straightforward: the input representation to a perfect operator is a target state representation iff (i) the verb is a target state verb and (ii) the feature value of ASP is -prog.

Among the non-target state verbs we find both event verbs (e.g. *run* or *dig*) and state verbs (such as *know*, *love* or *be happy*). There are also important differences between English perfects of these two types of verbs. The differences show with particular clarity when perfects are combined with *since* phrases, as shown by the examples in (3.94).

(3.94)a. Mary has known this since yesterday.
b. Mary knows this since yesterday.
c. Mary knew this since yesterday.
d. Since the first time they set eyes on each other he has loved her.
e. Since last week Mary has been sick.
f. Since nine o'clock this morning Mary has been running.
g. Since nine o'clock this morning Mary has run.

Since phrases show a very strong preference for perfects and for many speakers non-perfects are impossible in combination with such phrases. (For such speakers both (3.94.b) and (3.94.c) are ungrammatical – not only (3.94.b) but (3.94.c) as well.) But there is a notable distinction between combinations of *since* phrases with perfects of event descriptions and *since* phrases with perfects of state descriptions. For state descriptions the predominant interpretation is that the state described by the description that serves as input to the perfect operator holds from the time identified by the *since* phrase to the evaluation time. (For the sentences in (3.94), which are in the *Present* Perfect, the evaluation time is always the utterance time.) This use of the Perfect, which is restricted to perfects of state descriptions, is often referred to as the 'Universal Perfect'. In (3.94), examples of the Universal Perfect are (3.94.a,d,e). Sentences with *since* phrases and perfects of state descriptions also allow for 'existential' interpretations, according to which a state of the described kind has held for some time during the interval from the time identified by the *since* phrase to the evaluation time. This is shown by (3.94.e), which can be understood (a) as saying that Mary has been sick throughout the indicated period (Universal Perfect) or (b) as saying that Mary was sick during part of that time (the so-called 'Existential Perfect'). This same ambiguity is also found in (3.94.f), where the input description is that of a progressive state. (Why existential readings of sentences with *since* phrases and certain state descriptions, such as (3.94.a) or (3.94.d) are harder to get than for other such sentences is a complication to which I do not have anything to say.) But note that this ambiguity does not arise when the input description is an event description. For instance, (3.94.g) in which the input description to the perfect is an event description, only allows for an existential interpretation. On this point (3.94.g) clearly differs from its progressive counterpart (3.94.f).[39]

As an interpretational option the Universal Perfect is not restricted to the Present Perfect but can by found with the Past and the Future Perfect as well. But it is a feature of perfects that is quite specific to English and not shared with Germanic languages (e.g. German) or Romance languages (e.g. French). For a formal treatment of the kind we are concerned with in these notes the Universal Perfect constitutes a further complication that we will set aside.

A Lexical Entry for the Feature +perf

As noted, the result of applying the perfect operator to its input representation depends on two factors: (i) whether or not the input representation is a target state description and (ii) whether the tense feature is pres or some value different from pres. To repeat, when the tense feature value differs from pres (i.e. when it is past or fut) then the result state that enters into the specification of the perfect operations can be assumed to be a formal result state, irrespective of the particular properties of the input representation. When the tense feature is pres, however, then the form of the input representation matters. If the input representation is a target state event description, then the result state involved in the execution of the prefect operation must be this target state. But when the input representation does not come with a target state specification, then we will assume that the result state involved in the execution of the perfect operation is a formal result state. (For more discussion relevant to this last point see the final part of this section on the perfect, entitled 'Recency Effects of the Present Perfect in English'.)

The formulation of the lexical entry for +perf we will adopt follows this sub-

[39] Even on an existential reading (3.94.g) may not seem particularly felicitous. But there are certain contexts in which the sentence does appear acceptable, for instance when it is used to state that of the training program to which Mary is currently being subjected, in which running is one of the daily items, that part has been dealt with since non o 'clock. A universal interpretation of (3.94.g), on the other hand, does not seem possible in any context.

division into cases which we introduced in our informal discussions of the perfect: To be precise, the main division we will adopt is that between input descriptions of events that come with a target state and input descriptions of eventualities that do not. For the latter the Perf operator always introduces a new dref for a formal result state. For the former the target state becomes the result state that the operator promotes to referential argument of the output description in case the tense feature is pres. If the tense feature is different from press, then a new formal result state is introduced, just as is done for non-target state input descriptions irrespective of tense. The choice that has to be made in the case of target state input descriptions is expressed with the help of the operator $\overset{!}{\vee}$, in analogy with our earlier formulations of the lexical entries for the features past and fut.

In view of all the things we have been saying about the distinction between target states and formal result states you would expect at this point a lexical entry for the perfect that is sensitive to this distinction. But an entry that articulates this distinction will be useful only when it is possible for the construction algorithm that makes use of it to distinguish between inputs that are 'target state' descriptions (i.e. event descriptions that come with a specification of a target state for the events hey describe) and eventuality descriptions that are not target state. So the algorithm will have to mark the input representations to the perfect operator in some way for this distinction, and also must make explicit what the target state specification is for those input representations that are marked as input representations. Given the various limitations that we have already adopted it isn't as difficult to build such a module into our construction algorithm as it would be without those limitations. Within these given limits the only way in which the input representation to the perfect operator can be a target state description is when the lexical verb is a target state verb. Moreover, when the verb is a target state verb, this will lead to a target state description as input to the perfect operator if and only if no operation intervenes that turns the (target state) event description into a state description. In our present limited setting the only possible interventions of this sort are those triggered by the ASP values +prog, +hab and +dis. In other words, the input representation to the perfect operator will be a target state description iff (i) the lexical verb is a target state verb and (ii) the value of ASP is -prog.

Missing from a formal account of the distinction between perfects of target state inputs and perfects of non-target state inputs are thus only lexical entries for verbs that mark them as +/-target state and which provide explicit target state representations for those verbs that are +target state. We leave an implementation of this missing part as an exercise to the reader.

Once a mechanism for distinguishing between target state and non-target state input representations is in place, it is possible to make meaningful use of the lexical entry in (3.95), which is tuned to this distinction. In the absence of such a mechanism, however, much in this entry is otiose, and we might as well do with the simplifying entry in (3.96), which treats all input representations as if they were non-target state descriptions, but turning them into descriptions of formal result states. Both (3.95) and (3.96) are complemented by an entry for the feature value -perf which consists in passing up the input representation unchanged to the PerfP node.

(3.95) (lexical entry for the feature '+perf')

+perf (perf feature)

Sel. Restr: event description

Sem.Repr:

(i) if K is not a target state event description (i.e. if the input representation is marked -target state), then

$$<ev_{ref}, ... \mid K> \rightsquigarrow <s_{ref}, ev, ... \mid \begin{array}{|l|}\hline \\ \hline K \cup \begin{array}{|c|}\hline \\ \hline \text{res}(s,ev) \\ ev \supset\subset s \\ \hline\end{array} \\ \hline\end{array} >$$

(ii) if K is a target state description (i.e. if the input representation is marked +target state), then

$$<ev_{ref}, ... \mid K> \;\rightsquigarrow\; <s_{ref}, ev', ... \mid \begin{array}{|c|}\hline \\ \hline \begin{array}{c} K \cup \begin{array}{|c|}\hline \\ \hline \text{T-feature} = \text{pres} \\ \\ \text{Res}(s,ev',{}^\wedge ev.K) \\ ev \supset\subset s \\ \hline \end{array} \\ \overset{!}{\vee} \\ K \cup \begin{array}{|c|}\hline \\ \hline \text{T-feature} \neq \text{pres} \\ \\ \text{res}(s,ev) \\ ev \supset\subset s \\ \hline \end{array} \end{array} \\ \hline \end{array} >$$

Missing from part (ii) of this definition is the information that the result state s satisfies the target state description specified by the input representation. So long as no convention has been adopted for the form in which target state descriptions are specified as part of the input representations to +perf, this missing bit cannot be stated in formal terms. Once a convention has been adopted, however, adding the missing information to (ii) should be unproblematic.

(3.96) (simplified lexical entry for the feature '+perf')

+perf (perf feature)

Sel. Restr: event description

Sem.Repr: $<ev_{ref}, ... \mid K> \;\rightsquigarrow\; <s_{ref}, ev, ... \mid \begin{array}{|c|}\hline \\ \hline K \cup \begin{array}{|c|}\hline \\ \hline \text{res}(s,ev) \\ ev \supset\subset s \\ \hline \end{array} \\ \hline \end{array} >$

Exercise: Use the simplified entry for the perfect operator in (3.96) to construct DRSs for the following two sentences:

(i) Mary has been sick.
(ii) Mary will have run.

Recency Effects of the English Present Perfect

We conclude our discussion of the perfect with some further informal remarks about the special properties of Present Perfect in English.

The English Present Perfect has often been described as carrying an implication of 'current relevance'. With target state verbs current relevance typically consists in the current holding of the target state that comes with the described event and that (in the syntax-semantics interface architecture we are assuming) is given as part of the event's description. But perfects of non-target-state input descriptions to +perf also show current relevance effects. A classical example, about which there has been a good deal of discussion in the literature, is the sentence

(3.97) Nixon has died.

The relevant use of this sentence on which the literature has focused is an instance of what is often referred to as the 'hot news perfect'. This is a use that is restricted to situations in which the reported event has just happened and may still be expected to be new information to some or all of the audience. The result state effect of such utterances has to do with the information changes that they are able to bring about in the recipients for whom they are intended – in the case at hand, recipients for whom the subject Nixon is a familiar public figure but who as yet are unaware of his death.

An interesting feature of this example is that the verb *die* would seem to be a target state verb par excellence: the death of a person results in their no longer being alive and that is a state which will last forever. (At any rate, that is the understanding of pretty much all people I ever talk to, myself included.) But it seems to be precisely the fact that the result of dying lasts forever that makes this result state unsuitable as justification for Present Perfect uses of *die*. This is why the felicitous use of sentences like (3.96) is limited. For instance, suppose that a first year graduate student in philosophy, who has read some of the works of David Lewis, proposes to invite him for a talk at next year's colloquium. The student isn't aware that Lewis is no longer alive. (Lewis died, at what is nowadays the comparatively young age of 59, in 2001, nearly fifteen years ago.) To convey the inadequacy of his proposal by using the words 'Lewis has died' would be odd in this situation. Better would be something like "You know, don't you, that Lewis is no longer alive. He died in 2001." Had the proposal been made only a month after Lewis' death, the words might have been considered adequate.) In general, the most direct, invariable and obvious result of someone dying is that he or she is dead. But that won't do as justification for using the Present Perfect of *die* precisely because there is no end to such result states.

The exchange in (3.98) is another example that illustrates the role of current relevance.

(3.98) A: Why are you out of breath?

B: (i) I have been running.
(ii) I was running.
(iii) I have run.
(iv) I ran.

Of all the possible responses for B listed as (i)-(iv) (i) sounds the most natural to me. The justification of the Present Perfect in this reply is that one of the results of B's running is that he is still out of breath. That is a result of running which lasts only for a little while after you stop running. And it is that state that A has just drawn attention to by her question. So this is a context in which the relevant result state is already salient at the moment when B gives her answer and that facilitates its role as justification for B's use of the Present Perfect.

But this is only the beginning of a story that could account for why some of the replies (i)-(iv) are better than others and what it is that makes any of them bad. I won't try to tell such a story, and leave it to the readers to come up with their own assessment which of (i) - (iv) are good replies to A's question, and how they differ from the ones that do not seem right, and why. The matter is fairly complex. One source of the complexity has to do with a tendency of the Simple Past to be interpreted as presupposing that some particular time at which the described eventuality obtained or took place is given or can be recovered from context. For instance, (iv) suggests that B is referring to some particular event, perhaps that of getting to the location where the exchange between A and him is taking place. This 'anaphoric' dimension of the Simple Past is something that will be discussed at some length in Section 4.

Considerable efforts have been made over what at this point has been more than half a century to identify the different possible uses of the English Present Perfect and the reasons why some of its uses are fine while others are not. But a consensus on a conclusive set of answers still doesn't seem to be in sight.

Here is a summary of these last remarks on admissible uses of the Present Perfect:

(i) The use of the Present Perfect carries an implication of current relevance of the described eventuality – the eventuality described by the input representation to +perf must have produced an effect that still holds at the utterance time.

(ii) When the input description to a Present Perfect comes with a specification of a target state (and this target state is conceived as typically limited in duration) then, as a default, current relevance is satisfied iff the specified target state holds at the utterance time.

(iii) When the input description to a Present Perfect does not specify a suitable target state, then it must be possible to construct an 'ad hoc' justification from the content of the sentence in combination with the context in which it is being used. The result state that is accommodated as part of this justification must be a state that can be seen as a causal effect of the described eventuality and at the same time as relevant in the context in which the prefect sentence is being used.

An implementation of the recency requirement would involve a lexical entry for +perf which specifies that when the tense value is pres, then the output representation always involves the predicate Res. If the input representation is a target representation, then the referential argument s of Res should be a state of the type specified by the input representation, as discussed in connection with the entry (3.95). When the input representation is a non-target state eventuality description, then a coercion mechanism is activated that requires finding a result state that stands in a suitable causal relation to the second argument of Res and that satisfies appropriate relevance criteria (to which we have alluded, but about which we have not said anything of substance here).

3.9 Temporal Reference and Logical Operators

3.9.1 *every* and other True Quantifiers

In PART I of these Notes much of our attention was focused on quantification (together with the other two fundamental concepts of classical semantics, reference and predication). In this we have followed a tradition in Formal Semantics that goes back to Montague and that has its roots in the concerns of the great logicians of the late nineteenth and the first half of the twentieth century, who had come to see that quantification is what gives formal logic, in its canonical form of the Predicate Calculus, its astounding expressive power. Since Formal Semantics of natural language originated (with Montague's seminal work) as an application of the model-theoretic methods of mathematical logic to fragments of English, it is not surprising that an emphasis was put on fragments that had, at a minimum, the expressive power of first order predicate logic. An important aim of Montague's work was to establish that parts of English and other natural languages have this power, and that they have it by virtue of a syntax and semantics that are just as systematic as the syntax and semantics of the Predicate Calculus or the λ-Calculus. In fact, the natural language fragments these early studies of formal semantics focussed on do not only have the resources to express all that can be expressed in first order predicate logic; they can express complex first order propositions in what appear to be very similar ways, with quantifying DPs like 'every N' and 'some N' playing roles that closely resemble those played by the logical operators '$\forall$' and '$\exists$'.

Nevertheless, as formal semantics of natural language developed in the course of the past half century, we have become increasingly aware of the many respects in which natural languages differ from the predicate calculus and other systems of symbolic logic, both in their syntax and in their semantics. On the one hand, English – and much the same would seem to be true for most or all other human languages – has the capacity of expressing all of first order logic, and the only quantificational devices that it needs for this are DPs headed by the determiners *every* and *some/a*. But when we look more closely, we see that these DPs interact with other parts of the sentences in which they occur in ways that differ subtly but importantly from the structural interactions that can be observed within the predicate logic formulas to which the sentences are truth-conditionally equivalent. We touched on this in Section 3.4, where we drew the distinction between referential and non-referential arguments. Nouns and other predicate words, we concluded, bring one of their arguments – their referential argument – along and only when they have additional argument positions, as with relational nouns, those positions will be filled by separate 'argument terms', in a way reminiscent of how *all* argument positions are filled of the predicates of first order logic or the λ-Calculus. That distinction, we surmised, amounts to a radical difference between the organization of predicational information in natural languages as compared with the languages of symbolic logic.

For the most part the implications of this difference were left as topics for further exploration. But one implication was central to what we wanted to say: verbs too have referential and non-referential arguments, but for a verb the referential argument is always the eventuality it describes. That means on the one hand that verbs too instantiate the organization of predicational information that sets natural languages apart from the Predicate Calculus and its ilk, and on the other that semantic representations are permeated with eventualities. More precisely, all sentence representations now contain eventuality variables – in our terms: eventuality drefs, but for the present discussion the distinction between variables and discourse referents is immaterial – which are introduced by verbs but then have to be bound eventually, so that the representation will determine a well-defined proposition.

But how can eventuality drefs be bound? So far we have encountered only two kinds of possibilities: (i) by the property abstraction operator ^, as part of the interpretation of operators like PROG or HAB; (ii) as part of the final step in the construction of the semantic representations of main clauses, when all remaining drefs in stores are transferred to the Universes of the DRSs to their right. But there are many more possibilities than just these two. In particular new possibilities arise in sentences in which eventuality descriptions are interpreted as within the scope of quantifiers. These quantifiers can take various forms. Here we limit our attention to the nominal quantifiers that took up so much of our attention earlier on, viz. DPs with the determiners *every* and *some* or *a*.

The importance of interactions between quantifiers and eventuality drefs can hardly be overestimated. For one thing it plays a part in pretty much every sentence that has quantifiers at all. And yet it is one of the most systematically neglected topics in formal semantics. (This might not really be surprising if one keeps in mind that most approaches to the semantics and logic of natural language do not vouchsafe eventuality variables in the first place.) What we are facing here is thus a domain of investigation that is well-neigh virgin territory and yet one of the core phenomena of natural language semantics.[40]

In this and the next sections we will only touch on some of the basic aspects of this cluster of problems. But I hope that this will make the reader conscious of the pervasiveness of these interactions and of the wide range of different forms they may take.

Before we get down to the real work of this section, let me say something about the extensive work on quantification in natural language that has played such a prominent part in semantic research over the past fifty years and in which temporal reference plays virtually no role. Without a few words on this work the considerations that follow below might leave the impression that I do not think this work important, or even that I might think it misguided. Nothing could be further from the truth. As I see it, work on natural language quantification has been one of the most fruitful, productive and sophisticated areas of formal natural language semantics in its 50 year history. Nowhere else in natural language semantics have a combination of careful syntactic and semantic analysis with methods from mathematical logic led to a comparable multiplicity of linguistic insights and non-trivial formal results. (For a survey of such results, which stands out through its comprehensiveness, clarity and systematicity, and also because it contains a range of results and insights due to the authors themselves, see (Peters & Westerstahl 2006).) It should also be stressed that there is a simple and perfectly legitimate way to finesse matters of tense and aspect when studying quantification in natural language, which preserves all the important insights of the work alluded to: Restrict attention to sentences that are in the simple present tense and which only contain verbs that are state verbs. Such sentences can be treated, without any distortion of their quantificational structure, in the way they have been in almost all of the formal semantics literature: as claims about what is the case at the utterance time t. For the truth or falsity of this claim, all that matters about the predicates the sentence contains is what they are true of at t. In this setting it becomes unimportant whether state verbs are treated as predicates with an argument position for states or whether their only arguments are what we have been calling their non-referential arguments. For let V be a state verb belonging to the given sentence, let $\mathrm{P}_{V,s}$ be the predicate that treats V as a predicate with

[40]Not any more since the first draft of this section was written. A excellent proposal for a ‘minimal’ formal environment in the tradition of Montague Grammar that is suited for studying and describing interactions between eventualities and other elements in compositional sentence semantics has been provided by Champollion in his (Champollion 2014).

a state argument and P_V the predicate that treats V as a predicate without a state argument. Then V's contribution to the truth conditions of the claim made by its sentence are captured just as well by P_V as by $P_{V,s}$. Suppose for instance that V is a transitive verb. Then what matters for the semantic contribution it makes is what pairs $<a, b>$ P_V is true of at t. Another way of putting this is that all that matters is for which pairs $<a, b>$ there is a state s that includes t and is such that $P_{V,s}$ is true of $<s, a, b>$. If we make the assumption (to which we have already committed ourselves in any case at an earlier point) that whenever P_V is true of $<a, b>$ at t, then there is a state including t such that $P_{V,s}$ is true of $<s, a, b>$, then these two ways of analyzing V are equivalent: for such sentences there is no need to analyze V as a predicate of states. An analysis as P_V will do just as well.

Confining investigations into natural language quantification to sentences of the sort described in the last paragraph is not only legitimate, it is arguably the best way to bring into focus a cluster of properties of natural language quantifiers that would only have been obscured if their interactions with eventuality variables had been brought into the picture from the start. But the fact remains that natural language quantifiers do significantly interact with temporal reference most of the time. There may be some instances, such as the present tense sentences discussed above, where the eventualities contributed by verbs can be ignored without loss. But in speech about mundane mattersthese are the exception. In general there is no hope of getting the truth conditions of sentences with quantifiers right unless the interactions between those quantifiers and tenses, aspect operators and other devices of temporal reference are brought into the analysis as well.

Now let's get to work. We start with as simple a quantified sentence as you can find, viz. (3.99).

(3.99) Every philosopher slept.

Before turning to the actual DRS construction for (3.99) let us first have a semi-formal look at the role that tense plays in the interpretation of this sentence. Since the verb, *sleep* in this case, brings along its eventuality argument, this argument has to be bound somewhere within the resulting semantic representation. But where? Note that even with a sentence as simple as (3.99), in which there is only one overt quantifying phrase, we are facing questions of scope: Is the event dref e that is introduced by *sleep* in (3.99) to be bound within the scope of the quantifier *every* or outside of it? And, secondly, what are we to say about the scope relations between the quantifier *every philosopher* and the location time introduced by the past tense? Let us assume, as we have been doing so far, that e must be in the scope of the tense feature past, which introduces the time t that locates e somewhere in the past of n. This impulses some constraints on the possible scope del;at ions. But even with this constraint in place we are still faced with three scope possibilities:

(3.100)(i) Q $> t > e$.
(ii) $t >$ Q $> e$.
(iii) $t > e >$ Q.
(Here Q stands for the quantifier expressed by *every philosopher*.)

To see what these possibilities come to, let us spell out what each of them implies for the truth conditions of (3.99).

According to option (i), (3.99) says that for each philosopher x in the domain of quantification there is a time t_x in the past of n and an event e_x of x sleeping that is temporally included in t_x.

Option (ii) says that there was a time t in the past of n such that for each philosopher x there was an event e_x of x sleeping that was included in t.

Option (iii) has it that there was a t before n and an e included in t such that for every philosopher x e was an event of x sleeping.

What can we say about the truth conditions that are imposed by the three scope options? Which of them fits our intuitions about what 3.99) means most closely? That doesn't seem quite clear. And the reason for that would appear to be that in last analysis there isn't any substantive difference between the three conditions. Suppose for instance that (i) holds. Then the different times t_x at which the different philosophers x slept (i.e. which contained the respective sleeping events e_x) could be argued to be all included in a single past time t. This t will then have the property that for each x there is an event e_x of x sleeping that is included within it; which means that the truth conditions of option (ii) are satisfied as well. By much the same reasoning we can also argue that given (i) the different events e_x can be amalgamated into a single event e which has the property that for each philosopher x e is an event of x sleeping (in that it has a sub-event which is 'just' the sleeping of this x). So the truth conditions of option (iii) are satisfied as well. Reasoning in the opposite direction is possible too, leading from the truth conditions of option (iii) to those of option (ii) and to those of option (i).

Some of the assumptions that are used in these arguments arguably need further justification. (For example, what is the ontological principle from which we can infer that the different events e_x can be assembled into a single event e, or that an event e of all members of a certain set X sleeping has for each $x \in X$ a sub-event of x being asleep?) But these principles would seem hard to refute definitively, and so long as that hasn't been done it is impossible to firmly ascertain that there is a difference between the truth conditions determined by (i), (ii) and (iii). The matter is complicated further by an aspect to the interpretation of the Simple Past tense on which we haven't touched yet and that we will be able to address properly only in Section 4. When (3.99) is used as part of an ongoing conversation or text, then more often than not its past tense will be interpreted 'anaphorically', viz. as referring to some past time that has been introduced by some earlier sentence. In such interpretations the anaphoric link between the past tense of (3.99) and the earlier established time will overwrite, as it were, the scope relations in which the past tense stands to other scope-bearing elements of its sentence. For instance, even if we take it that the syntactic structure for (3.99) imposes the scope order given in (3.100.i), the anaphoric interpretation of its past tense will lead to truth conditions that are formally like those determined by option (ii).

Nevertheless, even if we cannot say anything at this point about anaphoric interpretations of (3.99), we are in a position to deal with its non-anaphoric interpretations, and a good deal can be learned from just doing that. Here we will look at only one of the possible interpretations of (3.99), the one which is suggested by the hierarchical structure of the LF that we will adopt as input to the construction of its semantic representation.

First, we have to decide what is the right LF for (3.99). In PART I we argued that Top Down DRT shares with other approaches to formal semantics the problem that for certain constructions involving quantifiers – those that involve inversely linked quantifiers, such as *one apple in every basket* – only an LF with raised quantifier phrases will lead to the intuitively correct semantics. This is true for Bottom Up DRT just as much as it is for Top Down DRT, so for these cases our present approach will require LF inputs with raised quantification phrases as well. In PART I it was proposed that since we cannot do without quantifier raising for some cases, we might as well use LFs with raised quantifying DPs in all cases. This is also the policy we now adopt for Bottom Up DRT.

This almost settles what LF we should assume for (3.99); but there is one further point that needs addressing. In PART I we assumed that raised DPs were adjoined to S. But given the assumptions we have made in the course of Section 3 about the role of Comp – the main clause complementizers, which are all that we have so far had occasion to deal with, trigger transfer of the drefs that remain in store to the Universe of the DRS following it, whereupon the store is dissolved – adjunction to S can no longer be what we want: DP adjunction has to be below Comp. But it should be *just* below Comp. Since subject DPs are 'specifiers of T'' (i.e. sisters to T') and since we assume that they are raised into adjunction positions, the only possible adjunction level that remains is TP. So that is where we will assume raised quantifying DPs end up, both when they are grammatical subjects and when they occupy some non-subject position in their clause.

For (3.99) this means that its LF should be as in (3.101) (where the AspP and PerfP levels have been suppressed, since in the present case no significant changes take place at these levels; the relevant feature values are -prog and -perf). Likewise we ignore the NumP projection in the DP *every philosopher*.

(3.101)

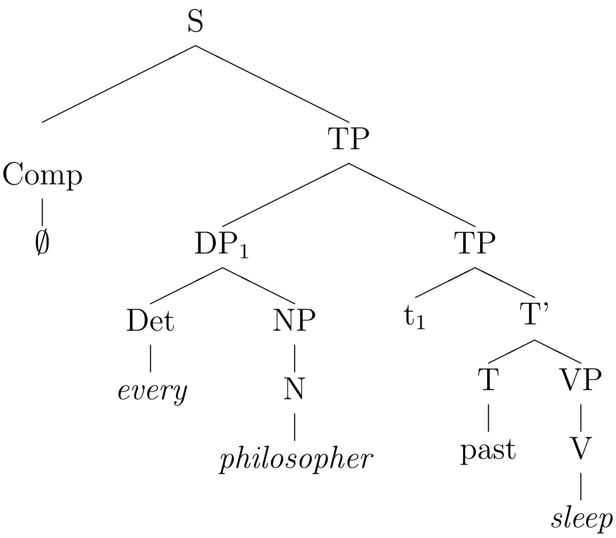

The computation of the semantic representation of T' in (3.101) we have seen before, so that is the stage where we start our discussion of the DRS construction for this LF:

(3.102)

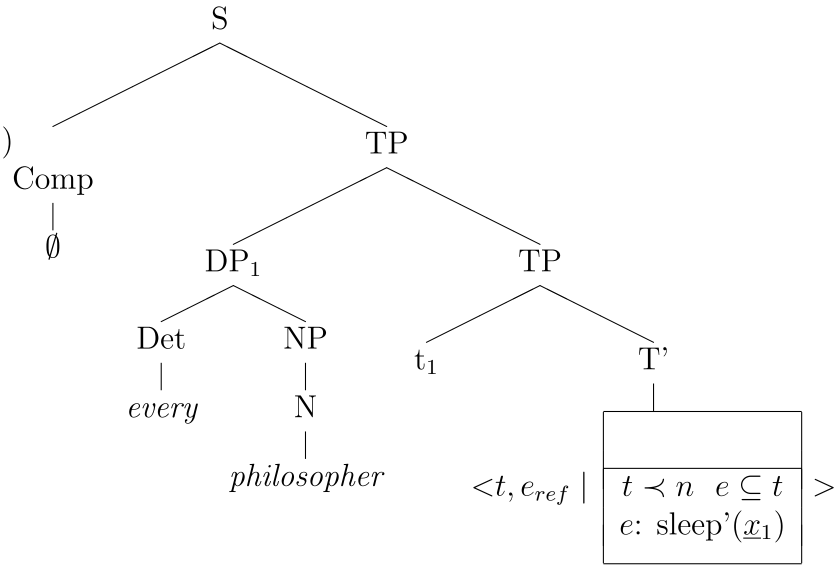

The first of the construction steps that are still to be carried out is the one dealing with the trace t_1. Note that the index of the trace matches two others, that on the raised DP of which it is the trace and that of the argument slot $\underline{x}$ of the predicate sleep' that the parser has identified as the slot of the DP. This 'double coindexation' motivates the following informal description of the step: introduce a new dref, co-index it with the DP and insert it into the coindexed slot. Our technical implementation of this informal description goes like this: (i) choose a new dref to interpret the trace (here we choose x); (ii) place this dref in the store, subscripted with the index from the trace; (iii) insert the dref in the coindexed argument slot. This leads to the lower TP representation in (3.103).

(3.103)

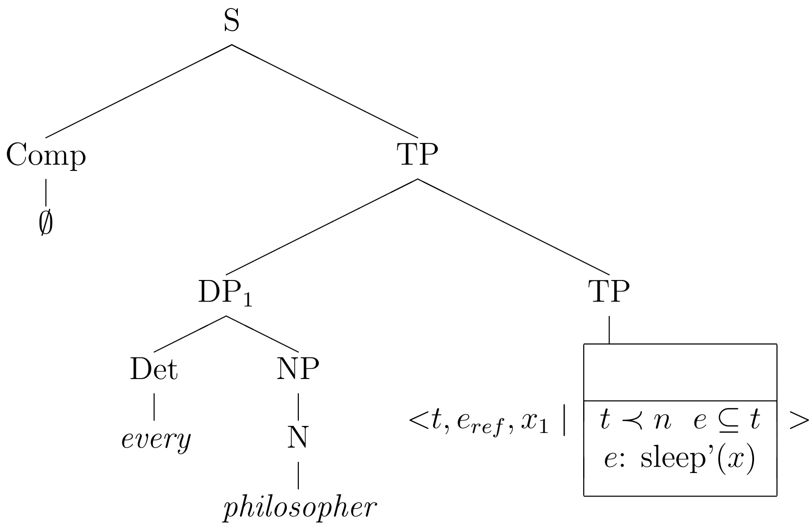

Next we turn to the semantic contribution of the raised DP. Here we face the same problem that we encountered when dealing with quantification in the Top Down setting of PART I. In order to verify that the DP is quantifying, and what kind of quantifier it is, we need to look down into the internal constituency of the DP in order to identify its Determiner. Neither information is available directly and explicitly at the DP, but it is clear that the parser that outputs the LF (3.101) for the input (3.99) could be redesigned so that this information would be directly available at the DP node itself (in the form of some kind of feature annotation). So, to repeat the point made in PART I, the need to inspect the syntactic structure of the DP at this point doesn't constitute a violation of compositionality that is a ground for concern; we will continue our policy of looking down into LF structure if and when we need to.

To carry out the different operations that are dictated by the quantifying determiner *every* which the construction algorithm identifies as the quantifier of the semantic representation that it must put in place, the algorithm needs as input representations (i) a semantic representation for the sister node to the Determiner node and (ii) a representation for the sister node to the DP. (3.103) gives the latter representation but not the former. The former representation – that for the NP *philosopher* is constructed according to principles that are well familiar at this point.[41] (3.104) shows the result of inserting this representation into (3.103).

[41] In this DRS construction we ignore the temporal dependence of the noun *philosopher*, so as not to have to deal with two many problems at once. But see the second exercise at the end of this subsection.

(3.104)

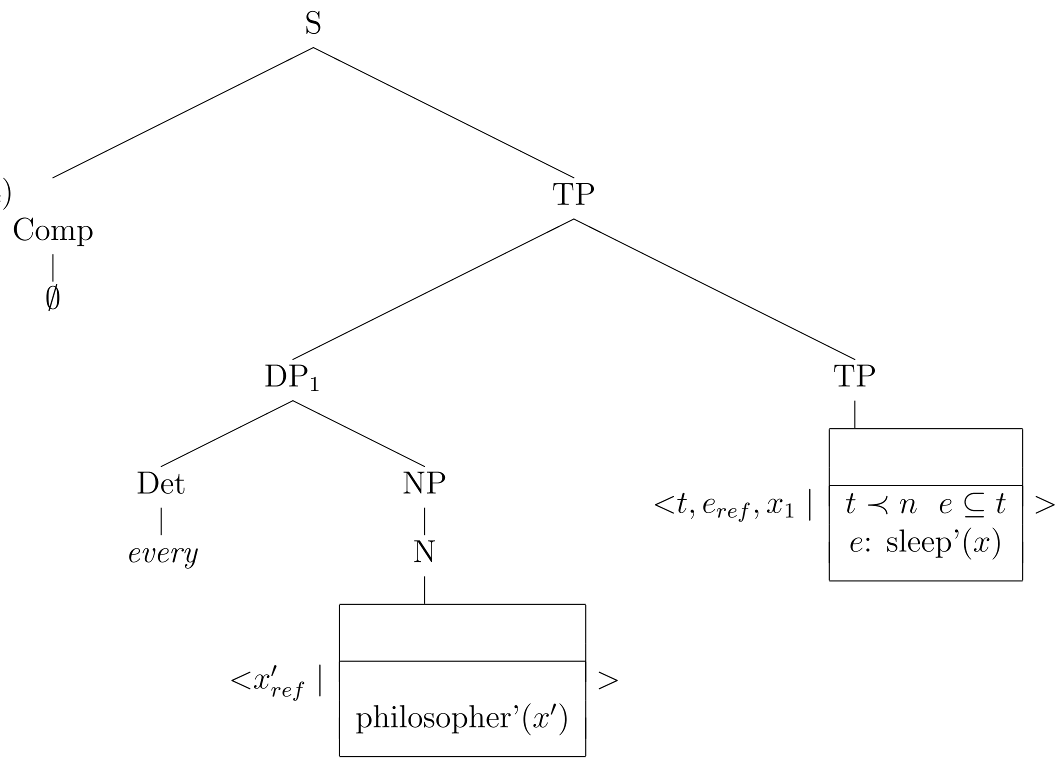

every now combines with the representation of its sister NP and with the representation of the DP's sister TP in a manner that is familiar in it essentials: *every* sets up a duplex condition for which the NP representation provides the restrictor and the TP representation the nuclear scope. The details are also pretty much as expected, except for one thing. Both the NP- and the TP-representation each have their own dref store. What are we to do with these stores? Put all the drefs they contain into a single store that we put in front of the output representation? Or keep the stores separate, attaching each to the (sub-)DRS that receives the other material from the input representation to which that store belongs? There are several reasons for adopting the second option, one of which will become visible presently. Before presenting the lexical entry for *every* and spelling out the details of how it is applied, we show the result of applying this entry in the case of (3.104).

(3.105)

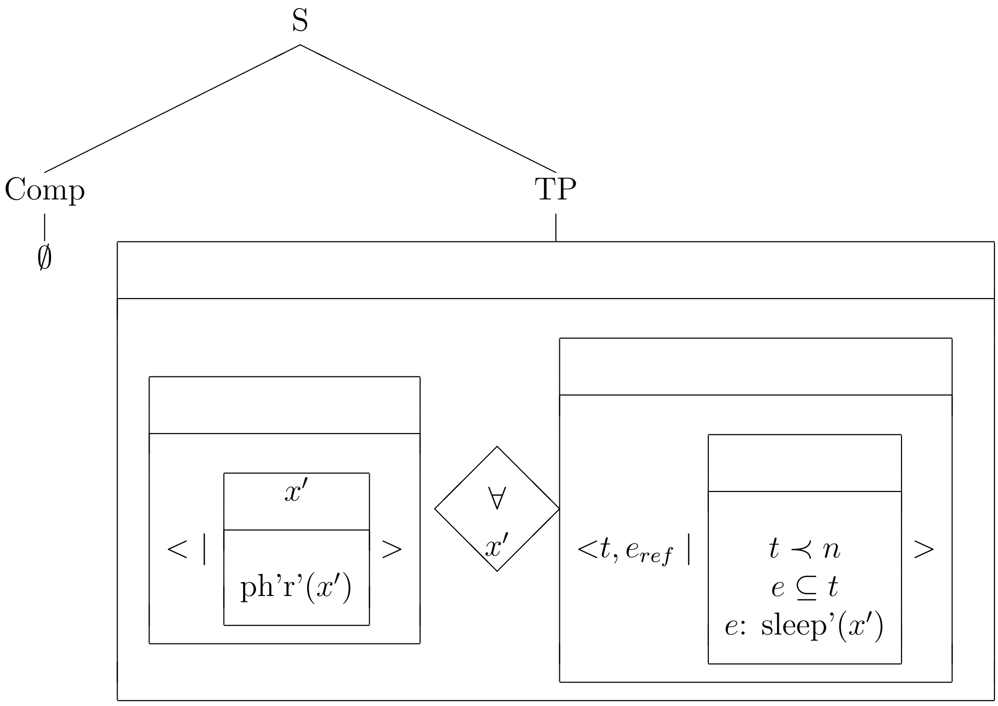

Some comments are in order about (3.105). The first one concerns the handling of the stores of the input representations to the construction rule for *every* of which (3.105) is the result. One reason for keeping the stores of restrictor and nuclear scope of the quantification local is that it should be possible to bind the drefs they contain locally. We will assume that when a dref gets transferred from a store to a DRS Universe its target must either be the Universe of the DRS K following the store or a higher Universe, belonging to a DRS which contains K somewhere as a sub-DRS. However, moving the dref to a higher Universe will be permitted only under special conditions (having to do with presupposition resolution); and since we will be in a position to state these conditions only in Section 4, we provisionally adopt the simplifying assumption that dref transfer to higher DRS Universes is prohibited.

On the other hand it will be a general principle throughout that drefs may never be 'lowered' – that is, moved from the store or Universe of a given DRS to the store or Universe of one of its sub-DRSs. For the case at hand this means that once we move the referential argument e of sleep' into the store of the DRS containing the duplex condition, there won't be any way of moving it back down into the Universe of the nuclear scope DRS. So if that is an option for where e may end up, it should be kept, for the time being, in a local store to that DRS, as shown in (3.105). The same consideration applies to the time dref t that was introduced by the past tense.

The second comment pertains to the referential argument of the representation of the sister node to the determiner. (In (3.105) this is the dref x', the referential argument of the NP representation. x' gets bound by the quantifier of the duplex condition and that means that it must be taken out of the NP store and transferred to the central diamond. But x' is also involved in some further operations. First, we stick to a practice that goes back to the point in time when duplex conditions were first introduced into DRT – that of adding the dref bound by the quantifier of a duplex condition also to the Universe of the restrictor DRS. (As argued in PART I, adding the dref to this Universe has no effect on the truth conditions of the duplex condition, and it is convenient in the definition and visualization of dref accessibility.)

In addition the dref x' must be 'unified' with the dref x that is coindexed with the DP containing the quantifying determiner. We will assume that this 'unification' takes the form of replacing x by x' everywhere in the representation K' of the sister to the DP node. (This presupposes that K' doesn't contain any occurrences of the replacing dref x' already, so that the substitution doesn't lead to any 'variable clashes'; but that won't happen so long as we stick to the practice of never introducing the same dref twice in the course of constructing a semantic representation, whether for a single sentence or for a discourse or text.)

All in all a fair number of operations have to be performed in getting from a representation like that in (3.104) to one like that in (3.105). Here is a summary:

(i) Removal of the dref α coindexed with the quantifying DP from the store of the representation $<ev_{ref}, ..\alpha, .. \mid K'>$ of the sister to the quantifying DP and replacement of this dref everywhere in K' by the referential argument α' of the DP.

(ii) Removal of the referential argument α' of the DP from the store of the representation $< \alpha'_{ref} .. \mid K>$ of the sister to the Determiner node and inserting it into the Universe of K.

(iii) Formation of a new DRS with an empty Universe and with only a single DRS condition. This condition is a duplex condition with (a) the modified representation $< .. \mid K>$ as restrictor DRS, (b) the quantifier '$\forall\alpha$'' as quantifier and (c) the modified representation $< ev_{ref}, ... \mid K'>$ of the sister to the DP node as nuclear scope.

A schematic representation of the result of the listed operations that are part of the construction rule triggered by *every* is shown in (3.106).

(3.106) Lexical entry for the determiner *every* (preliminary)

every (Determiner)

Sel. Restr:

Sem.Repr:

$<x_{ref}, .. \mid K> \;\rightsquigarrow$

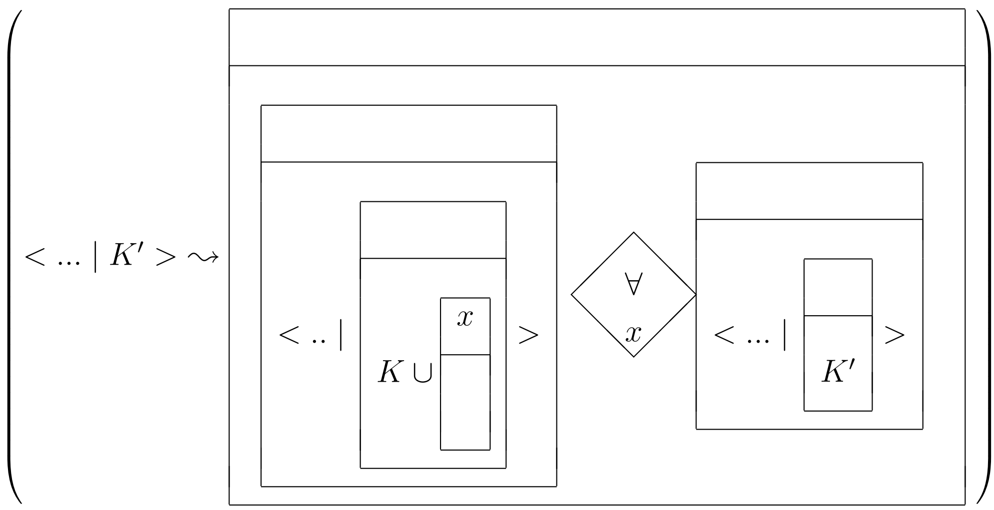

For a first revision of the construction rule in (3.106) see (3.110) below. A further revision will be made at the end of Section 3.10.1.

One more step is needed to convert (3.105) into a proper DRS. This is the step triggered by the Comp node, in which the remaining store elements are transferred to DRS Universes. We have seen many instances of this operation. But this is the first time that we are dealing with representations which have more than one store. In (3.105) there happens to be only one non-empty store (that of the nuclear scope representation), but that is enough to illustrate this new general issue that arises for representations with multiple stores: What is to be done with stores that precede sub-DRSs? The issue is settled in a simple way by the provisional principle we have just adopted, according to which a dref that is transferred from its store may only be transferred to the Universe of the DRS that immediately follows it. (But this principle is only provisional, and when in Section 4 we will be forced to give it up, we will find that the options that abandoning it will create are not easily controlled.) Application of our provincial processing principle to the two drefs in the nuclear scope store moves both of these to the Universe of the nuclear scope DRS.

Now that both subordinate stores are empty, they can be eliminated from the display of the representation and the DRSs that follow them can be merged with the DRSs containing them. The result of this is shown in (3.107).

(3.107)
$$\begin{array}{|c|}\hline \\ \hline \begin{array}{|c|}\hline x' \\ \hline \text{philosopher'}(x') \\ \hline \end{array} \quad \begin{array}{c} \forall \\ x' \end{array} \quad \begin{array}{|c|}\hline t \quad e \\ \hline t \prec n \\ e \subseteq t \\ e\text{: sleep'}(x') \\ \hline \end{array} \\ \hline \end{array}$$

The entry for *every* in (3.106) and its applications also requires some further comments. First, the entry differs from all previous entries presented in PART II in that it needs two inputs and not just one. Second, no selection restrictions are needed in this case, since the construction algorithm can be relied upon to impose the necessary constraints.[42] Thirdly, note that the unification involved in the creation of duplex conditions is fully determined by the LF: the only possible dref from the TP representation that qualifies

[42]There can of course be selection conflicts between the nominal head of the DP (the noun that is the head of its NP) and the argument position that the DP fills according to the syntactic analysis of the sentence. But this restriction is operative just as much for non-quantifying DPs. How the sortal properties of the DP, as specified by its nominal head, are checked for compatibility with the selection restrictions associated with the DP's argument position is a question we postpone till Section 4.

as unification partner for the referential argument of the quantifying DP is the dref introduced by the trace for that DP.

Lastly, there are the truth conditions determined by the representations that result through applications of (3.106), such as the representation in (3.107). What has been said about the truth-conditional contributions of duplex conditions in PART I remains in force. But a novel aspect to the duplex condition of (3.107) is that its nuclear scope contains the discourse referents for the event described by the verb and the location time introduced by tense. That these two drefs end up here and not in the store or DRS Universe of some other part of the representation is due to the LF we adopted as starting point for the construction of (3.107), in which the DP has scope over the lower TP. The effect of this on the contribution of the duplex condition to the truth conditions of (3.107) is that for every philosopher belonging to the domain of quantification there was some time in the past of n at which that philosopher slept. So, as far as these truth conditions are concerned, (3.105) would be true so long as each of the philosophers was asleep at some time in the past, even if those times were very far apart from each other. That is surely counterintuitive – the natural way in which to understand (3.107) is that there was some particular time at which all the philosophers from the given set were asleep. One way to obtain a representation which captures this intuition would be to start from an LF in which the constituent that contributes the location time has scope over the quantifying DP. But we have already given an argument against such an LF. Rather, as indicated above, the correct account of the intuition that (3.105) requires a single past time at which every philosopher slept is one that relies on the anaphoric dimension of non-present tenses – their capacity for resuming times that have been previously mentioned or that are salient in the context for some other reason. As noted, this anaphoric dimension of tenses and other devices of temporal reference cannot be dealt with in proper detail until presupposition management has been introduced into our construction algorithm, and thus will have to wait till Section 4. But a flavor of how the mechanisms that will be introduced there can lead to representations that capture our intuitive understanding of sentences like (3.105) is to consider sentences like (3.113), in which the tense is 'reinforced' by an adverb that denotes some particular time.

(3.108)On New Year's Eve 2001 every philosopher slept.

The DRS for this sentence will for each philosopher x locate the event e_x of x sleeping not only in the past of n but more specifically at New Year's Eve

2001, just as our intuitive understanding of this sentence demands.

Although the resources for a proper treatment of the anaphoric aspects of temporal reference are not yet in place, it is nevertheless appropriate at this point to discuss in a preliminary way a possibility that can be adequately implemented only via the presuppositional dimension of temporal reference devices. The representation-transforming operators that were introduced in earlier sections were all assumed to create eventuality descriptions as output representations. For instance, the output representations of the operator HAB are state descriptions – descriptions of states which hold over some period during which eventualities of the kind specified by the input descriptions occur 'habitually'. In view of the many similarities between habituals and quantification the question imposes itself whether a similar treatment – one which produces state eventuality descriptions as output representations – would not be appropriate for quantificational sentences as well. On this view the semantic contribution of a quantifying DP consists not just in the introduction of a duplex condition, but also involves the introduction of an eventuality dref, together with a condition which says that the type of the eventuality represented by this dref is given by the duplex condition. It is then this eventuality that should be located by tense and temporal adverbs, much as we have assumed about the states introduced by progressive and habitual aspect and by the perfect.

But what *kind* of eventualities should the eventualities be that are introduced by nominal quantifiers? Should they always be states, or could they also be events? And what precise form should their temporal location take? Before I propose explicit answers to these questions formulated within the framework we have been using up to this point, let me first say at a more informal and intuitively level what it is we are after. For the case of our example the intuition we want to capture is that the eventuality characterized by the duplex condition is a complex eventuality that is composed of all the individual sleep events that the sentence is talking about. One implication of this is that the duration of this eventuality must include the durations of all the individual sleep events and it is this this implication that will be central to our implementation of the intuitive idea.

This implementation involves a somewhat different way of relating quantification eventualities to the times that locate them than we have so far seen. What relation we want is best explained in connection with a temporal locating adverbial like *On New Year's Eve 2001* in (3.113). The content of (3.113) is that the individual sleep events all occurred within the time denoted by the phrase *New Year's Eve 2001.* To do justice to this aspect of the truth conditions of (3.113) we need to set the duration of the eventuality that summarizes all these individual events equal to the time denoted by this phrase.

I believe that what seems plain for the relation between the eventuality and the temporal adverb of (3.113) should also apply to the relation between the eventuality and the time t contributed by tense. Why this should be so will become clearer in Section 4. There the simple past tense will be treated as the trigger of a presupposition whose resolution will often take the form of identifying t as some particular past time that is part of the context and is contextually salient. In such cases the duration of the eventuality introduced by the quantification should also be identified as this salient time.

This then is our verdict about the temporal relations between quantification states and the times contributed by tense and by temporal locating adverbs: This relation is always that of 'temporal identity": the relation that holds between an eventuality ev and a time t iff $dur(ev) = t'$.

This still leaves the question of the kind of eventualities introduced by nominal quantifiers – are they events or are they states? Given what we have said about how these eventualities related to the location times supplied by tenses and temporal adverbs this question seems to have become of secondary importance. The choice made in earlier versions of DRT was that they are always states. For instance, in the case of sentence (3.99) the eventuality introduced by *every philosopher* is a state that the world is in at some time because of the circumstance that each of the philosophers slept somewhere within the period covered by the state. What really matters for us at this point is how these 'states' are temporally related to their location times: 'Quantification states' temporally coincide with their location times. It is this that is crucial for getting the intuitively correct truth conditions for sentences involving nominal quantifiers. As far as these truth conditions are concerned it makes no difference whether these eventualities are called 'states' or by some other name.

We will see in our discussion of negation below that a distinction between quantification states and quantification events is useful and important for another reason. One implication of that discussion will be that the semantic representations of the sentences considered here – (3.99) and (3.113) – should involve quantification events rather than quantification states. For the remainder of the discussion in this section of the interaction of quantifiers and tense we will nevertheless stick to the older treatment, according to which nominal quantifiers always introduce quantification states.

So much for the intuitive motivation for an analysis of nominal quantification as involving quantification states. Formally these states have to be characterized by the duplex conditions that we have already agreed to adopt as representational form for the semantic contributions made by nominal quantifiers. That is, we assume that the duplex condition provides us with the type characterization of the quantification state, and we use the usual ':'-notation for typing eventualities to express this.

To implement this idea in the form of construction operations that are triggered by nominal quantifiers and that will lead to the representations we want ought to have been straightforward, and no more complicated that what than it is for the states introduced by progressives, habitually interpreted simple tenses and perfects. But it isn't completely straightforward. The reason is a difference that may feel like an almost accidental consequence of the syntactic structures we have adopted. The states introduced by progressives, habitually interpreted tenses and perfects are introduced before the construction algorithm reaches the tense feature that is specified by the T node (and before it reaches the temporal locating adverb of the clause, in case there is one,which we have been assuming, rightly or wrongly, that it is adjoined above T). So there is no problem with temporally locating these progressive, habitual or result states, which have been introduced already, at the point when the temporal locations have to be represented. But with the contributions by quantifying phrases that have been raised to adjunction at the highest TP node the situation is the reverse: the temporal location triggered by T has already taken place when the nominal quantifier introduces its quantification state. But it is really this state that should be located by the tense feature and (if there is one) the temporal locating adverb. This means that the temporal location or locations that have already been put in place have to be reinterpreted in the light of the quantificational structure that the interpretation of the nominal quantifier contributes. Formally this requires manipulating the representation of the adjunction site in ways that feel ad hoc and as hard to square with any natural understanding of semantic compositionality.

In the case before us – the DRS construction for (3.99) – the temporal location that has already taken place is for the individual sleep events of the different philosophers. At the point when the DP *every philosopher* introduces its quantification state s, these temporal locations have to be bundled, as it were, into a single location of the state s, which can be thought of as subsuming all these individual events. Formally, this requires the retrieval from the TP representation that has already been constructed of the dref t and the Condition ‘$t \prec n$’ and lifting them to the level of the DRS that contains the typing Condition of s. The dref t is placed provisionally in the store of the main DRS and $t \prec n$’ into that DRS's Condition Set. Furthermore the Condition ‘$e \subseteq t'iseliminatedandreplacedbytheCondition$‘$dur(s) = t$’, which is also added to the Condition Set of the main DRS.

The result we get when applying this complex set of operations to (3.104) is presented in (3.109).

(3.109)

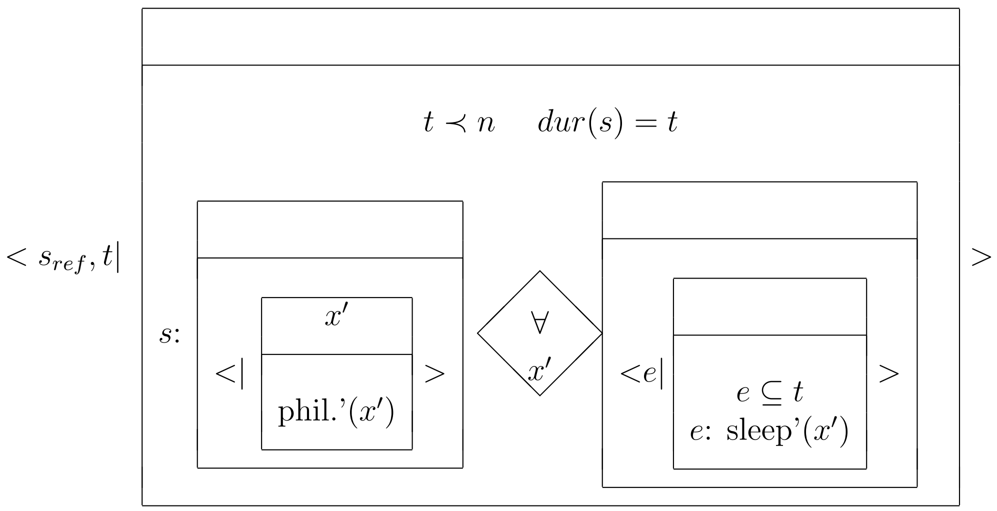

Because of the Condition ‘$dur(s) = t$’ in the Condition Set of the main DRS in (3.109) the DRS can be equivalently rewritten in the form (**??**), in which the Condition ‘$e \subseteq t$’ has been replaced by ‘$e \subseteq s$, capturing the intuitive idea that the events e are in some sense part of the state s’. (This is the form in which such representations are more often represented.)

(3.110)

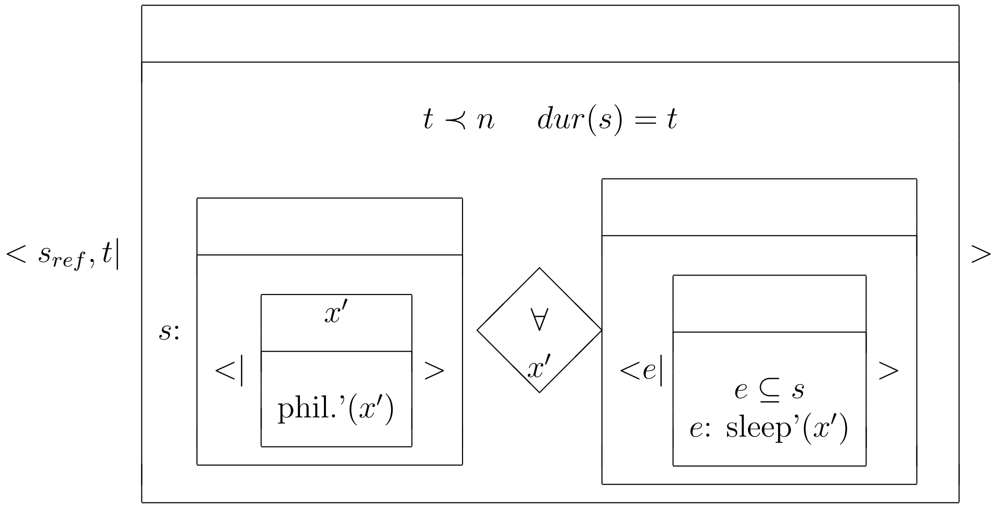

For a sentence like (3.113), which has a temporal adverb as well as a nominal quantifier, the problem of locating the quantification state occurs twice over, first in connection to tense and then in connection with the adverb. Temporal location by the adverb will also fall into place when presupposition is included in our framework. As already observed, in this modified framework definite DPs come with identification presuppositions. This is true in particular for *New Year's Eve 2001*, which is a kind of proper name. Resolution of the presupposition for this name will introduce the dref reprinting it – some time dref t', say – at the opt level, and with it its characterizing Condition, which web here conveniently abbreviate to "New-Year's-Eve 2001"(t'). We show the final DRS for this sentence (i.e. after the remaining drefs in the stores have been transferred to the Universes to their right).

(3.111)

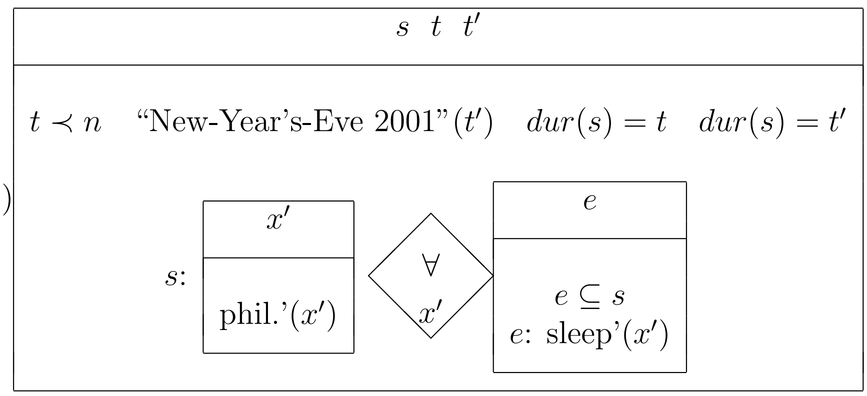

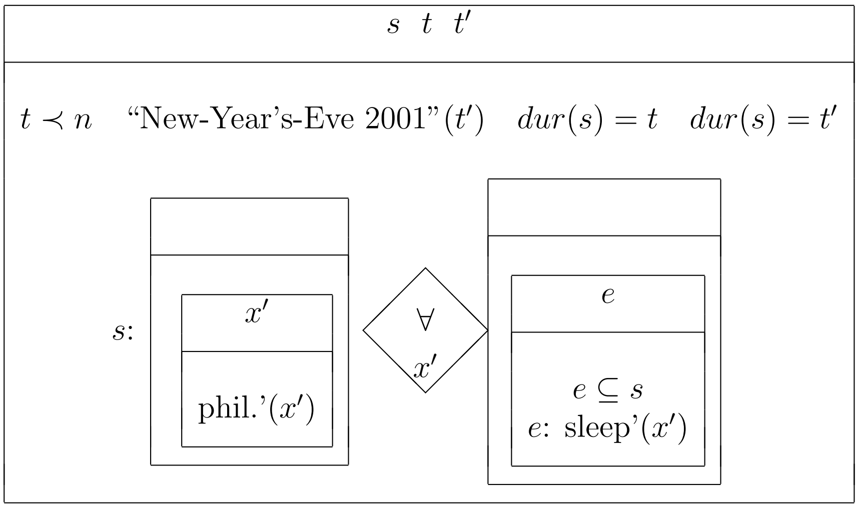

In the light of the considerations that led to the entry (3.106) for *every* should be modified. The new entry is given in (3.112).

(3.112) Lexical entry for the determiner *every* (modified)

every (Determiner)

Sel. Restr:

Sem.Repr:

$<x_{ref}, .. \mid K> \;\rightsquigarrow$

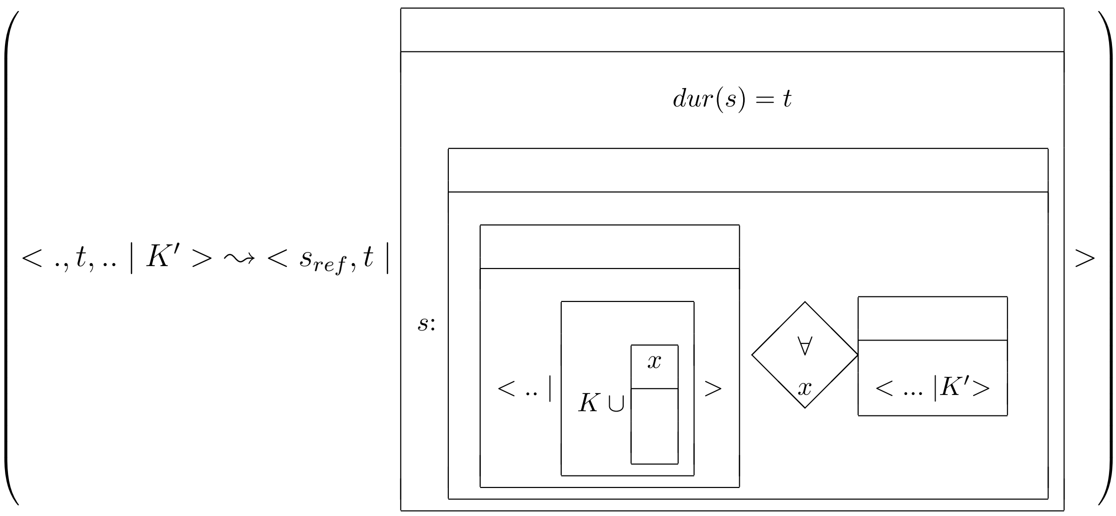

In Sections 3.11.4 and 3.11.5 it will be propose that the natural treatment of temporal quantification via adverbial quantifiers like *always* or *often* is also by way of representations of quantificational states. Once this view is adopted for adverbial quantifiers, it is then also natural to extend it to temporal quantifications that are expressed by nominal quantifying phrases like *every day* or *every 1-st of the month.* And once that step has been taken, the further extension to all cases of nominal quantification – irrespective of whether the quantification is over times, philosophers or whatever –will then be an even more natural one.

A remaining question is whether the eventuality that is introduced by the quantifying expression (and whose content is given by the duplex condition) should always be a state. In Section 3.10.1 we will find reasons that speak against making this general assumption. And in fact, those reasons suggest that in the case of (3.105) and (3.113) the eventuality introduced by the quantifier is – like the referential argument of the second of the two input representations, not a state but an event. (**??**) in which the remaining store elements have been transferred to the Universe of the main DRS displays the result. For the discussion of the general question what kinds of eventualities those introduced by quantifiers should be. see Section 3.10.1 .

Exercise: Using the principles illustrated in the representation construction above for sentence (3.99) construct semantic representations for the following sentences:

(3.113)a. Fred met every philosopher.

b. Fred has met every philosopher.

c. No linguist met every philosopher.

(In this last example treat *no* as a quantifying determiner which gives rise to duplex conditions whose central component consists of a quantifier-denoting term (for which it is natural to use the symbol ‘ $\forall\neg$’) together with the dref that the quantifier binds. State the verification conditions for duplex conditions with this quantifier.)

Exercise: Construct a DRS for (3.99) in which the noun *philosopher* is treated as a state description, with a lexical semantics of the form:

$$< s \mid \begin{array}{|l|}\hline \\ \hline s\text{: philosopher'}(\underline{\text{x}}) \\ \hline \end{array} >$$

Pay special attention to where the state s of this lexical semantic entry can/should be temporally located within the DRS for the sentence.

3.9.2 Indefinites

In our review in PART I we adopted the principle that indefinites can be analyzed both as existential quantifiers and as indefinite singular terms. This principle is independent from the choice between Top Down and Bottom Up construction and there is no reason to abandon it here. But within the setting of our current tense & aspect-sensitive Bottom Up construction method each of these two analyses presents its own new problems. These problems aren't fundamentally different from what we have encountered in our Bottom Up treatment of other constructions. But it will nonetheless be useful to go through the representation construction for a couple of examples with indefinites to see in detail how they can be treated within our present architecture. This will also give us a first impression of how a Bottom Up construction algorithm can deal with non-quantificational noun phrases.

We begin by having another look at the first sentence of our age-old example (2.1). (This time we only consider the first sentence, since it is only the interaction between the past tense and the indefinite DP of this sentence that is of importance for the present discussion, and also because we have no adequate treatment at this point for the anaphoric pronouns of the second sentence.)

(2.1) Pedro owns a donkey. He beats it.

Our first task is to decide on the syntactic inputs from which the semantic representations for our sentence are to be derived. The discussion of indefinites in PART I led us to two different analyses for them, as singular terms and as quantifying phrases, and to two distinct structures at LF, one in which the indefinite DP is quantifier raised and one in which it remains in situ. We retain these assumptions about LF structure. So for the first sentence of (2.1) we get the two LFs in (3.114) – (3.114.a) for the term analysis of its indefinite and (3.114.b) for the quantifier analysis.

(3.114)

a. [$_S$ [$_{Comp}$ $\emptyset$] [$_{TP}$ [$_{DP}$ *Pedro*] [$_{T'}$ [$_T$ pres] [$_{VP}$ [$_V$ *own*] [$_{DP}$ [$_{Det}$ *a*] [$_{NP}$ [$_N$ *donkey*]]]]]]]

b. [$_S$ [$_{Comp}$ $\emptyset$] [$_{TP}$ [$_{DP_2}$ [$_{Det}$ *a*] [$_{NP}$ [$_N$ *donkey*]]] [$_{TP}$ [$_{DP}$ *Pedro*] [$_{T'}$ [$_T$ pres] [$_{VP}$ [$_V$ *own*] t_2]]]]]

First the DRS construction for the LF in (3.114.a). We begin by replacing the lexical items *Pedro*, *own* and *donkey* in (3.114.a) by their semantic representations. Moreover, we move the semantic representation of the N-node of the direct object straight up to the NP node that dominates it. These operations lead to (3.115).

(3.115)

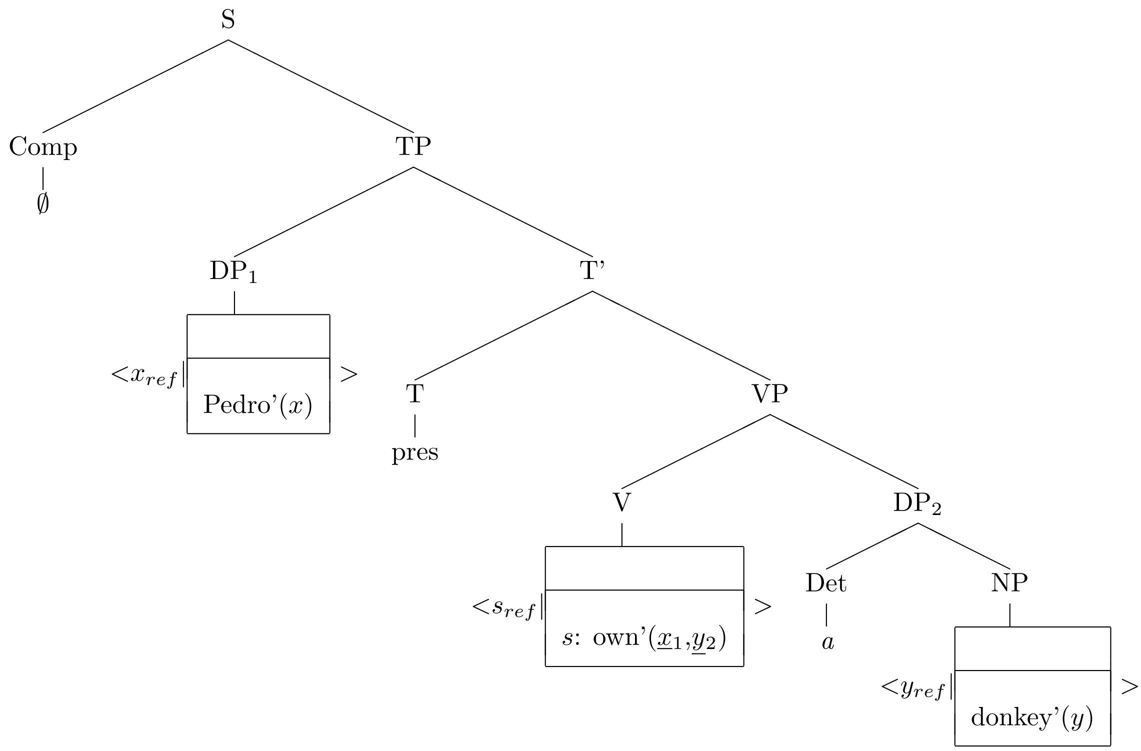

In view of the decision to raise quantifying DPs while leaving DPs analyzed as terms in situ, the position of the DP *a donkey* in (3.115) is an indication that this DP is being treated as a term. So the semantics must treat it as a term as well. Just as we did for the quantifying determiner of the last section, we need a lexical entry for the word *a* – here as the determiner of indefinite DPs interpreted as terms – that spells out how (in particular) the VP representation of (3.115) is computed from the representations of V and DP$_2$. The result of applying this lexical entry to (3.115) is shown in (3.116).

(3.116)

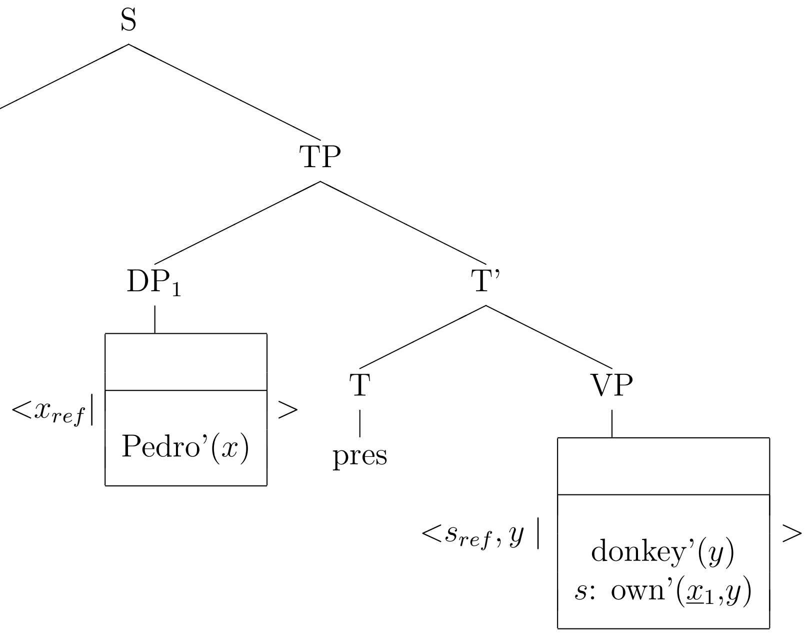

Combining the state description provided by the VP in (3.116) with the feature pres is unproblematic and familiar, and so is the combination of the result of this operation with the representation of the subject DP *Pedro*. These two operations lead to the TP representation in (3.117).

(3.117)

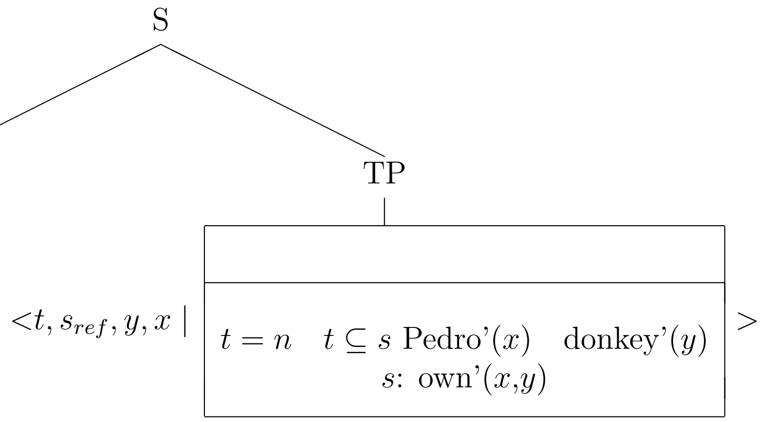

This is a TP representation of the kind familiar from earlier sections, in which there is just one store. The final step consists in eliminating that store, by transferring all its drefs to the Universe of the DRS it precedes. The result is shown in (3.118).

(3.118)

$$\begin{array}{|c|}\hline t \quad s \quad x \quad y' \\ \hline t = n \quad t \subseteq s \quad \text{Pedro'}(x) \quad \text{donkey'}(y') \\ s\text{: own'}(x,y') \\ \hline \end{array}$$

In the LF (3.114.b) the indefinite DP *a donkey* has been raised and adjoined to TP, indicating that it is treated as a quantifying phrase. The construction rule for indefinites analyzed as quantifiers is closely similar to that for *every*-phrases: *a* is treated as the introducer of an eventuality characterized by a duplex condition, but now with an existential instead of a universal quantifier. The lexical entry that is needed to construct a complete DRS from this LF is a variant (with an existential instead of a universal quantifier inside the diamond of the duplex condition) of the revised entry for *every* in (3.110).

(3.119) shows the point in the representation construction from (3.114.b) at which the representations for the lower TP node and the NP of the DP *a donkey* have been constructed. Note that the trace in (3.114.b) is handled in the same way as we handled the trace that waqs left behind by the quantifying DP *every philosopher* in (3.99).

(3.119)

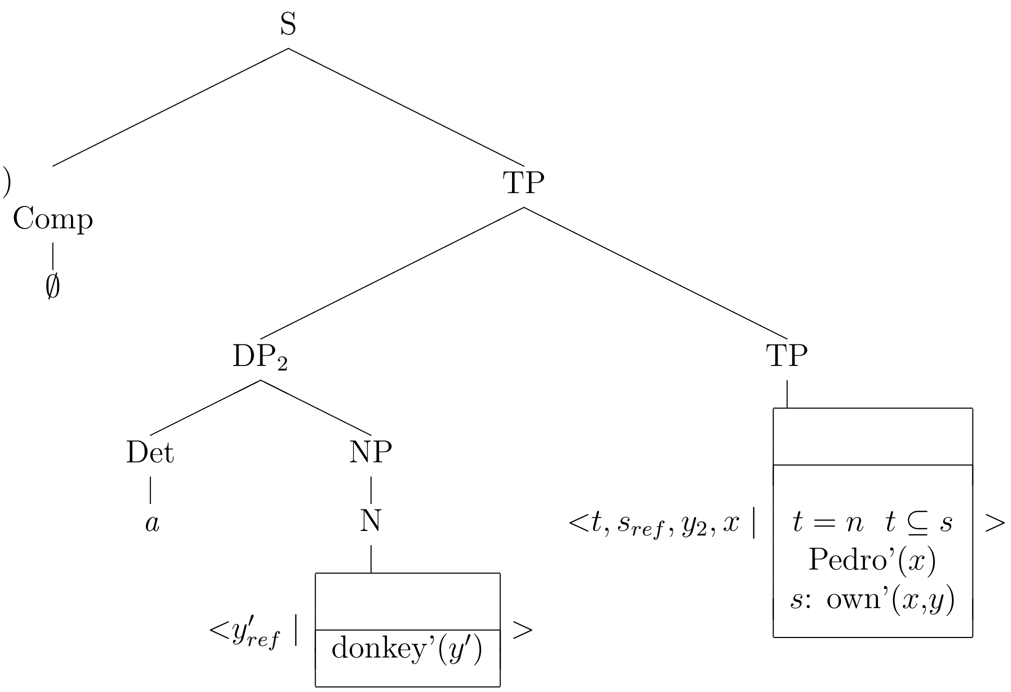

The quantifying determiner *a* now makes its contribution by introducing an existential duplex condition. Once again we run into the problem of having to reinterpret the locating function of the time t introduced by the tense. We perform the same lifting of the dref t introduced by the tense feature as we did when ealing with the DP *every philosopher* in the last section. The DRS in (3.120.a) is the result of these operations together with the final transfer of discourse referents from store to DRS Universe.

(3.120)a.

$$\begin{array}{|c|}\hline s' \quad t \\ \hline \\ t \subseteq n \quad t \subseteq s' \\ \\ s'\text{:} \ \begin{array}{|c|}\hline y' \\ \hline \text{donkey'}(y') \\ \hline \end{array} \ \left\langle \begin{array}{c} \exists \\ y' \end{array} \right\rangle \ \begin{array}{|c|}\hline s \\ \hline s' \subseteq s \\ \text{Pedro'}(x) \\ s\text{: own'}(x,y') \\ \hline \end{array} \\ \\ \hline \end{array}$$

b.

$$\begin{array}{|c|}\hline t \quad s \quad x \quad y' \\ \hline \\ t \subseteq n \\ t \subseteq s \\ \text{Pedro'}(x) \quad \text{donkey'}(y') \\ s\text{: own'}(x,y') \\ \hline \end{array}$$

Because its duplex condition is existential, the DRS in (3.120.a) can be simplified: The existence of a state s' at time t to the effect that the duplex condition in (3.120.a) is satisfied amounts to the same thing as the existence of a state s holding at t of the type described in the nuclear scope. When (3.120.a) is simplified in this spirit, we get as result the DRS in (3.120.b). Note that this is just the representation that we obtain for the first sentence of (2.1) when we analyze the indefinite DP as an indefinite term. Thus, just as we saw in PART I, for a simple sentence like the first sentence of (2.1) it ultimately makes no difference to the semantics whether we choose to treat its indefinite as a quantifier or as a term – either choice leads eventually to the same semantic representation. But we will find in Section 4 that these two ways of analyzing indefinite DPs do not always lead to equivalent representations. Term treatments of indefinites can lead to interpretations that cannot be obtained when they are analyzed as quantifiers.

Exercise: The transition from (3.120.a) to (3.120.b) is justified by a Meaning Postulate for existential duplex conditions. Formulate this Meaning Postulate.

3.9.3 Another syntactic construction involving traces: Relative Clauses

We looked at relative clause formation in PART I, but not so far in PART II. However, pretty much everything that is needed to deal with relative clauses in the setting of our present Bottom Up architecture is already in place. This is true in particular for the semantic handling of traces, which the LFs for relative clauses share with the LFs for sentences with quantifying DPs. We illustrate the semantic processing of NPs with relative clauses at the hand of example (3.121). The LF of (3.121) is given in (3.122). (Compare this with (2.23) in PART I.)

(3.121) Pedro knew a farmer who owned a donkey.

(3.122)

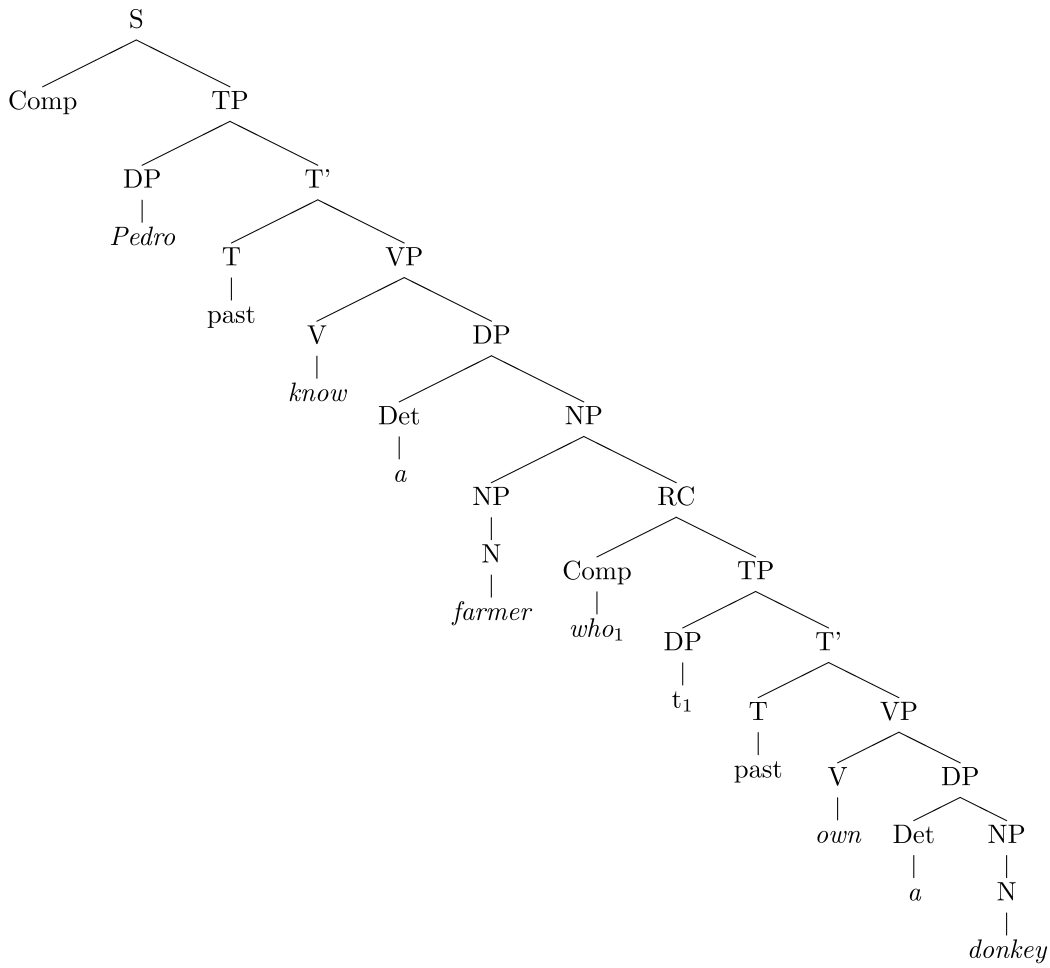

Our primary concern is with the direct object DP of (3.122). (3.123) shows the point where the semantic representation for the T' node of the relative clause has been constructed.

(3.123)

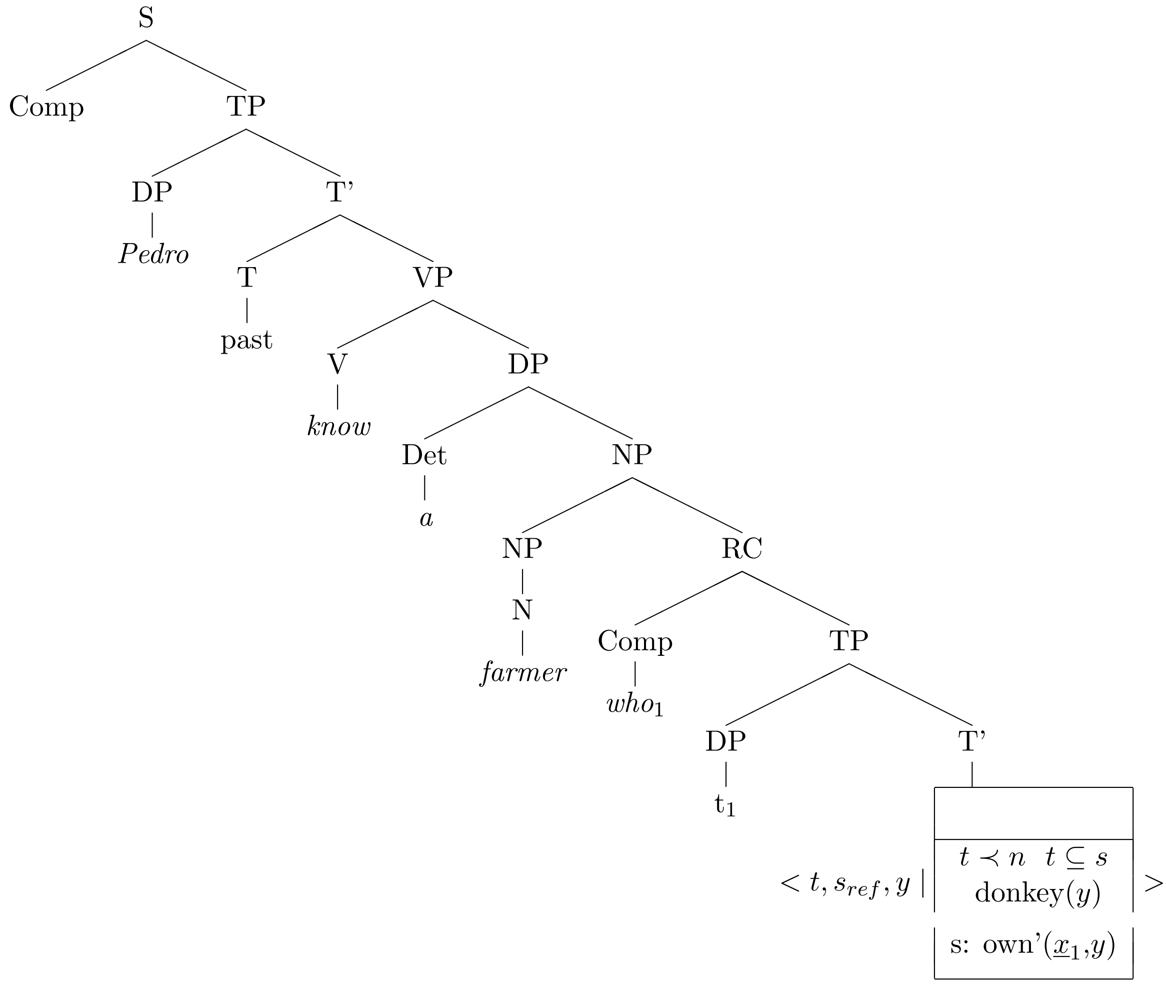

The next step in the representation construction must combine the T' representation with that of the DP. This is an instance of an operation that is by now well familiar to us – that of argument insertion – but there is a new twist to this operation because the DP is now a trace. How should traces be handled by our bottom up construction algorithm? Our guide in deciding this matter is the way in which traces are handled in certain versions of Montague Grammar (see e.g. (Heim & Kratzer 1998)): the trace is turned into a variable that bears the same index; this variable fills the argument slot of the DP and can, at the point where the relative pronoun makes its contribution, be resumed (e.g. for lambda abstraction, as in (Heim & Kratzer 1998)). In our set-up discourse referents do the work of variables. So if we want to follow the Montague Grammar treatments alluded to as closely as our framework allows, we should let the trace introduce a new dref. This dref gets inserted for the argument slot with which the trace is coindexed. In addition we place the dref in the store and index it with the index of the trace, so that it too becomes coindexed with the relative pronoun, something that will be needed in the next step.

The result of these operations, with x as the chosen dref, is shown in (3.124).

(3.124)

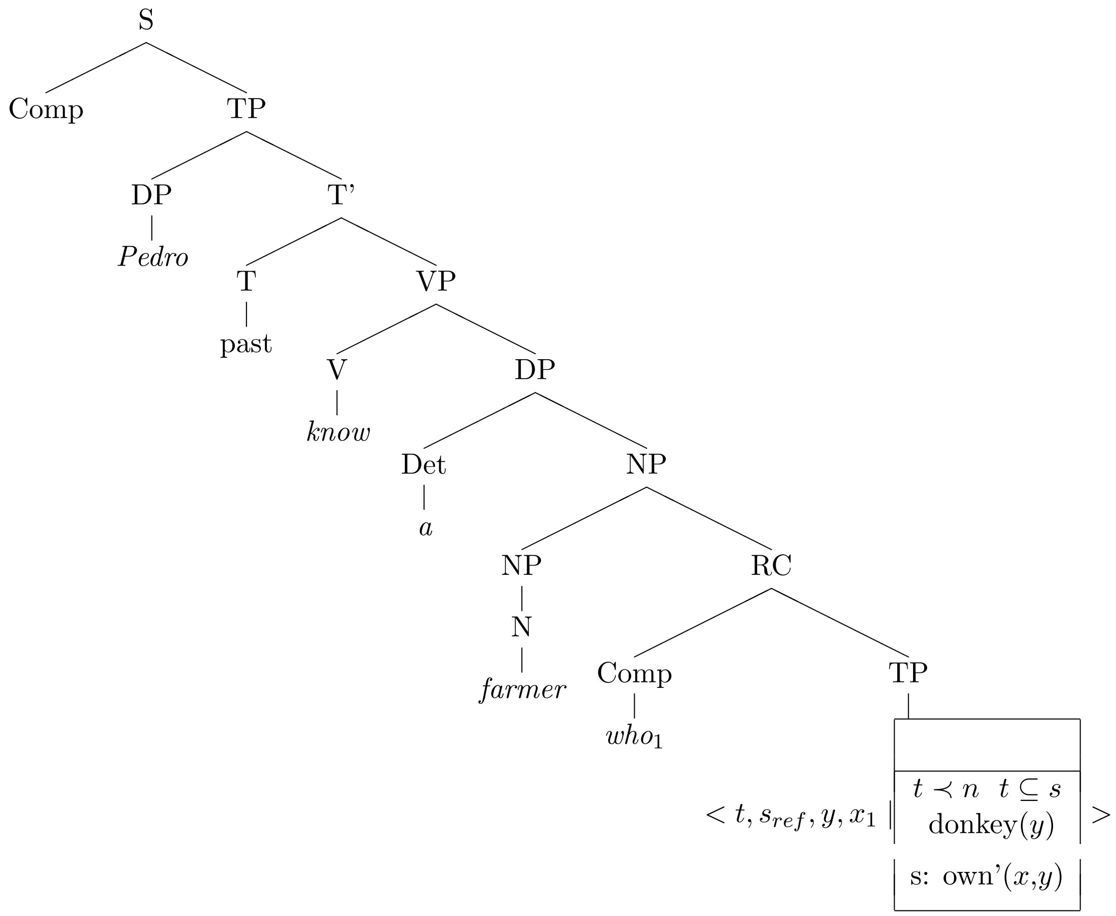

How do we move from the representation of the TP node to that of the RC node? The first thing to observe in this connection is that the structure of a relative clause as it is presented here is essentially the same as the structure of main clauses that we have been assuming all along (from the moment we have been building DRSs bottom up): A structure in which the highest node has two daughters, one labeled 'Comp' and the other labeled 'TP'. In fact, the natural way to view relative clauses is as sentence structures of a special sort, on a par with other types of subordinate clauses on the one hand and with main clauses on the other.[43] We capitalize on this similarity by borrowing as much as we can from our treatment of main clauses, in which the last step consists of transferring the remaining drefs in the store of the TP representation to the main Universe of the DRS following the store.

But of course in the case of a relative clause that can't be all. The output representation of the operation we are trying to define, that of the RC node, must have a form in which it can be combined with that of the RC's adjunction site, one that will always be an NP. In formulating operations that get us to such a form we take our inspiration once more from treatments of relative clauses that can be found in Montague Grammar (again, see (Heim & Kratzer 1998)). In some such treatments (including the one we have looked at in our selective survey of (Heim & Kratzer 1998)) the relative pronoun gives rise to a λ-abstraction over the variable introduced by its trace. This creates a λ term that functions like a 1-places predicate, which can then be combined with the NP to which the RC has been adjoined. (This NP also has the status of a 1-place predicate.) We cannot reproduce such a treatment to the letter, as the system in which we are working has neither variables nor λ-abstraction. But we can get the effect of making the RC representation play the part of a 1-place predicate by promoting the dref introduced by the trace of the relative pronoun to one with the status of the referential argument of the RC representation.

As it stands, this last principle – of making the trace dref into the referential argument – isn't compatible with the principle that all store elements should be transferred to a DRS universe. So we have to compromise. The compromise is that all drefs in the store of the TP representation of a relative clause are transferred to the Universe of the DRS following the store *with the exception* of the trace dref, which survives as only member of the store of the RC representation. (Transferring the other drefs in the store of the TP representation at this point won't have any adverse effects later on, since they are no longer needed in further construction operations.)

The result of applying these operations to the TP representation in (3.124) leads to the structure in (3.125).

[43] Perhaps it would have been in better keeping with this perspective if we had used a label like 'MC' for the highest node of a main clause, instead of 'S'. But the use of the labels 'S' and 'RC' has a certain tradition. Main clauses are the sentences *par excellence* in that each complete sentence must have a main clause, but many sentences consist of a main clause and nothing more.

(3.125)

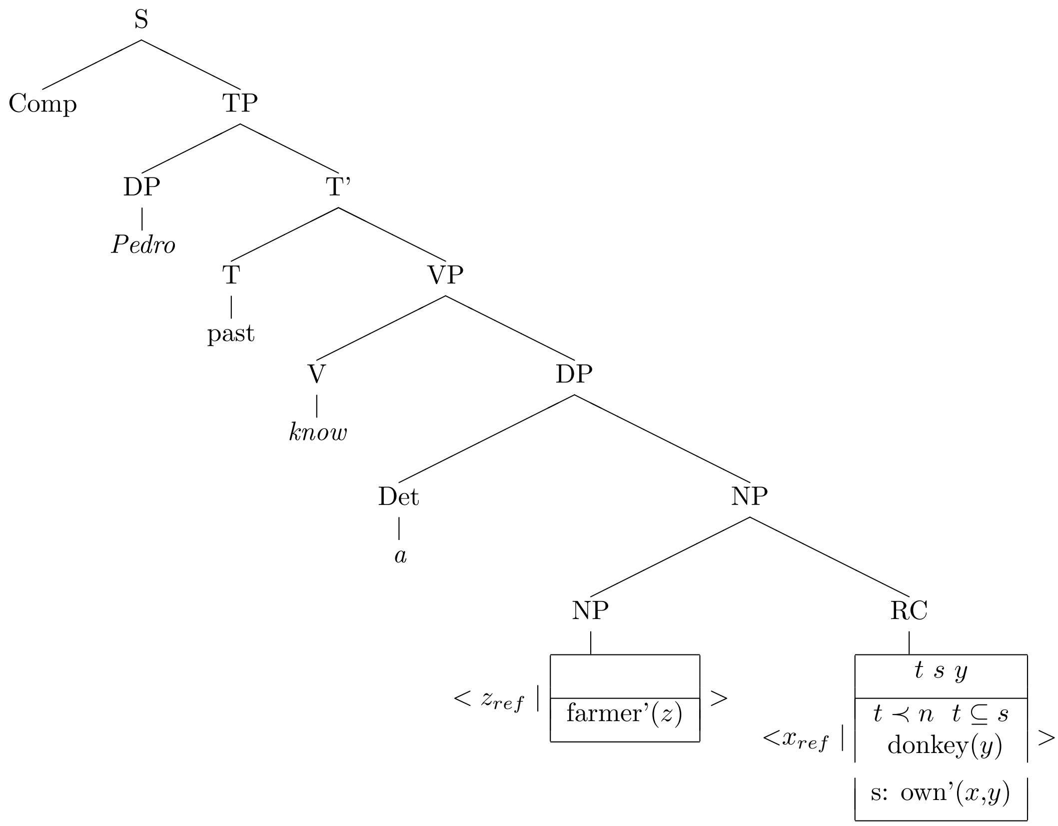

The step that follows now is the one that combines the RC representation with that of the NP to which the RC has been adjoined. As in all other instances of adjunction that are considered in these notes, the semantics of adjunction is implemented as a form of unification. More precisely, it consists in unifying the referential arguments of the two representations, followed by merging the representations that result from this unification. Unification of the two referential arguments is implemented by letting the referential argument of the adjunction site replace the referential argument of the adjunct in all its occurrences and thus eliminating this second referential argument from the resulting representation altogether. In the case before us the referential argument z of the adjunction site representation replaces the referential argument x of the RC, see (3.126).

(3.126)

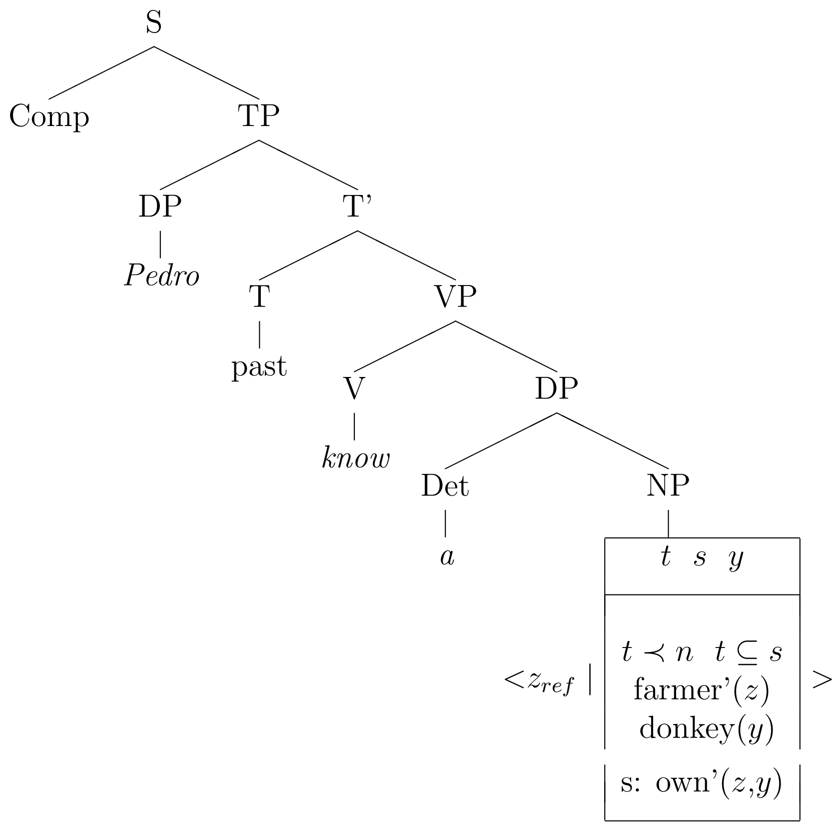

From this point onwards the construction follows a path that has been traced repeatedly and in detail in previous constructions. We leave this remainder of the construction to the reader and just show what the final result looks like.

(3.127)

$$\begin{array}{|c|}\hline t' \quad s' \quad p \quad t \quad s \quad y \quad z \\ \hline t' \prec n \;\; t' \subseteq s' \;\; t \prec n \;\; t \subseteq s \\ \text{Pedro'}(p) \quad \text{farmer'}(z) \\ \text{donkey}(y) \\ \text{s': know'}(p,z) \\ \text{s: own'}(z,y) \\ \hline \end{array}$$

<u>Exercise:</u>

(i) Fill in the remaining steps of the DRS-construction for (3.121) that are needed to get from (3.126) to (3.127).

(ii) Construct the DRS shown in (3.127) while treating the indefinite as an existentially quantifying DP.

3.9.4 Negation

In a language like English pretty much every negation-free indicative sentence has a negated counterpart, which denies what the unnegated sentence asserts. That is the simple story, and it is correct as far as it goes. But it doesn't go very far, and one main reason for that it that it ignores what negation can do to aspect. Once aspect is taken seriously, and in particular when you are alert to the distinction between event-describing and state-describing sentences, you cannot fail to notice that there are important differences between what negation does to sentences with the one kind of aspect and what it does to sentences with the other kind.

The easier case, it seems, is presented by negations of state descriptions. Consider the sentences in (3.128).

(3.128)a. Johnny is happy.

b. Johnny lives in Paris.

c. Johnny isn't happy.

d. Johnny doesn't live in Paris.

According to our analysis (3.128.a) and (3.128.b) both assert that a certain state holds at the utterance time n, something that could hardly be controversial. Much the same can be said about their negations (3.128.c) and (3.128.d): they also say that certain states hold at n, and these states are the exact opposites of those described by the first two sentences. More precisely, for each state type (or description determining that type) there is a corresponding complement type (or complement state description) which is instantiated whenever the first type is not. One of the uses of negation is to transform a given state description S into its complement description. We will informally refer to this latter description as '$not - S$'.

What exactly does it mean to say that 'the state description not-S is instantiated whenever S is not'? Suppose that t is any time and that S has no instantiation at t, in the sense that there is no state s of type S such that $t \subseteq s$. Should we conclude that there must then be a state s' instantiating $not-S$ such that $t \subseteq s'$? Clearly not in general. For the times that we use in our treatment of tense and aspect typically have temporal extension, and when t is such a time, then it is clearly possible that it overlaps partly with a state s of type S and partly with a state s' of type $not-S$, in which case we will have neither $t \subseteq s$ nor $t \subseteq s'$.[44]

We will just have to accept this. In general complementary state descriptions S and $not-S$ do not satisfy the disjunction in (3.129).

(3.129)Either there is a state s of type S such that $t \subseteq s$ or there is a state s' of type $not-S$ such that $t \subseteq s'$.

A weaker principle than (3.129) is (3.130).

(3.130)Suppose that for no $t' \subseteq t$ there is a state s instantiating S such that $t' \subseteq s$. Then there is a state s' instantiating $not-S$ such that $t \subseteq s'$.

Even (3.130) is problematic when t is included in a period of transition from S to $not-S$ or vice versa. But it has far fewer exceptions than (3.129) and tends to be satisfied when t is short enough. In particular, (3.130) is often satisfied when t is the time n of a present tense utterance. That is why sentences like 'Either Johnny lives in Paris or he doesn't.' and 'Either Johnny is happy or he isn't.' strike us for the most part as obvious and trivial truths. (If such sentences convey any non-trivial meaning at all, it is in a meta-logical capacity, as saying that they are among the cases in which the principle of bivalence – the principle that either a sentence is true or else the negation of that sentence is true – does after all apply.)

To sum up this part of the discussion:

[44]The best we can hope for is that the disjunction 'either there is a state s of type S such that $t \subseteq s$ or there is a state s' of type $not{-}S$ such that $t \subseteq s'$ holds whenever t is short enough. In particular, we might hope that the disjunction holds when t is instant of time (i.e. when it is duration-less). Even that assumption is not self-evident. For what should we say about a time when Johnny goes from happy to not-happy? Is there then a last instant of his being happy? Or a first instant of his being unhappy? Or is there a transition period (however brief) during which Johnny is neither clearly happy nor clearly unhappy, and which can only be described as 'a time when Johnny changed from happy to unhappy'? These are among the perennial problems of the logic of time and change.

• The effect of negation on a state describing sentence is to transform the state description of that sentence into the complement state description.

• A state description S and its complement description $not{-}S$ are mutually exclusive in that it cannot be true for any t that there is both a state s of type S such that $t \subseteq s$ and also a state s' of type $not{-}S$ such that $t \subseteq s'$.

• But S and $not{-}S$ are not jointly exhaustive: the disjunction in (3.129) is not always satisfied; nor can we be sure that the weaker (3.130) holds without exception.

• However, for short intervals of time t both (3.130) and (3.129) have a good chance of being true and the ways in which we express ourselves in language often presuppose that they are true.

In what follows we set the possibility of counterexamples to (3.130) aside and adopt the supposition of (3.130) – for no $t' \subseteq t$ there is a state s instantiating S such that $t' \subseteq s$ – as the definition of the extension of the state description $not{-}S$. In other words, (i) for any model M and any state description S of our representation language the extension of $not{-}S$ will contain a state s' of duration t for every t such that for no $t' \subseteq t$ there is a state s instantiating S such that $t' \subseteq s$; and (ii) all the states in the extension of $not{-}S$ are of this kind; that is, for each such state s' there is a t as described under (i) such that $\mathrm{dur}(s') = t$.

A different problem arises for negations of state describing sentences with a tense other than the Present Tense. (3.131) presents past tense versions of the negations in (3.128).

(3.131)a. Johnny wasn't happy.

b. Johnny didn't live in Paris.

Such sentences raise the question: At what time or times in the past is the complement state description – that of being a state of John not being happy or of his not being in Paris – supposed to have been instantiated? This of course is a problem that also arises for the corresponding unnegated sentences. In fact, it is a general problem for past tense sentences. It arises both for sentences that describe states and for sentences that describe events, although – as we will see more clearly in Sentence 4 – not quite in the same way. As we have noted before – and as we will be in a position to make formally precise only in Section 4 –, simple past tense sentences tend to be understood as presupposing some particular past time t that is salient in the discourse context. Unnegated sentences are taken to locate the eventualities they describe at t. Negated sentences say of t that an eventuality of the kind described by their unnegated counterpart did not hold or occur at t. But what exactly does that last statement amount to?

For negations of state describing sentences the explication is straightforward and more or less dictated by what we have already said: such sentences simply assert that the states they describe held at the past times they presuppose, and those states are the complements to the ones described by the corresponding unnegated sentences. But for negations of event descriptions the matter is different. They assert the absence at or during some time t of any events of the kinds described by the event descriptions they negate. Examples of such past tense negated event descriptions, all of which illustrate this point, are given in (3.132).

(3.132)a. Johnny didn't cry.

b. Mary didn't meet a senator.

c. Mary didn't greet a senator.

All these sentences invite a paraphrase along the lines of: 'At some given time in the past there was no event of the kind described'. The insight that this is how negated event sentences work – that they assert that there was no event of the described kind at or over some independently determined interval of time – was captured early on in the history of formal semantics in a short but seminal article by Partee (Partee 1973). In the dominant theory of tense and temporal reference at the time the past tense was treated as an existential quantifier over all past times. Partee's telling example in that paper, the sentence (3.133) as said by someone to her spouse when they have just set off from their home for a long drive, cannot be analyzed adequately when it is assumed that the past tense expresses such an unrestricted quantifier over past times.

(3.133) I didn't turn off the stove.

Such an account of the pas tense would allow for the following two interpretations, one for when negation is assigned wide scope and one for when it is assigned narrow scope with respect to the past tense: (i) 'It is not the case that there was a time when I turned off the stove' (wide scope); (ii) 'There was some time in the past when I did not turn off the stove' (narrow scope). Clearly neither of these paraphrases is any good. The first makes a claim that is almost certainly false – the speaker, it may be assumed, turned off the stove on many occasions – and the second is sort of tautological – surely the speaker can't have spent her entire life turning off the stove. Neither paraphrase captures the intuitive content of (3.133) – the contingent statement that the speaker didn't turn off the stove during the comparatively brief period when they were closing up the house.

Note that the semantics we have so far given for the past tense suffers from the very defect that Partee exposed. For our semantics for the feature past is a purely existential one: there has to be some time in the past of n that locates the described eventuality. But the only way in which we can improve our defective account is by doing justice to the presuppositional dimension of tenses and that will be possible only as part of revamping our syntax-semantics interface into one can deal with presuppositions, as we will do in Section 4.

Partee's observation in relation to (3.133) was that an existential treatment of the past tense leads to the wrong truth conditions whether negation is assumed to take scope over it or under it. This way of arguing the case is appropriate because there is no good way of telling from the morpho-syntactic structure of the sentence what the scope relation between tense and negation really is. This is an instance of a notorious problem about negation of a more general kind. Quite often negation is expressed by a suffix on the verb, or by some other device that occurs in close vicinity to the verb and has the morphological features of a verb modifier. But as a rule these forms of morphological realization should not be taken to entail that the scope of the thus realized negation operator that is morphologically realized in some such way is restricted to the verb, and that all other sentence constituents take scope over it. In the case of negation morpho-syntactic scope and semantic scope are very different things, and the former is but a poor guide to the latter.

However, if we want to say anything specific about the construction of semantic representations for sentences that contain negations, then some decision needs to be made about the scope relations that negations stand in to other scope-taking constituents, even if that decision isn't the final word. So here are the – partly stipulative – assumptions that we are going to make: The only instances of negation we are going to consider are operators that act on the representations of those syntactic constituents that are sisters to T; given the commitments about syntactic structure we have made these are the semantic representations of the nodes labeled PerfP. We assume moreover that

Neg has its own projection level, situated between the level of Perf and that of T. An illustration of what these assumptions come to is the following syntactic structure in (3.134.b) for the sentence in (3.134.a). (The sentences we will look at later in this section will have the Prog and Perf values -prog and -perf. In the syntactic structures for these sentences we will skip the Perfect and Progressive projection levels when dealing with these sentences, with the effect that the eventuality description on which the negation operates is that of the VP.)

(3.134)

a. Fred hasn't paid the rent.

b.

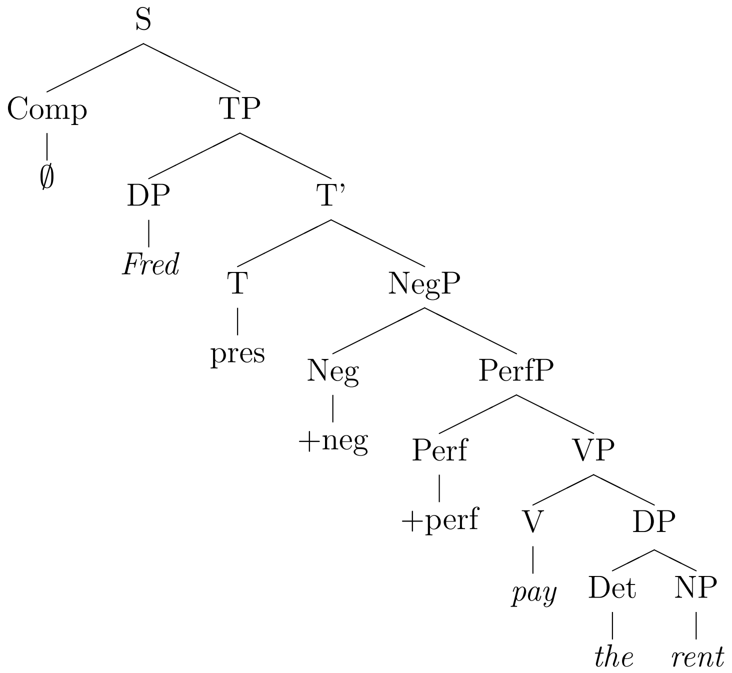

As this syntactic structure suggests, the Neg operator, triggered by the feature value +neg, gets the PerfP representation as input and has to transform this representation into an output representation. What should this output representation be like? This is a tricky question. On the one hand one might think that the result of applying negation to an eventuality description should be a state description – the description of a state that holds throughout a period of time by virtue of there being within that time no eventuality of the kind described by the input representation. This intuition seems especially strong in relation to sentences for which the input representation to Neg is an event description, like those in (3.132). For instance, it is natural to think of the phrase *didn't cry* in (3.132.a) as describing a state that held throughout the relevant past location time. When applied to event descriptions, the intuition suggests, negation has the effect of an aspect shift: it turns event descriptions into state descriptions. (For discussion see e.g. (de Swart & Molendijk 1999).)

Unfortunately there is a problem with this view of negation as transforming all eventuality descriptions into descriptions of states. Or at any rate, this is a problem for us. It has to do with our identification of the opposition between Perfective and Imperfective Aspect with that between event descriptions and state descriptions. The problem shows up in a comparison of a pair of sample discourses considered in the introduction to PART II, (3.2), with variants in which the second sentences have been replaced by their negations. The two discourses from (3.2) are repeated below as (3.135.a,b) and their variants are given in (3.135.c,d).

(3.135)

a. When Alan opened his eyes he saw his wife who was standing by his bedside.
   She smiled.

b. When Alan opened his eyes he saw his wife who was standing by his bedside.
   She was smiling.

c. When Alan opened his eyes he saw his wife who was standing by his bedside.
   She didn't smile.

d. When Alan opened his eyes he saw his wife who was standing by his bedside.
   She wasn't smiling.

In our first inspection of (3.135.a,b) we noted that (3.135.a) conveys that Alan's wife smiled in reaction to Alan opening his eyes, while (3.135.b) suggests that when Alan opened his eyes she was smiling already. The point of our revisiting these examples is that there is a similar distinction between (3.135.c) and (3.135.d): (3.135.c) is concerned with the way Alan's wife reacted to the event e of him opening his eyes, (3.135.d) with the conditions prevailing at the time when e took place. The only difference with (3.135.a,b) is that (3.135.c,d) speak of an absence of smiles rather than their presence.

That negation preserves the difference between (3.135.a) and (3.135.b) is something that any systematic way of dealing with the semantics of negated event and state descriptions should be able to account for. In traditional terms the relevant difference between (3.135.c) and (3.135.d) would be described in terms of the Perfective-Imperfective opposition: (3.135.c) has Perfective and (3.135.d) Imperfective aspect, and that difference is responsible for the different ways in which these two sentences interact with the first sentence of (3.135.c,d). Further data support this assumption. Compare for instance the difference between the present tense sentences in (3.136).

(3.136)

a. Johnny isn't crying.

b. Johnny doesn't cry.

The prominent interpretation of (3.136.a) is that at the utterance time n Johnny is not in a crying state. This is a purely contingent, episodic description of what is going on at the utterance time, which need not be taken to carry any implications about Johnny's general dispositions or tendencies. (3.136.b), on the other hand, only seems to have such a generic interpretation, it describes Johnny as someone who isn't a cry-baby. This too is a difference that is commonly associated with the Perfective-Imperfective opposition: Perfective aspect is incompatible with the use of the present tense as a way of describing what is the case at the utterance time (that use which in Section 3.7.2 we agreed to call the 'Standard Use' of the Present Tense). For unnegated sentences the relevant distinction – between 'Johnny cries' and 'Johnny is crying' – was discussed at some length in Section 3.7.2. (3.136.a,b) show that negation preserves this distinction.

These observations suggest the following general hypothesis:

(3.137) Perfective and Imperfective aspect are both preserved by negation.

Given our earlier decision to identify the perfective/imperfective distinction with the distinction between event and state descriptions, this leaves us with a choice between two options. Either we abandon the assumption of a strict correlation between the two oppositions, or we bite the bullet that our assumptions force us to bite and treat the eventualities described by the negations of event descriptions as events. My proposal is that we go for the second option. This means in particular that the second sentence of (3.135.c) will be analyzed as the description of an event, to the effect that while this event is going on there is no occurrence of an event of the type of Alan's wife smiling. In context – and, more specifically, in the context of a larger discourse – this event description will then be handled in the same way as other event descriptions, and in particular like those that do not involve negation. (Note well, it isn't possible to explain at this point how treating the negations of event descriptions as event descriptions is going to help in accounting for the discourse effects illustrated by (3.135) and (3.136). This is a problem that we will only be able to address in Section 4.)

How plausible is the treatment of negations of event descriptions as event descriptions? Is it reasonable, for instance, to speak of an *event* of someone not smiling? My own feeling about the value of 'intuitions' concerning matters in this area is that not too much weight should be attached to them. But to the extent that intuition deserves to be taken seriously, the idea of an event of x not smiling, i.e. of x *withholding* a smile, does seem to me to be quite in keeping with our pre-theoretic conception of what events can be like. Non-action – the not performing of an action, the not engaging in a process that one might have performed or engaged in on a given occasion – often deserves to be seen as a form of action too. The decision to sit still or to keep your mouth shut often needs to be actively sustained, and to be sustained actively for as long as you stick to it, and the effort that takes is often greater than that of getting up, or speaking your mind.

There is also another consideration that enters into the question what aspectual status should be assigned to the results of applying negation to eventuality descriptions. This is how the eventualities described by the negations of eventuality descriptions are temporally located (by tenses and temporal adverbs). Recall the discussion of the temporal location of quantification states in Section 3.9.1. There we noted that these states must be located by tense and temporal aspect via temporal coincidence and not by one-directional temporal inclusion. This same principle – location via temporal coincidence – also applies to the eventualities that are described by the result of negation. A good example to illustrate the point is the aleady mentioned (3.133), repeated below.

(3.133) I didn't turn off the stove.

As Partee noted in relation to this example, the utterance of it that she considered was about some particular time in the past (the time when the speaker and her partner were closing up the house before setting off on their trip). But as we observed in our discussion of this example, what this means in more detail is that there was no event of the speaker turning off the stove within the period of time when the two were closing up. So if we want to represent the not turning off the stove that the sentence talks about as an eventuality, then its duration should be identified with its location time (provided by tense or temporal adverb). In other words, just as with the quantification states of Section 3.9.1, temporal location has to take the form of the Condition '$dur(ev) = t$'.

Once this principle for determining the temporal location of a negation eventuality has been adopted, then the question whether the eventualities that are to be located in this way should be states or events can be resolved on the basis of other considerations. In this regard we find ourselves in the same situation that we found ourselves in when discussing the quantification states introduced by nominal quantifiers in Secti0on 3.9.1. But the difference is that we now have the other considerations that may determine whether the eventuality at issue i a state or an event. Our observations in connection with (3.135) indicate that the eventualities described by the negations of eventuality descriptions must be of the same kind as the eventualities described by the non-negated phrase (i.e. by the input to the negation operator). Thus the eventuality described by the negation of (3.135.c) should be an event and the one described by the negation of (3.135.d) a state.

3.9.4.1 Negation and Quantifiers

Besides these questions concerning the interaction between negation, aspect and tense there are also the interactions, much more widely documented in the semantics literature, between negation and the classical prototypes of scope-taking sentence constituents: the quantifying expressions, with quantifying DPs as their most prominent representatives. What follows are just a few observations on this topic, centered on the examples in (3.138).

(3.138)a. Mary didn't know every philosopher in the room.

b. Mary didn't know any philosopher in the room.

c. Mary didn't know some philosopher in the room.

d. Mary didn't know a philosopher in the room.

e. Mary didn't give a passing grade to every student in the class.

f. Every philosopher in the room didn't know Mary.

g. Every philosopher isn't a charlatan.

h. Not every philosopher is a charlatan.

Before we engage upon a brief discussion of these examples, let me be plain about its purpose. The few remarks that follow are intended to do no more than point towards some aspects of the interactions between negation and quantifiers, as constituting a range of phenomena that deserve closer scrutiny, and that actually have received closer scrutiny in settings different from the one we have adopted in these notes. These accounts could no doubt be adapted to the special features of our framework. But this is a task that must be left for some other occasion.

A first point to observe in connection with the examples in (3.138) is that our decision where to locate negations in the syntactic structures of negated sentences – viz. at a level between the Perf level and the T level – is supported by at least some of these examples. Compare in particular sentences (3.138.a) and (3.138.f). According to the syntactic proposal we have assumed for negated sentences the subject DP in (3.138.f) is outside the syntactic scope of negation. And if a semantic representation is constructed from this syntactic structure, then the subject will also be outside the scope of negation in this representation and the negation thus within the scope of the universal quantifier that the subject DP contributes. In other words, the reading that the sentence is assigned this way is to the effect that for every x such that x is a philosopher in the room it is not the case that x knew Mary. In contrast, when the same *every*-DP occurs in direct object position, as it does in (3.138.a), then it will be within the scope of negation. That too seems to be in agreement with intuition, according to which (3.138.a) has the reading that it is not the case that Mary knew every philosopher in the room.

But we need to be careful with the conclusions we draw from argumentations of this sort. One consideration we ignored in the argument above has to do with the possibility of Quantifier Raising. If we assume that the object DP of (3.138.a) is QR'ed out of its direct object position in (3.138.a), we arrive at a prediction for the truth conditions of (3.138.a) according to which the quantifier *every* has scope over the negation, which is precisely not what our intuitions about this sentence seem to be telling us. And there is also a problem of a different kind. What I said about intuitions concerning the meaning of (3.138.f) wasn't quite right. This sentence not only has the interpretation I claimed it has, but also one according to which its universal quantifier is within the scope of its negation. This alternative interpretation becomes especially prominent when the sentence is spoken and focal stress is put on the word *every*. This effect is arguably even stronger for the next sentence in the list, (3.138.g).

The moral of this last remark is that information structure also has a part to play in determining the scope relations between negations and quantifiers. But there are other complicating factors as well. One such factor is made visible by the four sentences (3.138.a-d). The difference between (3.138.a) and (3.138.b) is intuitively clear. The truth conditions of (3.138.a) are those of a sentence that contains a universal quantifier in the scope of a negation. (3.138.b) differs in that its truth conditions are those of a universal quantifier that scopes over negation. Accounting for these truth conditions within the setting we have adopted is not without its problems. In the LF for (3.138.a) the raised DP *every philosopher in the room* will have scope over the negation. True, the trace of the DP is within the scope of negation. But what justifies in assuming that it is the trace rather than the raised DP which has left it behind that is decisive for the scope relation to negation? And assuming that such a justification can be given, how can we compute a representation that captures the intuitive truth conditions of (3.138.a) from such a logical form?

An account of the intuitive truth conditions of (3.138.b) involves very different considerations, the likes of which we have not so far encountered in the Notes. They have to do with the fact that *any* is a Negative Polarity Item: it may only occur in 'negative environments'. What the 'negative environments' are that NPIs need has to be spelled out carefully, and it has to be spelled out differently for different NPIs. But two points are beyond controversy: (i) explicit negations create negative environments; and (ii) when in (3.138.b) the negation is removed, then the resulting environment is not a negative environment. This is why (3.138.b) becomes ungrammatical (or at the very least very marked), when the negation is removed – —*any* is no longer licensed– and also why (3.138.b) as it is given only has the one reading it has: The only way to construe (3.138.b) as a grammatical sentence is to construe *any* as in the scope of *not*, for only in this way can the string be justified as a grammatical sentence. The semantics that this construal places upon (3.138.b) is then the result of the semantics of *any*, which is that of a kind of emphatic existential. The semantics of (3.138.b) is that of a sentence with an existential quantifier within the scope of negation, which is equivalent to a universal quantifier that has the negation within its scope.

Some is in a way the opposite of *any*. Its semantics is also that of an existential. But it differs from *any* in that it avoids negative environments. For

(3.138.c) this means that *some* has to be construed as *not* within the scope of the negation. As a result the sentence is assigned a reading according to which there was some philosopher in the room that Mary didn't know. ((3.138.d) sounds a little strange. For me personally the sentence seems ambiguous. The indefinite article *a* seems ambiguous between an existential that can take wide scope over negation, just as *some* can, which yields a reading like that of (3.138.b). On the other hand *a* also seems to allow for a kind of emphatic negative polarity reading in this sentence, much like *any*. On this reading the sentence gets the same truth conditions as (3.138.a).)

Judgments seem to vary with regard to whether (3.138.a) can be given the 'strong' reading, according to which every philosopher in the room was such that Mary didn't know her or him, besides the 'weak' reading, according to which not every philosopher in the room was such that Mary knew him. An explanation for why some speakers get the strong reading may be that they can read the negation as applying directly to the verb, so that *not know* becomes a kind of compound verb that functions like a quasi-lexical unit (like the French verb *ignorer*). A similar explanation seems to apply to (3.138.e), even though here the 'quasi-lexical verb' is a bigger constituent, which includes the indefinite direct object *a passing grade* as well as the verb.

On the other hand, scope ambiguities have long been attested for sentences with *every*-DPs in subject position and it would appear that for these sentences the ambiguity is more prominent: both readings are quite easy to get. An example is (3.138.f). But here too there seems to be an asymmetry between the two readings. The reading in which the negation has scope over the *every*-phrase is one that requires a special, marked intonation, with a strong accent on *every* and rising prosody at the end of the sentence. Curiously this 'not-for-all' reading seems to be the only readily available reading for (3.138.g) – I have no explanation for why this should be.

These somewhat haphazardly selected examples and observations are meant to provide a glimpse of the many factors that shape the truth conditions of sentences that contain negations and quantifiers. But note that none of the factors that have just been discussed have anything to do with the central concern of this part of the notes – what roles that are played in sentence and discourse interpretation by tense and aspect. A treatment of the interactions between negations and quantifiers within the framework we are developing, in which all verbal predications are treated as eventuality descriptions, is an even more complex undertaking than the analyses of quantifier-negation interactions in which tense and aspect are ignored. Such a treatment will of course be needed eventually, since in actual speech and writing interactions between quantifiers and negations are found together with all tenses (and not just in conjunction with the Simple Present tense). To my knowledge no such treatment exists at the present time. [**check Champollion!**] In any case, it is a task beyond what we can accomplish in these notes.

3.9.4.2 DRS Constructions for negated Sentences

The two sentences for which we will provide DRS constructions are both very simple. One of them is (3.132.a), repeated here as (3.139.a), and the other, repeated here as (3.139.b), is the negated present perfect sentence that was shown with its syntactic structure in (3.134).

(3.139)a. Johnny didn't cry.

b. Fred hasn't paid the rent.

We start with the DRS construction for (3.139.a). In the light of what we have been saying about the syntax we adopt for negated sentences it should be clear that the LF for this sentence can be given as in (3.140). (As announced earlier, we have in (3.140.a), where the value of the Perf feature is -perf and that of the Prog feature is -prog, suppressed both Perf and Prog projection levels.)

(3.140)

a.

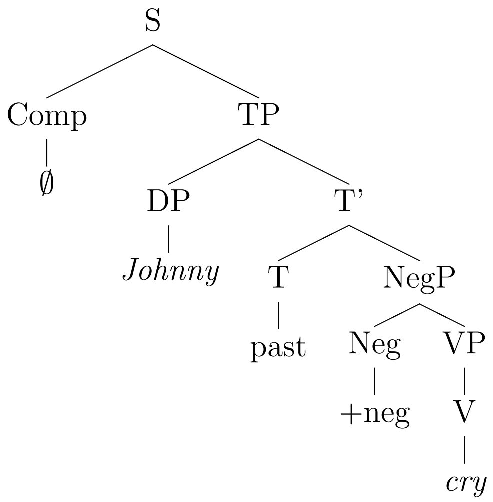

b.

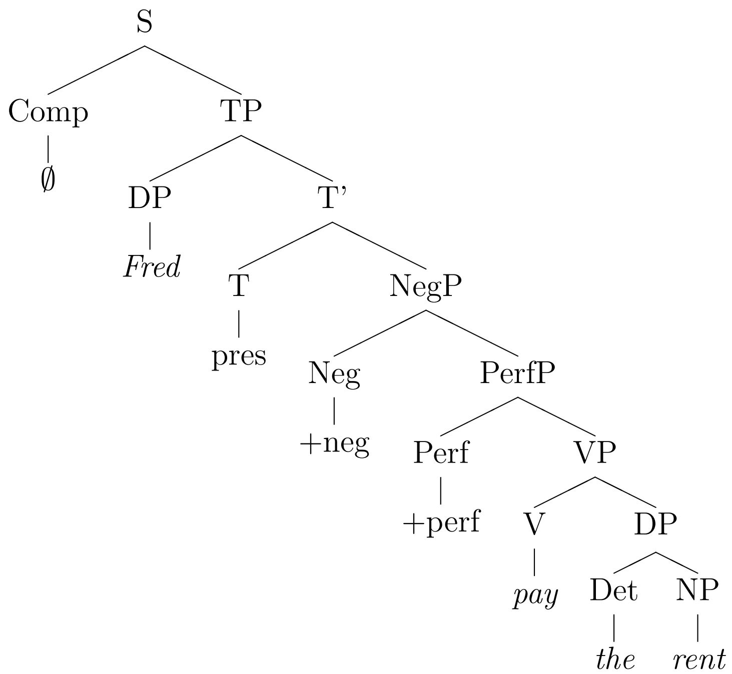

The construction steps needed to reach the representation for the VP of (3.140.a) are familiar and we start with the VP representation in place, as shown in (3.141).

(3.141)

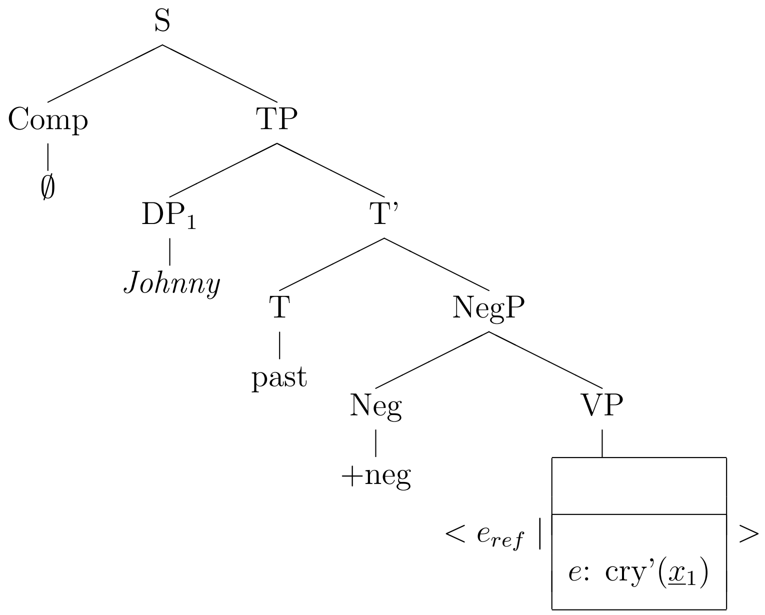

(3.142) presents the result of applying negation to the VP representation in (3.141). As, noted, temporal location is now not as much of a problem as it is with the temporal location of quantification eventualities, where the syntactic structures we have adopted dictate that a bottom up construction execute some kind of temporal location before it deals with the quantification. In the present construction the construction algorithm will have dealt with negation before it gets to temporal location, so there should be no need for the revisionary operations that we found ourselves condemned to in Section 3.9.1. But one part of the complications that we first noted in Section 3.9.1 remains: the eventuality descried by the result of applying negation will have to be located by way of temporal equality. In fact, this complication affects two aspects of representation construction. On the one hand the result of applying negation (or for that matter of constructing the contribution of a quantifier) must be recognizable as the description of an eventuality that is to be located like a 'quantification eventuality'. On the other the lexical entries of the tenses must be prepared for locating quantification eventualities in the right way. This is a new contingency with which the entries for the tenses we have given are not equipped to deal with. Recall for instance, the entry we gave for the Simple Past tense. This entry offers two location options, one for when the input is an event description and one for when it is a state description. What it needs in addition in order to deal correctly with the case before us is a third option that deals with quantification eventualities.

To make this possible we have to make new provisions both on the side of the operator and on that of the operandum. The operandum needs to make clear whether or not the eventuality it describes is a quantification eventuality. This information has to be marked on the output of any operation that produces descriptions of such eventualities, and marking the description as such should obviously be the task of the operator that produces such outputs. We need to decide on some way of marking the distinction. About the simplest way to do this, and one consistent with ways of marking distinctions that we have been employing so far, is to add the information in the form of a subscript on the described eventuality. We use the subscript $_{quan}$ for quantification eventualities. Absence of this subscript on the referential argument of an eventuality description indicates that what is described is not a quantification eventuality.

With this new convention in place it is now more or less clear how our entries for the past and the future tense should be modified. They should now offer a 3-place $\overset{!}{\vee}$ disjunction, in which the first two disjuncts are as before and a third disjunct which is activated if and only if the referential argument of the input representation bears the subscript $_{quan}$. This third disjunct specifies that the duration of the referential argument of the input coincides with the time that is introduced by the locating constituent These specifications doe not fix the actual formulation of the entry in every detail. But I assume that this can be left to the individual reader.

With these provisions in place, the result of applying the negation operator to the VP representation of (3.141) holds no further surprises.

(3.142)

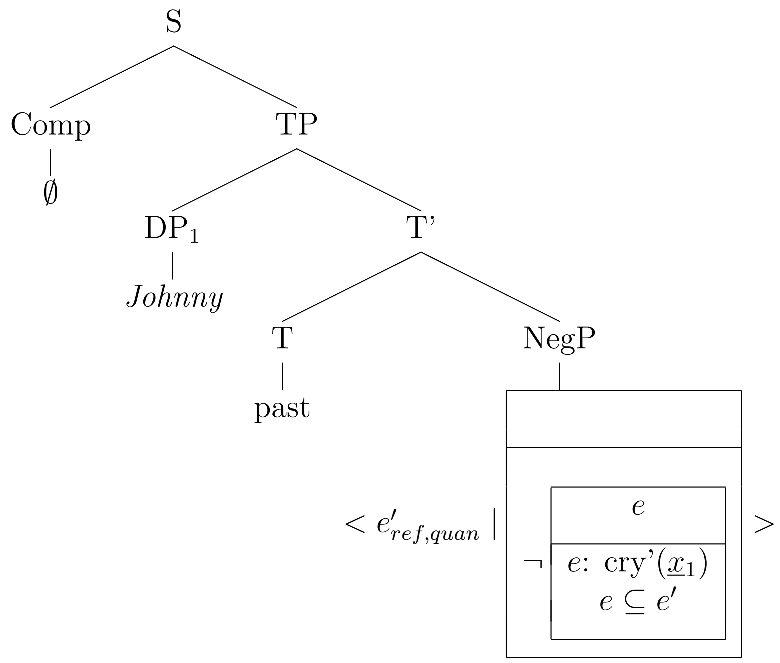

In the next step the NegP representation is combined with the semantics of the feature past. This involves application of the revised entry for this feature. The presence of the subscript $_{quan}$ on the event dref e' of the input selects the third disjunct and the result is as in (3.143).

(3.143)

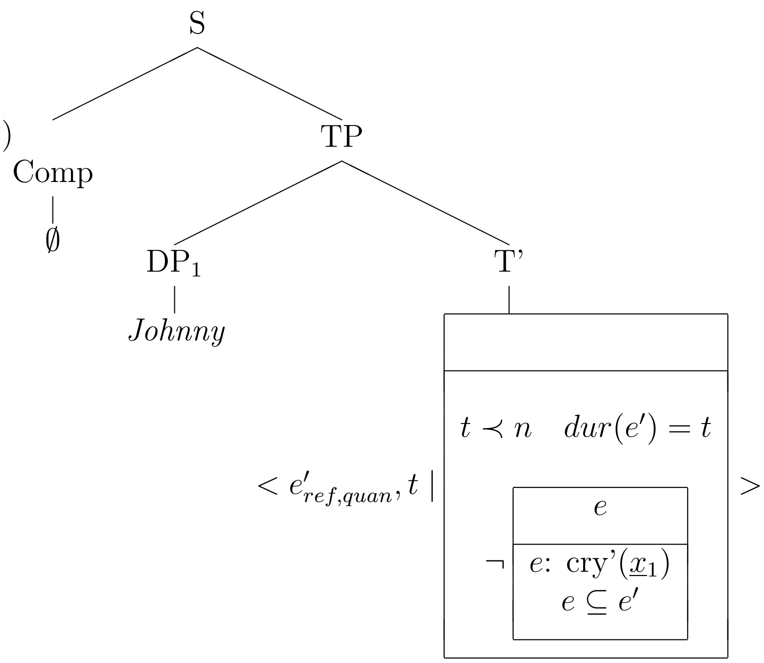

The remaining steps are familiar. (3.144) gives the final result.

(3.144)
$$
\begin{array}{|c|}
\hline
e' \quad t \quad j \\
\hline
t \prec n \quad dur(e') = t \quad \text{Johnny'}(j) \\
\neg \begin{array}{|c|}
\hline
e \\
\hline
e\text{: cry'}(j) \\
e \subseteq e' \\
\hline
\end{array} \\
\hline
\end{array}
$$

The DRS construction for (3.139.b) proceeds in much the same way as that for (3.140.a). But there is one important difference, which has to do with present tense uses of descriptions of states that are naturally understood as holding over longer periods of time, which can extend arbitrarily far beyond the utterance time in either direction. This could lead to problems in those cases where a president tense sentence that describes a quantification state contains a temporal adverb that denotes a time which clearly extends beyond the time of utterance. The following sentence is an example.

(3.145) This year every day of the week coincides with the first of a month.

(As a matter of fact this is true of every calendar year, but that is irrelevant to what the sentence is meant to illustrate.)

The most prominent interpretation of this sentence – when it is offered out of the blue like it is here – is that according to which it makes a statement about the year in which it is uttered: It says that each of the seven days of the week – Sunday, Monday, etc – is the day of the week of the first of one of the twelve months that make up this year. To get the right semantic representation for this reading of the sentence the state description that semantically represents the VP *coincide with the first of a month* must be temporally located (a) as holding at n (in virtue of the present tense) and (b) as being true of the time denoted by *this year*. The technical complication that this might seem to represent is the result of two assumptions we have made: (i) to treat the contribution of the (standard use of the) present tense as one which identifies the location time t with the utterance time n and (ii) to treat quantification eventualities as temporally coinciding with their location times. For (3.145) this would mean that the duration of the state described by the VP must on the one hand coincide with the utterance time and on the other with the denotation of *this year*. Evidently that is impossible unless the utterance time is assumed to extend over the entire calendar year within which the utterance is made. We haven't said anything that definitively rules this possibility out. But it is clearly a very counterintuitive one and it is made extremely implausible by variants of (3.145) in which the sentence is followed by something like *and this month there are five Sundays*. Are we to assume that between the first and the second conjunct of this enlarged sentence the utterance time is reconceived – from lasting the entire year to lasting just the current month? If more reflection or more examples are needed to convince you that this won't work will in general, those are left to you.

I take the more plausible conclusion from these considerations to be that the contribution that is made by the present tense to the location of quantification eventualities is different from the contributions made by temporal adverbs. For one thing, the present tense imposes the selection restriction that its input representation must be a state description. This is just as before and now also applies to inputs that are descriptions of quantification eventualities. Furthermore, all that the present tense does is to impose on a state description that it receives as input that the state it describes temporally includes (the time t that is identified with) n – this applies to quantification states just as it applies to any other states. A temporal adverb on the other hand locates quantification states as having durations that coincide with the time it contributes. Here too things remain as they were: for all quantification eventualities location by a temporal adverb takes the form of temporal coincidence, irrespective of whether the quantification eventuality is an event or a state and irrespective of tense. In short, everything we have said up to this point about the president tense and about adverbial quantification can remain as it is – luckily, you might say. but note that things work out because the entry for the present tense as we have it pays no attention to the subscript $_{quan}$, while in adverbial location the presence of $_{quan}$ overrules the distinction between events and states.

We show the construction of the semantic representation for sentence (3.139.b), from the LF in (3.140.b), starting at the point where the representation for the PerfP node is in place:

(3.146)

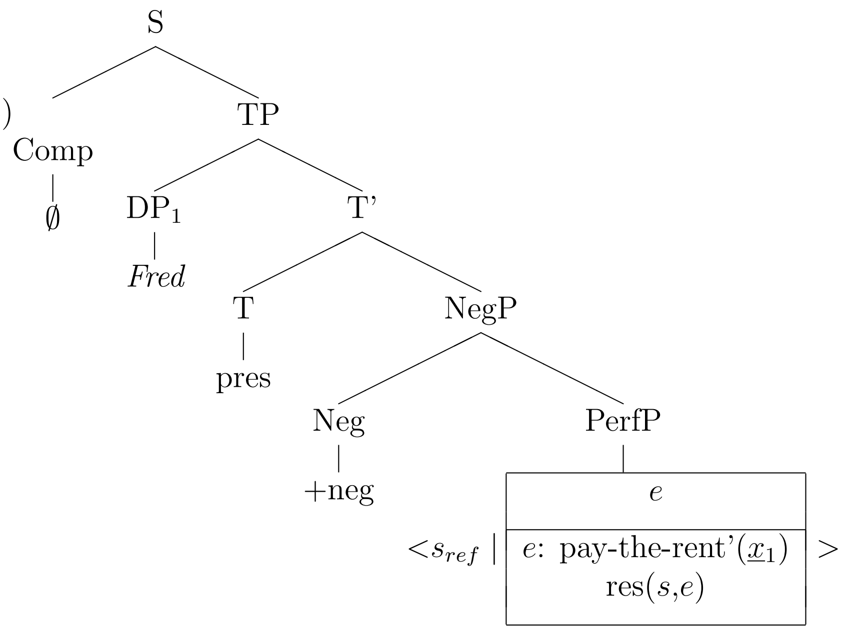

The next step in the construction forms the representation of the negation of this description.

(3.147)

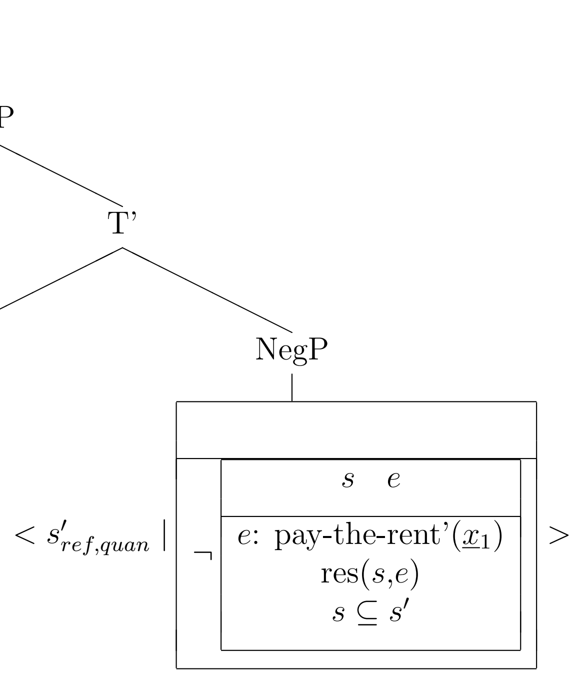

The next step is the result of combining the NegP representation in (3.147) with the feature pres. Nothing happens here that is formally different form what we have seen. but the point is that this is because press ignores the subscript $_{quan}$. The result of the combination is shown in (3.148).

(3.148)

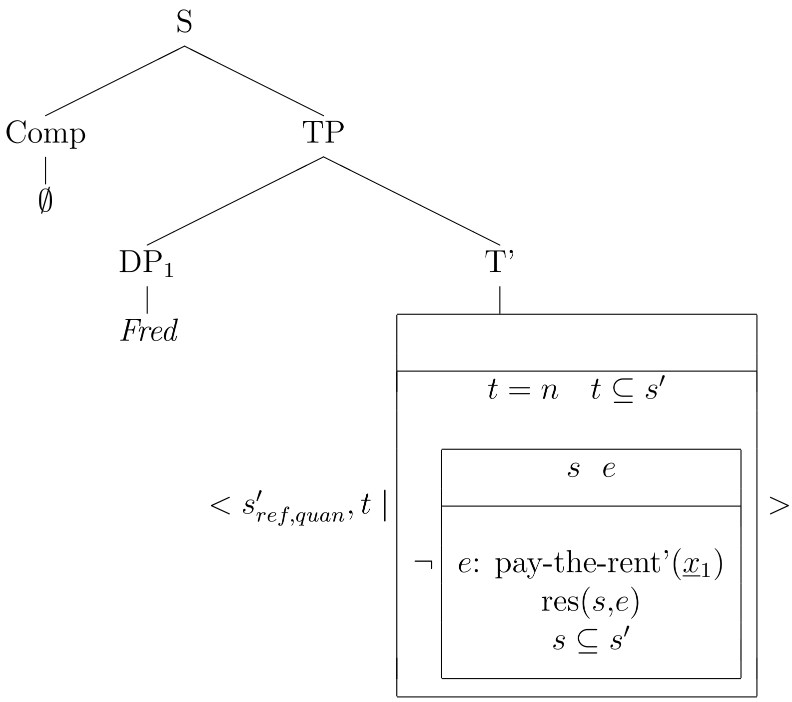

The remaining two steps are familiar. They lead to the DRS in (3.149).

(3.149)
$$\begin{array}{|c|}\hline s' \;\; t \;\; f \\ \hline t = n \quad t \subseteq s' \quad \text{Fred'}(f) \\ \\ \neg \begin{array}{|c|}\hline s \;\; e \\ \hline \\ e\text{: pay-the-rent'}(f) \\ \text{res}(s,e) \\ s \subseteq s' \\ \\ \hline \end{array} \\ \hline \end{array}$$

<u>Exercises</u>

1. Complete the DRS construction for (3.140.b).

2. Construct the DRS for the following sentences:

(i) Fred didn't climb Everest.
(ii) Fred won't climb Everest.
(iii) Fred hasn't been climbing Everest.

3. Construct a DRS for sentence (3.145)
To conclude this section about negation we give the lexical entry for the feature value +neg. (This is the entry that we already been applying in the two DRS constructions shown above.) The semantics of +neg is specified using the same format that we have also used to formulate the entry for the feature 'past' in (3.22) (as well as for other 1-place operators).

(3.150) (lexical entry for the feature +neg)

+neg

Sel. Restr: eventuality description

Sem.Repr: $<ev_{ref}, ... \mid K> \rightsquigarrow$

$$<ev'_{ref,quan} \mid \boxed{\neg \boxed{< ... \mid (K \bigcup \begin{array}{|c|} \hline ev \\ \hline ev \subseteq ev' \\ \hline \end{array}) >}} >$$

Input-output constraint:
if ev is an event, then ev' is an event;
if ev is a state then ev' is a state.[45]

3.9.5 The 2-Place Operators of Standard Propositional Logic

All of the standard binary connectives of classical propositional logic – ¬, &, ∨, → and ↔ – need revisiting now that tense and aspect have been made an integral part of our semantic analysis. So far we have only dealt with negation. In this subsection we have very brief looks at the remaining ones. We start with →, since conditionals have played an important part

[45]Note that if the store of the input representation contains other discourse referents besides the referential argument $ev_{ref,quan}$, these other drefs are placed in a new store that immediately precedes the newly formed DRS $K \bigcup \begin{array}{|c|} \hline ev \\ \hline ev \subseteq ev' \\ \hline \end{array}$. Although this complication does not arise in the two examples above, where the referential argument ev is the only dref in the store of the input representation to the negation operator, it does arise often enough. (We will see an example of this presently, when we construct the DRS for sentence (3.151) below.) The reason for this way of handling the additional drefs of the input store can be fully explained only on the basis of the extensions of our approach that will be presented in Section 4.

in our discussion of donkey phenomena as one of the main motivations for developing the original version of DRT. On &, $\leftrightarrow$ and $\vee$ we will be very brief.

3.9.5.1 Conditionals

We focus on just one way of expressing conditionals in English, in which a main clause combines with a entence-initial *if*-clause. And we will be looking at just one example of this form.

(3.151) If Pedro bought a donkey, he didn't buy a mule.

The LF we assume for this sentence is similar to the one adopted in PART I for the conditional donkey sentence 'If Pedro owns a donkey, he beats it.' in that the *if*-clause is left-adjoined to the main clause. A difference, however, is that we now take the *if*-clause to be adjoined to TP, for reasons that will become clear when we construct the semantics for this LF. (A further difference, which is imposed by commitments already made in PART II, is that both *if*-clause and main clause now each have their additional projection levels between VP and S, viz. AspP, PerfP, NegP and TP. In the presentation of the LF for (3.151), in (3.152) below, most of these levels have – consistently with our general practice – been suppressed, since their feature values are semantically otiose in that they dictate the mere passing up of the representations of their sister nodes to their mother nodes.) In general, however, the projection levels between TP and VP (or between NegP and VP, as the case may be) are important,because it is there that much of the aspectual information is located that determines how *if*-clause and main clause are temporally and by implication also rhetorically related. Such interactions will play no active part in the present section, in which we will make do with the LF for (3.151) shown in (3.152). Nevertheless we display, in (**????**.a,b), the full LFs for the *if*-clause and the main clause of (3.151), as a way of going on record for what information may be available about either clause that will have to be coordinated with such information provided by the other clause.

(3.152)

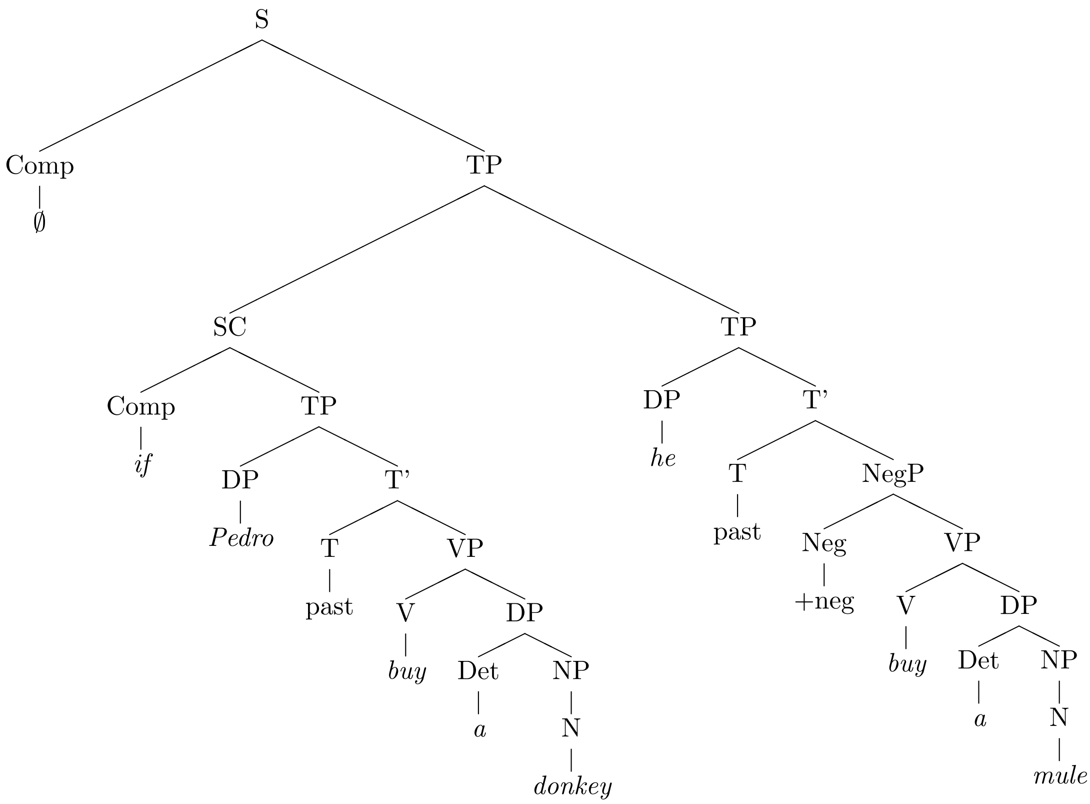

The representations of the TP of the *if*-clause of (3.152) and of the lower TP of its main clause are constructed by applying construction principles whose applications we have seen. We display these representations without further ado; see (3.153).

(3.153)

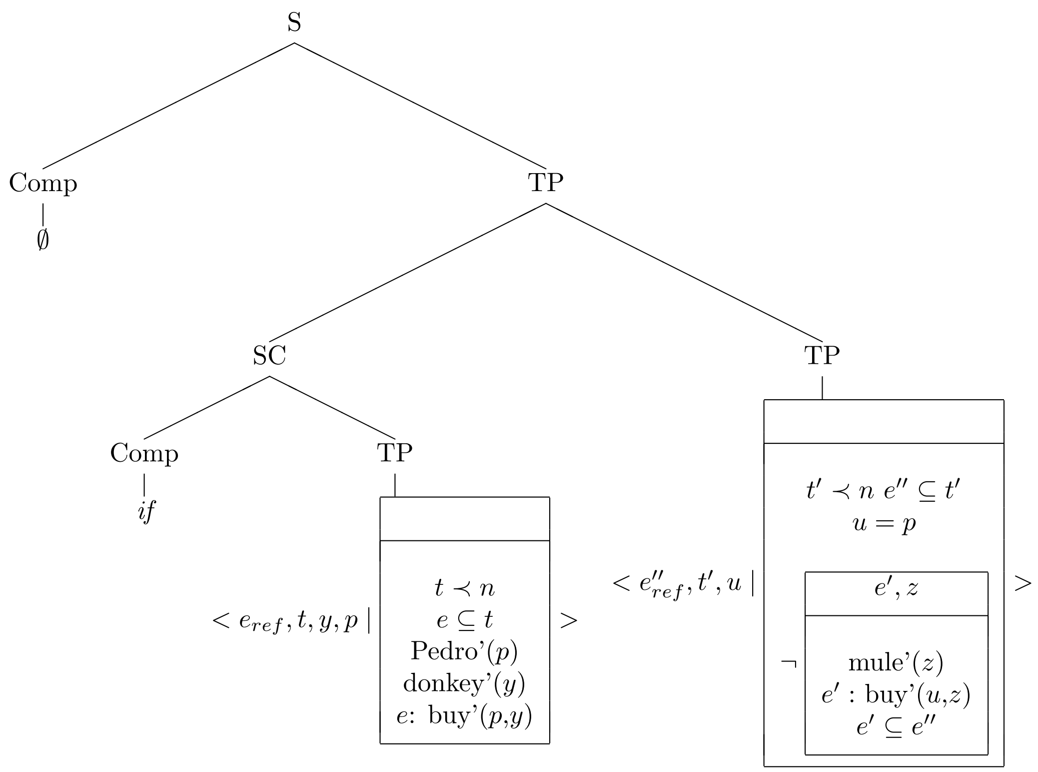

(It has already been noted that as things stand we don't have a way of dealing with anaphoric pronouns. So the treatment that *he* has received in the DRS construction for the lower main clause TP in (3.153) is to be regarded as provisional – a kind of rain check on the treatment of proper names, pronouns and other definite DPs that will be developed in Section 4.)

The next construction step that has to be carried out in the continuation of the DRS construction in (3.153) is triggered by the word *if* that fills the Comp position of the SC. *if* tells us that the representation of the subordinate clause which it heads and that of the main clause to which the SC is adjoined are to be combined as antecedent and consequent of a conditional DRS Condition and that this DRS Condition has to be placed into the Condition Set of a new DRS that is otherwise empty. As in the treatment of quantification in Section 3.9.1 we keep the stores of antecedent and consequent 'local', rather than merging them into a single store for the conditional DRS Condition as a whole. (As far as store management goes, the only difference with the representations of nominal quantifications is that the representation of a conditional has no explicit 'variable binder' like the central diamond in a duplex condition, so there is no dref that this binder selects as 'bindee'. A consequence of this is that in the formation of a conditional-representation no dref gets moved from the store of either one of the component representations to some DRS Universe.)

(3.154)

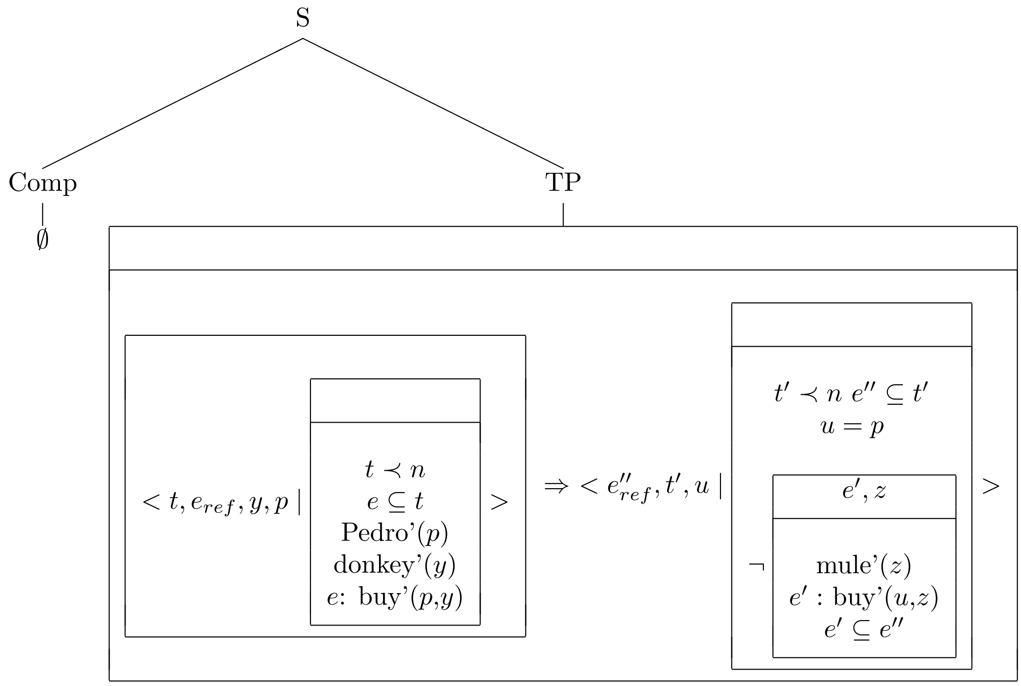

The remaining step, triggered by the main clause Comp, is the familiar one that deals with the drefs that are still in stores. (As with the quantificational sentences discussed in Section 3.9.1, the representation in (3.154) has more than one store. So the instruction about what is to be done with the remaining store elements has to specify the placement options for the drefs in each of these stores. We noted earlier that these options are limited to (i) the Universe of the DRS K following the store in question, and (ii) the Universe of any 'higher' DRS of which K is an immediate or indirect sub-DRS. But in some cases that still leaves a lot of room for choice.)

The question which drefs should go where in those case where this is scope for choice is once more one that can be addressed properly only in the setting we will develop in Section 4. Needed is the presuppositional treatment of the different types of definite DPs. Incorporation of these presuppositional treatments of definites into the framework will in particular yield the conclusion that the discourse referent p for *Pedro* becomes part of the Universe of the main DRS. And this treatment will also lead in a natural away to the placement of the Condition 'Pedro'(p)' in the Condition Set of the main DRS. There will moving it to the main Condition Set where it ought to be. Here we will jus t show the final result without trying to make a further case for the operations that are needed to go from (3.154) to it.

The transfer of the dref y introduced in the processing of the indefinite DP *a donkey* is also a matter for which no satisfactory account is possible at this point. One option for the drefs introduced by indefinites is adding them to the Universe of the DRS to the right of their store. That is the Universe that y should go to if the resulting DRS is to capture the truth conditions of donkey sentences that we have been concerned with from the time that they made their appearance in these notes. But drefs introduced by indefinites may also be transferred to higher Universes, in those case where an indefinite is used 'specifically', in one of the several senses in which the term 'specific indefinite' has been used in the literature. Part of the literature on indefinites has suggested that their specific uses are marked and that the default use is the one that in our set-up is captured by moving the representing dref from its store to the DRS to its right. Though I am not convinced that this perspective is completely correct, I propose that for present purposes we do as if it were. This makes the rule that drefs for indefinites are to be moved from their store to the Universe of the DRS following the store into a kind of default rule. We will from no on proceed in accordance with this assumption: In the absence of contrary evidence this is where the drefs for indefinites go when they are moved from the stores in which they occur.

With these assumptions about the store-to-DRS transfers for the drefs in (3.154) (and the ad hoc move of lifting the Condition 'Pedro'(p)' into th Condition Set of the main DRS) the last step turns this representation into the DRS in (3.155).

(3.155)
$$\begin{array}{|c|}\hline p \\ \hline \text{Pedro'}(p) \\ \begin{array}{|c|}\hline t\ \ e\ \ y \\ \hline t \prec n \\ e \subseteq t \\ \text{donkey'}(y) \\ e\text{: buy'}(p,y) \\ \hline \end{array} \Rightarrow \begin{array}{|c|}\hline t'\ \ e''\ \ u \\ \hline t' \prec n\ dur(e'') = t' \\ u = p \\ \neg \begin{array}{|c|}\hline e'\ z \\ \hline \text{mule'}(z) \\ e' : \text{buy'}(u,z) \\ e' \subseteq e'' \\ \hline \end{array} \\ \hline \end{array} \\ \hline \end{array}$$

For good measure we add a lexical entry for the use of *if* as indicator of an antecedent-consequent relation between the clause it heads and the clause to which that clause is adjoined. Like the entry for *every* in (3.106), the entry for *if* has two input representations, the first for the representation of the *if*-clause and the second for the adjunction site of the *if*-clause.

One glaring absence in this representation of (3.151) is that of any information about the temporal relation between the events e and e''. Intuitively speaking, (3.151) can be used to express two clearly distance propositions. The first says that if there was a purchase of a donkey then there wasn't a purchase of a mule: two different types of events such that if an event of the first type occurred, then an event of the second type did not. But the sentence can also used for a second purpose – that of expressing, in a kind of roundabout way, that donkeys are different from mules: if an event e was an event of Pedro buying a donkey then it wasn't event of Pedro buying a mule for the simple reason that a donkey isn't a mule. On this second interpretation the temporal relation between the event of the main clause and that of the subordinate clause is clear: they are simultaneous because they are one and the same event. But also on the first interpetstion one feels that some kid of temporal closeness between main clause event and *if*-clause event must be somehow part of way the sentence is trying to say - some kind of temporal proximity seems to be involved. But it is hard to pin what that is, or what is responsible for it. Certainly these matters cannot be dealt with in any adequate way with the tolls we have.

We conclude this subsection on *if*-clause-main clause conditionals with a lexical entry for the complementizer *if*. In the schematic representation of the semantic operation expressed by *if* the DRS K is the semantic representation of the *if*-clause and K' the semantic representation of the clause to which it is adjoined.

(3.156) (lexical entry for the Subordinate Clause Complementizer *if*)

if (SC Complementizer)

Sel. Restr: —

Sem.Repr:

$$<\{..\}_{ant} \mid K> \rightsquigarrow$$

$$\left(< \{..\}_{con} \mid K' > \rightsquigarrow \begin{array}{|c|} \hline \\ \hline < \{..\}_{ant} \mid \begin{array}{|c|}\hline \\ \hline K \\ \hline\end{array} > \Rightarrow < \{..\}_{con} \mid \begin{array}{|c|}\hline \\ \hline K' \\ \hline\end{array} > \\ \hline \end{array} \right)$$

3.9.5.2 Conjunctions

From the perspective of Discourse Representation Theory conjunction is a curiously ambivalent notion. On the one hand there is the representation formalism of DRT. Here conjunction is a structural concept: the conjunction of two or more DRS Conditions is represented by them belonging to one and the same DRS Condition Set; and the conjunction of two DRSs can be represented simply as their merge. No special complex DRS Condition type, of the sort needed to represent negations or conditionals is needed in this case.

Representing conjunction as joint set membership or as DRS merge has obvious consequences for the logic of conjunction: Conjunctions $\&(A_1,...,A_n)$ of n conjuncts are invariant under permutation: $\&(A_1,...,A_n)$ is true iff

&(A_{i_1},...,A_{i_n}) is true, where A_{i_1},...,A_{i_n} is any permutation of A_1,...,A_n. For binary conjunctions A & B this entails that A & B is logically equivalent to B & A and that (A & B) & C is logically equivalent A & (B & C); that is, the operation & is *commutative* and *associative*, as terminology has it.

On the other hand, DRT was originally designed to account for phenomena of pronominal and temporal anaphora. The examples of these we have seen so far were either donkey sentences (like in (3.157.b) or donkey discourses (as in (3.157.c) and many of the examples in Section 3.1). The original version of DRT makes the prediction that the pronoun *it* can be interpreted as anaphoric to the indefinite *a donkey* in (3.157.b) and (3.157.c), but not in (3.157.e) and (3.157.f). Three of these predictions are correct, but the one concerning (3.157.e) is problematic: In this sentence interpreting the pronoun as anaphoric to the indefinite doesn't seem out of the question. But the theory correctly captures the left-right asymmetry illustrated by the pair (3.157.c) - (3.157.f); and this asymmetry is mimicked perfectly by the pair (3.157.a) - (3.157.d): the pronoun cannot precede its indefinite antecedent. It is this asymmetry that renders the original DRT treatment inadequate. The representations that the theory uses in its account of anaphora should not represent conjunctions as the merges of the representations of their conjuncts.[46]

(3.157)a. Pedro owns a donkey and he keeps it in his backyard.

b. If Pedro owns a donkey, he keeps it in his backyard.

c. Pedro owns a donkey. He keeps it in his backyard.

d. He keeps it in his backyard and Pedro owns a donkey.

[46] The conditional sentences in (3.157) are the odd ones out. Here anaphora is not subject to the same left-right asymmetry that can be observed for conjunctions and sentence successions. No doubt the reason for this has to do with the logical asymmetry between *if*-clause (antecedent) and main clause (consequent), which is the same irrespective of whether the *if*-clause precedes the main clause or follows it. Nevertheless, pronominal anaphora is subject to left-right effects here too, but because of the non-alignment between surface left-right ordering and the antecedent-consequent asymmetry the acceptability patterns for conditionals turn out to be quite complicated and more complicated than predicted by original DRT. A further complication for anaphoric pronouns with indefinite antecedents is that sentences with indefinite arguments are sometimes interpretable as expressing generic quantifications in which the indefinite shares the scope of the generic quantifier. This possibility does not only arise for sentences that have the form of conditionals, but it is particularly common for them; and it adds a further complexity to the question when anaphoric relations between indefinites and pronouns in *if*-clause-main clause sentences are possible. For more examples and more discussion see Section 4.3.3.

e. If he keeps it in his backyard, Pedro owns a donkey.

f. He keeps it in his backyard. Pedro owns a donkey.

The left-right asymmetry exemplified by (3.157.a,d) is also found – unsurprisingly – in conjunctions of more than two conjuncts; and with such longer conjunctions we also find that bracketing is not arbitrary: 'A & B & C' should not be analyzed as A & (B & C) but as (A & B) & C. I leave it to the readers to find their own examples for this.

For the construction of DRSs as logical forms for conjunctive sentences like (3.157.a) these observations have the following implications. In the course of the construction we will have to keep the conjuncts of a conjunction separate, with their order – that in which the conjuncts appear in the sentence that is being interpreted – explicitly encoded. The order encoding will have to survive until the point has been reached where anaphoric interpretations of pronouns have taken place. Once anaphoric matters have been settled, however, there is no further need to keep the representations of the conjuncts apart from each other. So at that point they can be merged with each other and with the DRS that contains them.

The unfortunate part of the decision to say something about clausal conjunctions at this particular point of the development in these Notes of the bottom up construction method is that we are not yet in a position to deal with anaphora. So a convincing demonstration of why conjunctions have to be processed in the way just hinted at will have to wait. We will be able to return to it only in Section 4.[47] All we can do right now is to show what the semantic construction method for conjunctions will have to like, so that the mechanisms for anaphora resolution that will be discussed in Section 4 can be applied to conjunctions as well as to conditionals and to cases of cross-sentential anaphora.

The sentence we will use to illustrate the method is given in (3.158).

(3.158)Maria owns a horse and Pedro owns a donkey.

To construct a semantic representation for this sentence we will, as always, need a syntactic structure for it, from which the construction of the semantics can proceed. The syntactic structures for the conjuncts of (3.158) are fixed by what we have been doing so far. What we are still missing is a syntactic

[47]This part of the Notes is still to be added.

representation of the conjunction. But as far as this is concerned the Notes contain a kind of precedent. (2.82) gave the syntactic structure of the complex DP *Fred and Mary*. As we pointed out in the discussion surrounding this example, such conjunctive DPs, which denote the mereological sum of the denotations of their conjuncts, are semantically very different from the clausal conjunctions of which (3.158) is an example. But as far as syntax is concerned, there is no need to make a difference between these two ways in which conjunctions can be used. So we assume that the syntactic structure of (3.158) is as in (3.159).

(3.159)

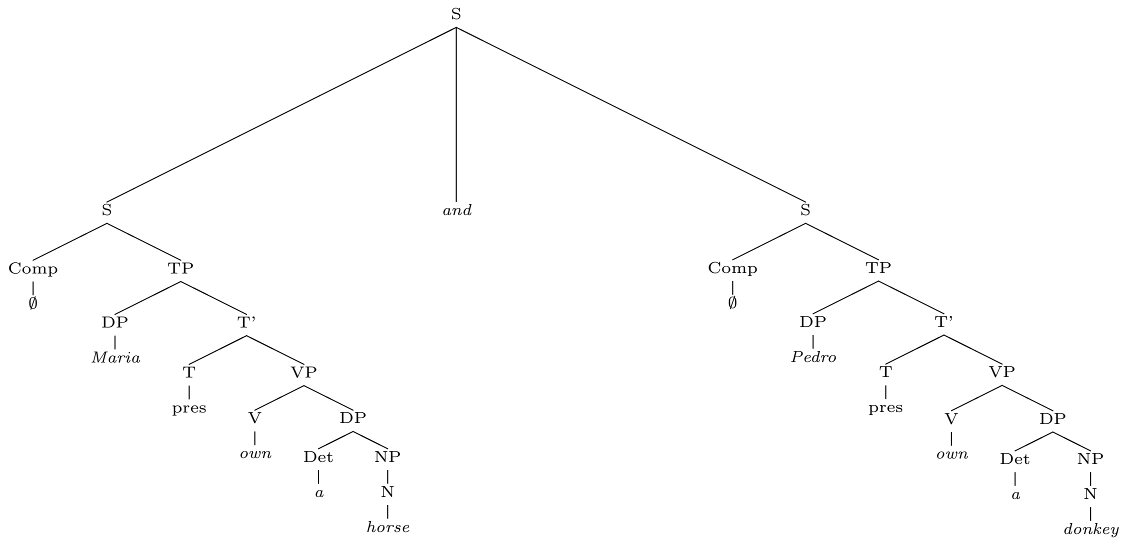

The semantics for (3.159) is also fixed by the assumptions adopted in Section 3, except for the interpretation of *and*. As regards the semantics of *and*: it is here that the special provisions have to be made which keep the conjuncts separate for as long as that is necessary. We will treat clausal conjunctions semantically as complex DRS Conditions, in which DRSs for the conjuncts are combined into a single Condition by means of an operator that captures the meaning of *and* and plays, at the level of DRS-syntax, the same role as the conditional operator $\Rightarrow$. We choose as symbol for this operator the semi-colon ';'.[48]

To see what this comes to, we first need to build the representations of the conjuncts of (3.158). These can be computed in the familiar way. (Compare

[48] Following the lead of Dynamic Semantics in the style of Groenendijk and Stokhof. See e.g. (Groenendijk & Stokhof 1990).

for instance the DRS construction for ‘Frieda closed the shop’ in Section 3.5.1.). These computations lead to the DRSs shown in (3.160.a,b).

(3.160)a.
$$\begin{array}{|c|}\hline t_1 \quad s_1 \quad x \quad y \\ \hline t_1 = n \\ t_1 \subseteq s_1 \\ \text{Maria'}(x) \\ \text{horse'}(z) \\ s_1\text{: own'}(x,y) \\ \hline \end{array}$$

b.
$$\begin{array}{|c|}\hline t_2 \quad s_2 \quad u \quad v \\ \hline t_2 = n \\ t_2 \subseteq s_2 \\ \text{Pedro'}(u) \\ \text{donkey'}(v) \\ s_2\text{: own'}(u,v) \\ \hline \end{array}$$

Inserting these DRSs under the S-nodes of the conjuncts in (3.159) leads to the structure in (3.161).

(3.161)

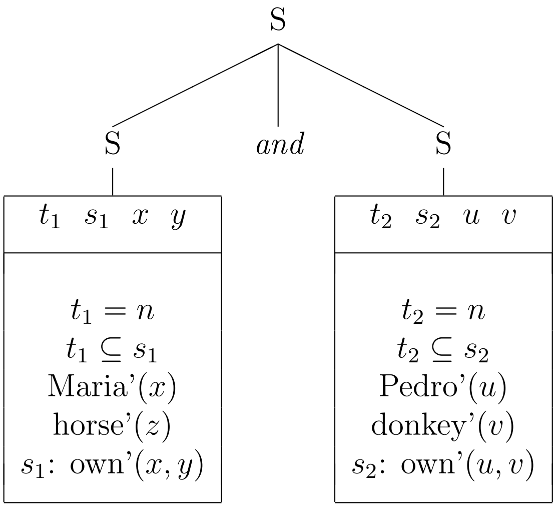

The processing rule for *and* that must be applied to (3.161) is reminiscent of the rule triggered by *if* that we assumed in our treatment of conditionals that take the form of main clauses modified by subordinate clauses whose Complementizer is *if* (see the part of the present section headed **'Conditionals'**). But there is one difference: the syntactic structure we assumed for *if*-conditionals was that of a modified TP node for the main clause. The structure assumed in (3.159) is one in which *and* coordinates to S-labeled structures. Given the way in which we have set up the syntax-semantics interface, this has repercussions at the level of the semantics in that in (3.161) all drefs have been transferred from their stores to DRS Universes. (Note well in this connection that as things have been set up there is for each conjunct only one DRS Universe to which transfer is possible.) This might raise the worry that in the end some drefs, those that have been introduced for the proper names *Maria* and *Pedro*, will not end up where they should. We will see that for the present example things will work out as wanted. But in other cases the treatment described here of conjunctions that involve proper names will lead to semantic representations that are wrong in that the representing drefs of proper names end up in the wrong place. This problem would be hard to fix with the tools that are available right now. But it will disappear when proper names are treated as presupposition triggers in Section 4 (see Section 4.3.1).

The construction principle that we adopt for *and* is by and large a variation of the rule for *if* formulated in (3.156). The DRSs K_1 and K_2 for the first and the second conjunct are combined into a DRS Condition of the form $K_1; K_2$ and this Condition is then made into the sole member of the Condition Set of a new DRS. The result of applying this principle to (3.161) is shown in (3.162).

(3.162)

$$\begin{array}{|c|}\hline \\ \hline \begin{array}{|c|}\hline t_1 \quad s_1 \quad x \quad y \\ \hline t_1 = n \\ t_1 \subseteq s_1 \\ \text{Maria'}(x) \\ \text{horse'}(z) \\ s_1\text{: own'}(x, y) \\ \hline \end{array} ; \begin{array}{|c|}\hline t_2 \quad s_2 \quad u \quad v \\ \hline t_2 = n \\ t_2 \subseteq s_2 \\ \text{Pedro'}(u) \\ \text{donkey'}(v) \\ s_2\text{: own'}(u, v) \\ \hline \end{array} \\ \hline \end{array}$$

It is at this point of the construction that anaphoric relations between the first and second conjunct would have to be resolved. For the example we are dealing with no such operations are needed, so we can proceed straight away to the next step, which simplifies the structure in (3.162) to the one in (3.163) below, in which K_1 and K_2 are merged with each other and with the DRS containing the ;-Condition of which they are the two constituents.

(3.163)

$$\begin{array}{|c|}\hline t_1 \quad s_1 \quad x \quad y \quad t_2 \quad s_2 \quad u \quad v \\ \hline \\ t_1 = n \quad t_1 \subseteq s_1 \quad t_2 = n \quad t_2 \subseteq s_2 \\ \text{Maria'}(x) \quad \text{horse'}(z) \quad \text{Pedro'}(u) \quad \text{donkey'}(v) \\ s_1\text{: own'}(x, y) \quad s_2\text{: own'}(u, v) \\ \hline \end{array}$$

Note: Following standard practice in DRT we should give explicit verification conditions for DRS Conditions of the form $K; K'$. These are straightforward: a function f verifies $K; K'$ in a model M iff there is an extension g of f that is also defined of the Universe of K and that verifies K in M and an extension h of g that is also defined on the Universe of K' and that verifies K' in M. These verification conditions for Conditions of the form $K; K'$ render DRSs like (3.162) logically equivalent to their simplifications exemplified by (3.163).

Exercise Formulate the rule for *and* that has been applied in the transition from (3.161) to (3.162) in the format of the *if*-rule given in (3.156).

We conclude this discussion of sentence-level conjunctions with two general observations. The first is implicit in the introductory paragraphs to this discussion of conjunction: in English (and natural languages more generally) binary conjunctions are a special case of conjunctions involving n conjuncts for, in principle, any n $\geq$ 2. The most common way in which conjunctions with $>$ 2 conjuncts are expressed in English is for *and* to occur between the last and one but last conjunct while all other conjuncts are separated by commas. For instance, a conjunction invoking three clauses S1, S2, S3 is most naturally expressed as 'S1, S2 and S3'. The semantic processing of such conjunctions should involve 'bracketing from left to right', in the following sense: The semantic representation of 'S1, S2 *and* S3' should yield a complex DRS Condition that can be loosely described as having the form $(K_1;K_2);K_3$, where K_1, K_2, K_3 are the semantic representations of S1, S2, S3. More precisely, an English sentence that has the form of a conjunction

'S1, S2 and S3' will at the relevant representation stage lead to a DRS of the form given in (3.164).

$$(3.164)\quad \begin{array}{|c|}\hline \quad \\ \hline \begin{array}{|c|}\hline \quad \\ \hline K_1;K_2 \\ \hline\end{array}\ ;K_3 \\ \hline\end{array}$$

The second observation is of an even more general nature and goes beyond the limits of these Notes, which deal with syntax-semantics interface questions fairly narrowly conceived. Natural languages are replete with constructions that combine two clauses into one clause that is true only if both these clauses are. Because of this truth-conditional property any such construction can be considered a 'conjunction' of sorts. But that doesn't prevent such constructions from playing a wide variety of semantic and pragmatic roles, and often those roles can be described as each other's opposite. Think for instance of the two constructions in English which consist of a main clause and either (i) a *because*-clause or (ii) an *although*-clause. *because*-clauses and *although*-clauses play opposite roles insofar as a *because*-clause gives a reason for the truth of the main clause, whereas an *although*-clause signals that the main clause is true in spite of the fact that the *although*-clause is true as well. A similar contrast can be observed between coordinations with *and so* and *but*, respectively: *and so* marks the first conjunct as a reason for the truth of the second conjunct, while *but* often implies that the truth of the second conjunct is unexpected in light of the truth of the first conjunct.

Such epistemic, rhetorical and other discourse relations between the conjuncts of 'conjunctions' (in the very general sense of 'conjunction' in which we are using the word in this and the last paragraph) and of corresponding relations between successive complete sentences that make up multi-sentence discourses and texts, have become a field of semantic and pragmatic research in its own right.[49] The only reason to point to this aspect of linguistic meaning here is that any theory which takes the rhetorical dimension of conjunctions seriously must have access to the separate semantic contents of the conjuncts. If such a theory is to be built on top of the syntax-semantics interface theory developed in these Notes, then the semantic structures made

[49] An important direction within this field is the SDRT ('Segmented Discourse Representation Theory') of Asher and Lascarides (see in particular (Asher & Lascarides 2003) and (Asher & Lascarides 2007)).

available by this foundation must make the semantic representations of the conjuncts available in some such form as (3.162) and not throw these contents together in the way they were in (3.163). Furthermore, in a theory that takes these rhetorical dimensions of meaning into account, the semantic representations of the conjuncts will remain separate because they will play their own separate roles as arguments to the rhetorical and other discourse-structural predicates.

3.9.5.3 Disjunctions

In classical logic conjunction and disjunction appear as very similar: they differ in their truth conditions – A & B is true iff both A and B are true, A $\vee$ B is true iff at least one of A and B is true – but otherwise they are on a par: binary sentence operators each of which is both commutative and associative. From the perspective of DRT-based natural language semantics, however, conjunction and disjunction are quite different. In our discussion of conjunctions we noted that in DRT-based representation languages conjunction is structurally representable, as DRS merge or, in the case of DRS Conditions, as shared membership in a DRS Condition Set. But nevertheless, we saw, the conjuncts of natural language conjunctions have to be represented separately because of anaphoric and other semantic or pragmatic relations between them. Disjunctions are different on both counts. First, within DRT disjunction is not a structural operation. Like the conditional its representation requires its own type of complex DRS Condition, one which combines two DRSs into a single Condition. For the operator of this DRS Condition it has become standard to use the same symbol that has been traditionally employed to denote disjunction in classical logic, viz. '$\vee$'. The verification conditions for DRS Conditions of the form $K_A \vee K_B$ are those for t inclusive disjunction of classical logic: f verifies $K_A \vee K_B$ iff there exists an extension g *supseteq* f that verifies K_A or there an extension g *supseteq* f that verifies K_B (and this is to be understood as not excluding the possibility that there are extensions of both kinds).[50]

From the point of view of anaphora, disjunctions are also notably different from conjunctions. In general in a Condition of the form $K_A \vee K_B$, neither K_A nor K_B can serve are available as anaphoric antecedents to the other

[50]Remember: '$\vee$' is not to be confused with the ambiguity operator '$\overset{!}{\vee}$' which we first encountered when drawing up a lexical entry for the feature past in Section 3.3.) In other words, a disjunction 'A or B' is represented by the DRS Condition '$K_A \vee K_B$', where K_A and K_B are the representations of A and B.

DRS. This is easy to verify by looking at various disjunctive sentences of English. But it is also and independently suggested by the verification conditions for Conditions of the form $K_A \vee K_B$. According to these the two disjuncts K_A and K_B are *alternative* possibilities; neither can be construed as presupposing or building on the truth of the other.

This implies that neither the drefs from the Universe of the second disjunct K_B are accessible from the first disjunct K_A nor those from the Universe of the first disjunct K_A accessible from the Universe of the second disjunct K_B. This statement is the extension of DRT accessibility that we have been using to complex DRS Conditions of the form $K_A \vee K_B$ (see (Kamp & Reyle 1993), Ch. 1).

Although we cannot implement any assumptions about how anaphora in disjunctions works until a treatment of presupposition has been put into place – we are in the same predicament here as we were in relation to conjunctions in the above discussion of conjunctions – here a couple of observations to show that matters are more complex than this extension of the definition of accessibility implies. First, there are examples like (3.165.a), where it looks like an indefinite in one disjunct can serve as antecedent for a pronoun in the other.

This verdict may seem wrong in the light of examples like (3.165.a).

(3.165)a. Pedro owns a motorcycle or he rents it.

b. Pedro owns it or he rents a motorcycle.

It seems quite unproblematic to interpret the pronoun *it* in the second disjunct of (3.165.a) as anaphoric to the indefinite *a motorcycle* in the first disjunct. And for some speakers even (3.165.b) seems possible. The reason, I suggest, why these interpretations are possible is that indefinite often have non-local scope[51] In relation to (3.165) the wide scope options discussed in the literature come to this: (3.165) can be interpreted as statement to the effect that there is a motorcycle such that Pedro either owns it or, alternatively, that he rents. Why the indefinite DP that is the existential constituent of this existential proposition can occur as a syntactic constituent of one of the disjuncts continues to be topic of debate.

[51]Recall also the remarks on indefinite in generic sentences in footnote 46. There is a rich literature on 'non-local' interpretations of indefinites, including (Fodor & Sag 1982), (Farkas 2002), (Abusch 1993), (Chierchia 2001), (Schwarz 2001), (Endriss 2009), (Kamp & Bende-Farkas 2019).

That (3.165.a) is preferred to (3.165.b) shouldn't be surprising. More surprising is that there are speakers at all who accept (3.165.b). In both (3.165.a) and (3.165.b) the presence of *it* seems to help making the wide scope interpretation of the indefinite possible. But given that the definite has to occur in one of the disjuncts and the pronoun in the other the realization of this in (3.165.a) has the advantage over the realization in (3.165.b) that the pronoun comes after its antecedent – there is a general preference for 'anaphora' over 'kataphora'.

Although we are not in a position here to show how DRSs that assign 'non-local scope' interpretations to indefinite DPs are constructed, we can show the final results of such constructions. (3.166) is the DRS for (3.165.a).[52]

(3.166)
$$\begin{array}{|c|}\hline p, y \\ \hline \text{Pedro'}(p) \quad \text{motorcycle'}(y) \\ \begin{array}{|c|}\hline s_1 \quad t_1 \\ \hline t_1 = n \quad t_1 \subseteq s_1 \\ s_1 : \text{own'}(p,y) \\ \hline \end{array} \vee \begin{array}{|c|}\hline s_2 \quad t_2 \quad u \quad v \\ \hline u = p \quad v = t \\ t_2 = n \quad t_2 \subseteq s_2 \\ s_2 : \text{rent'}(u,v) \\ \hline \end{array} \\ \hline \end{array}$$

A second problem with the definition of accessibility for disjunctive DRS Conditions is illustrated by sentences like (3.167).

(3.167)Either Pedro doesn't own a donkey or he is hiding it.

[52]A proper account of non-local scope indefinites – indefinites that are interpreted as having wider scope than their syntactic position would license for a quantifying DP – would have to deal with a number of different issues: (i) How are the representing drefs of wide scope indefinites to be 'lifted' to the higher DRS Universe that would give the indefinite a case of non-local scope? (ii) How do we make sure that the Conditions that represent the descriptive content of the indefinite end up in the Condition Set of the DRS whose Universe receives the representing dref? (iii) What are the possible positions to which the dref of an indefinite occurring in a given syntactic position may be lifted? (iv) How are we to decide which indefinites are to be given wide scope and how wide a scope they should be given? None of these questions are easy to answer, for formal as well as empirical reasons. Part of the answer is a treatment of wide scope indefinites as triggering 'specificity presuppositions'. Because of this, the problems of wide scope indefinites could be dealt with only in Section 4, after the relevant machinery for dealing with various kinds of presuppositions has been put into place. We will not do this in these Notes, however. For some discussions of specificity in a DRT-related context see (Kamp & Bende-Farkas 2019)

That anaphora is possible in examples like (3.167) goes back to an observation by Karttunen (see (Karttunen 1974)). The basic insight is that in sentences of the form 'not-A or B' the interpretation of B has access to the interpretation of A. The English paraphrase 'A or else B' of 'A or B' helps to understand the underlying semantics and logic of this interpretational option: 'not-A or else B' can be further paraphrased as 'not-A, or in case that not not-A, B'; or simplified, 'not-A, or in case that A, B'. These paraphrases make it plausible (although they of course do not prove) that when the first disjunct of a disjunction is the negation of a sentence A, then the interpretation of the second disjoint B may exploit the interpretation of A.

To formalize this interpretation principle for sentences of the form 'not-A or B' as part of the construction algorithm we are developing in PART II we need a construction rule of a type that we have not yet encountered. It involves the transformation of complex DRSs after they have been constructed. (More such rules are needed in connection with other syntactic constructions, but here we will only consider the one we need to deal with sentences like (3.167).) A preliminary formulation of the rule we need is given in (3.168).

(3.168)
$$\begin{array}{|l|}\hline \\ \hline \begin{array}{|c|}\hline \\ \hline \neg K_A \\ \hline \end{array} \vee K_B \\ \hline \end{array} \rightsquigarrow \begin{array}{|l|}\hline \\ \hline \begin{array}{|c|}\hline \\ \hline \neg K_A \\ \hline \end{array} \vee K_A \bigoplus K_B \\ \hline \end{array}$$

The actual formulation of the rule is more complicated because of the treatment of negation we adopted in Section 3.10.1, which involves the introduction of a new 'negation eventuality'. The following illustration of how the new rule works makes use of the correct formulation, which is left as an exercise to the reader.

The example to show how the rule can be applied is the following variant (3.158) of sentence (3.169) from our discussion of conjunctions.

(3.169)Maria doesn't own a horse or Pedro owns a donkey.

We assume that the two disjuncts have the LFs in (3.170.a,b) and thus that (3.169) as a whole has the LF shown in (3.170.c). Applying the construction algorithm as we have it at this point leads to the DRSs in (3.171.a,b).

(3.170)a.

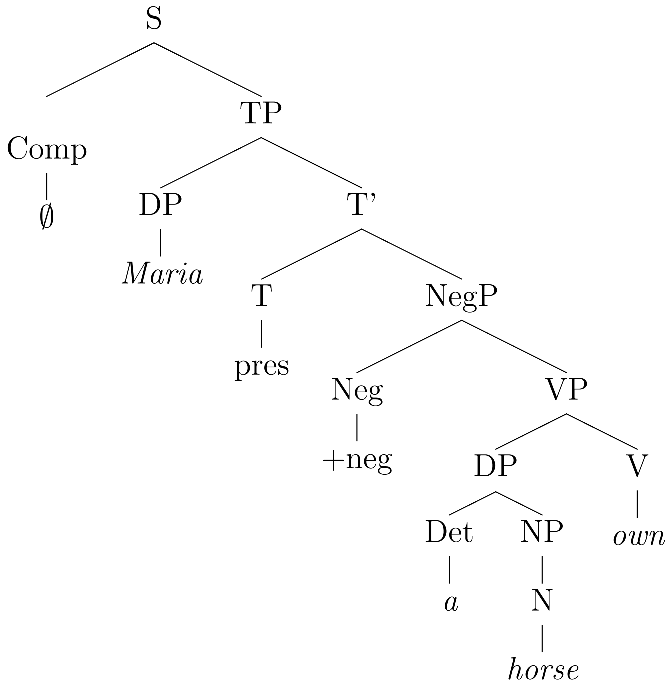
S
Comp
∅
TP
DP
Maria
T'
T
pres
NegP
Neg
+neg
VP
DP
Det
a
NP
N
horse
V
own

b.

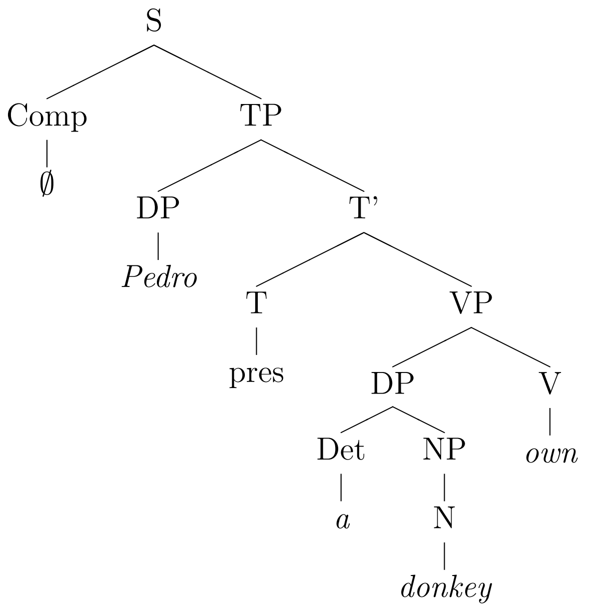
S
Comp
∅
TP
DP
Pedro
T'
T
pres
VP
DP
Det
a
NP
N
donkey
V
own

c.

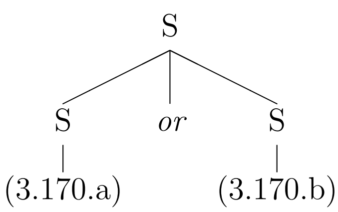
S
S
or
S
(3.170.a)
(3.170.b)

(3.171)a.

$$\begin{array}{|c|}\hline s_1' \quad t_1 \quad m \\ \hline n \subseteq t_1 \quad dur(s_1') = t_1 \\ \text{Maria'}(m) \\ \neg \begin{array}{|c|}\hline s_1 \quad y \\ \hline \text{horse'}(y) \\ s_1\text{: own'}(m,y) \\ s_1 \subseteq s_1' \\ \hline \end{array} \\ \hline \end{array}$$

b.

$$\begin{array}{|c|}\hline s_2 \quad t_2 \quad p \quad z \\ \hline n \subseteq t_2 \quad t_2 \subseteq s_2 \\ \text{Pedro'}(p) \quad \text{donkey'}(z) \\ s_2\text{: own'}(p,z) \\ \hline \end{array}$$

Let us assume that the DRS construction for (3.169) has reached the point shown in (3.172).

(3.172)

$$\begin{array}{ccc} & \text{S} & \\ \text{S} & \textit{or} & \text{S} \\ \begin{array}{|c|}\hline s_1' \quad t_1 \quad m \\ \hline n \subseteq t_1 \quad dur(s_1') = t_1 \\ \text{Maria'}(m) \\ \neg \begin{array}{|c|}\hline s_1 \quad y \\ \hline \text{horse'}(y) \\ s_1\text{: own'}(m,y) \\ s_1 \subseteq s_1' \\ \hline \end{array} \\ \hline \end{array} & & \begin{array}{|c|}\hline s_2 \quad t_2 \quad p \quad z \\ \hline n \subseteq t_2 \quad t_2 \subseteq s_2 \\ \text{Pedro'}(p) \quad \text{donkey'}(z) \\ s_2\text{: own'}(p,z) \\ \hline \end{array} \end{array}$$

At this point the rule for *or* must apply, which turns (3.172) into a DRS with a single $\vee$-Condition in its Condition set:

$$
(3.173)\quad
\begin{array}{|c|}
\hline
\\
\hline
\begin{array}{|c|}
\hline
s_1' \quad t_1 \quad m \\
\hline
n \subseteq t_1 \quad dur(s_1') = t_1 \\
\text{Maria'}(m) \\
\neg \begin{array}{|c|}
\hline
s_1 \quad y \\
\hline
\text{horse'}(y) \\
s_1\text{: own'}(m,y) \\
s_1 \subseteq s_1' \\
\hline
\end{array} \\
\hline
\end{array}
\vee
\begin{array}{|c|}
\hline
s_2 \quad t_2 \quad p \quad z \\
\hline
n \subseteq t_2 \quad t_2 \subseteq s_2 \\
\text{Pedro'}(p) \quad \text{donkey'}(z) \\
s_2\text{: own'}(p,z) \\
\hline
\end{array} \\
\hline
\end{array}
$$

This representation isn't quite what we want insofar as the drefs m and p for the proper names *Maria* and *Pedro* and the associated Conditions 'Maria'(m)' and 'Pedro'(p)' have ended up in the respective disjuncts and not in the Universe and Condition Set of the main DRS. This problem, of which we have encountered a number of instance by now, will be resolved in Section 4, where proper names are treated as triggers of their identification presuppositions (see Section 4.3.1). Let us assume for the sake of argument that this mechanism is already in place and that the representation is as in (3.174), rather than (3.173).

$$
(3.174)\quad
\begin{array}{|c|}
\hline
m \quad p \\
\hline
\text{Maria'}(m) \quad \text{Pedro'}(p) \\
\begin{array}{|c|}
\hline
s_1' \quad t_1 \quad m \\
\hline
n \subseteq t_1 \quad dur(s_1') = t_1 \\
\neg \begin{array}{|c|}
\hline
s_1 \quad y \\
\hline
\text{horse'}(y) \\
s_1\text{: own'}(m,y) \\
s_1 \subseteq s_1' \\
\hline
\end{array} \\
\hline
\end{array}
\vee
\begin{array}{|c|}
\hline
s_2 \quad t_2 \quad z \\
\hline
n \subseteq t_2 \quad t_2 \subseteq s_2 \\
\text{donkey'}(z) \\
s_2\text{: own'}(p,z) \\
\hline
\end{array} \\
\hline
\end{array}
$$

The point has now been reached to apply the rule that this example is meant to illustrate: the non-negated version of the left disjunct may be added to (i.e. merged with) the DRS representing the right disjunct. The result is shown in (3.175).

$$(3.175)\quad \begin{array}{|c|} \hline m \quad p \\ \hline \text{Maria'}(m) \quad \text{Pedro'}(p) \\ \begin{array}{|c|} \hline s_1' \quad t_1 \quad m \\ \hline n \subseteq t_1 \quad dur(s_1') = t_1 \\ \neg \begin{array}{|c|} \hline s_1 \quad y \\ \hline \text{horse'}(y) \\ s_1\text{: own'}(m,y) \\ s_1 \subseteq s_1' \\ \hline \end{array} \\ \hline \end{array} \vee \begin{array}{|c|} \hline s_1 \quad t_1 \quad y \\ s_2 \quad t_2 \quad z \\ \hline n \subseteq t_1 \quad t_1 \subseteq s_1 \\ \text{horse'}(y) \\ s_1\text{: own'}(m,y) \\ \\ n \subseteq t_2 \quad t_2 \subseteq s_2 \\ \text{donkey'}(z) \\ s_2\text{: own'}(p,z) \\ \hline \end{array} \\ \hline \end{array}$$

(3.175) is obtained from (3.174), I claimed, by adding the 'non-negated version' of the disjunct to the left of the $\vee$-Condition to the disjunct on the right. But what exactly is the 'non-negated version' of the disjunct on the left? Intuitively this ought to be a question with an easy answer and in a way it is. (I trust that it is not too hard to see that it is a 'non-negated version' of the disjunct on the left that has been added to the disjunct on the right in (3.175).) But unfortunately, this notion isn't as easy to spell out in formal terms as one might have hoped. The reason is that we have treated negation as syntactically situated somewhere in the middle along the projection line of the verb, and not at the very end, as a modifier of TP or S. One effect of this is that while the dref s_1' for the state that is described as not obtaining by the first disjunct is within the scope of negation (as it should be), but also as one that is contained within a state represented by s_1' which is above negation and locates the negation state in relation to n. There are, as I have tried to make plausible in the section on negation, good reasons for assigning neg the low syntactic position that we have assumed for it in general and that, consistently with this general principle, it has been given in (3.171.a). But in connection with the relation between a clause A and its negation not-A it causes awkward complications. One cannot help feeling that these complications shouldn't be there at all, but i do not see how to get rid of them.

One way in which it might be thought the problem could be solved formally is to prune the LF for the negated clause – in our present example: the LF in (3.170.a) – by eliminating the negation from it; or, what should come to the same thing, to take instead the LF for A, and construct the DRS for that. It is that DRS, then, that the execution of the rule we are discussing should merge with the DRS on the right. But this solution also has its problems. They arise whenever DRS construction is non-deterministic, in that certain construction rules can be applied in more than one way to the input they receive, or perhaps because construction rules can be applied in more than one different order. That is, if in our formulation of the rule we go back to the LFs of the left hand side disjunct not-A and the clause A of which not-A is the negation, then we must make sure that all indeterminacies in the DRS constructions of A and not-A are resolved in the same way. A correct formulation of this kind of 'processing parallelism will vary as a function of the other construction rules of which the construction algorithm has been: for each of these rules whose application contains non-deterministic elements it will be necessary to analyze in detail what its indeterminism exactly comes to and what it is for two applications of the rule to resolve its indeterminism in the same way.

In comparison, the alternative way to proceed – that of subjecting the DRS for the LF of not-A that forms the left disjunct of the given $\vee$ Condition to pruning of all the elements that were introduced into it in the processing of its negation – still seems the better option. But spelling this out in proper detail also requires close attention to detail, having to do with the particular choices we have made in implementing negation as part of DRS construction. Since the details of 'adding K_A to the second disjunct' are of practical relevnce only when the result can be used for anaphora resolution and anaphora cannot be dealt with at this point it would not be useful to go into further detail here.

As concluding observation on the accessibility of A to B in disjunctions of the form 'not-A or B' let me mention the perhaps best-known example of this (yet another example due to Partee):

(3.176)Either there is no bathroom in this house or it is in a funny place.

There is a difference between this example and our earlier example (3.167). The first disjunct of (3.167) has the form of the overt negation of a clause with an indefinite direct object, *a donkey.* In (3.176) the negation is expressed as part of the copula complement *no bathroom.* A first reaction to this variant might be that its left disjunct is logically equivalent to sentence (3.177).

(3.177)It is not the case that there is a bathroom in this house.

So, if we make the plausible assumption that prefixing a sentence with *it is not the case that* is one way of forming its negation, then it might also seem reasonable to interpret 'the DRS which represents the clause of which *there is no bathroom in this house* is the negation' as the DRS which follows *it is not the case that* in (3.177). On this interpretation, applying the rule to the DRS for (3.176) would come to adding the DRS for 'there is a bathroom in this house'. The Universe of that DRS would have a dref for the 'missing' bathroom, which would thus become accessible to the pronoun *it.*

However, as we have seen earlier when discussing Partee's ball example (see in particular Section 2.1 and Section 2.4); when it comes to matters of anaphora, arguments that trade on the logical equivalence of English sentences cannot be trusted. That DPs of the form 'no + NP' can provide anaphoric antecedents for pronouns in other disjuncts cannot have an explanation that is based simply and solely on logical equivalence.

But then, what could be a better argument? Here is a tentative suggestion. The etymology according to which 'no' is a contraction of 'not' and 'a', would seem to have retained some of its power in contemporary English grammar: There must be a level of linguistic representation at which DPs of the form '*no* + NP' are decomposed into negation and indefinite. But this is no more than a hint. I have no suggestion to make about the point or level at which this decomposition should make its entry into the analysis of clauses containing DPs of this form. Fuerther work would be needed to give substance to this hint.

I leave this example as yet another type of challenge for a formally and conceptually precise natural language syntax-semantics interface of the kind pursued in these Notes.

This concludes the present discussion of disjunction. It has been a brief story and one that has been full of loose ends. The reason for this, to repeat, has been that one again we have focused on what makes the realization of a logical operation in natural language more problematic than can be recognized from its logical definition (i.e. the semantic definition of the disjunction operator in classical logic). And for proper attention to these problems, which have to do with the interaction between disjunction and anaphora, the discussion has come too early, at a point where we have left the top down method of DRS construction behind us, but haven't yet made enough progress with the bottom up method to be able to handle pronominal and other anaphora. So why not wait until the development of the bottom up method has progressed far enough? I have no conclusive answer to this question, except that once the representation of quantification, negation and conditionals have been included in our DRS language, it seems natural to also the further complex DRS Conditions that are needed to represent the remaining operators that are part of the standard presentations of propositional and predicate logic.

3.9.5.4 The Polymorphism of Conjunction and Disjunction

There is one final issue to be mentioned here which concerns conjunction as well as disjunction and which could in principle have been mentioned earlier, in the part devoted to conjunction. I will explain below why I have waited with the issue until this point.

This is the issue. Both conjunction and disjunction can not only combine clauses into clauses, but also VPs into VPs, NPs into NPs, PPs into PPs, DPs into DPs and on. (3.178) lists a few examples.

(3.178)a. Frieda didn't close the shop and go home (as she normally does at 6.30).

b. Frieda didn't submit an essay or give a presentation.

c. His lawyer and man for shady deals has just been arrested.

d. She came to the party with a colleague or former student.

e. The bar was behind the church and next to the post office

f. She will be waiting on the terrace or in the bar.

g. A Croatian and a Serbian lost in the quarterfinals.

h. A Croatian or a Serbian won the first prize.

i. A Croatian or Serbian won the first prize.

Some of these sentences are equivalent to sentences in which the *and* or *or* they contain has been fully distributed, in such a way that the expressions that *and* or *or* combine are complete clauses. For instance, (3.178.f) is equivalent to 'She will be waiting on the terrace or she will be waiting in the bar.'

and (3.178.g) to 'A Croatian lost in the quarterfinals and a Serbian lost in the quarterfinals'. But in general the sentences that can be obtained by such a distribution procedure are not equivalent to the starting sentences. For a telling example consider (3.178.b). Clearly, 'Frieda didn't submit an essay or give a presentation.' isn't logically equivalent to 'Frieda didn't submit an essay or Frieda didn't give a presentation.' This shows that the possibility of using *and* or *or* to combine expressions of other syntactic categories than S or TP cannot be treated as a type of syntactic ellipsis, which allows for reconstruction of tacit material at some level of syntactic representation.[53]

The conclusion must be that (i) the LFs for sentences like those in (3.178) must represent their 'ands' and 'ors' as connecting the non-sentential expressions that they connect overtly, and (ii) extend the syntax-semantics interface in such a way that it can deal with non-sentential conjunctions and disjunctions. (i) is a task for syntax and is, consistently with the practice we have been following all along, left to the syntacticians. We will just assume that the parser delivers suitable LFs for such sentences, as in the one example we will consider here.

But (ii) is within our province and some indication of how it can be dealt with should be given. The matter is not completely trivial. What follows won't be systematic account of the problems that a proper treatment of non-sentential conjunctions and disjunctions within our framework will have to deal with. But we will look at a couple of examples that bring out the central hurdle and that will make it possible to show how it can be overcome. We start with sentence (3.178.b), which as we just saw is not equivalent to the result of mechanicallly distributing its *or*.

We assume that (3.178.b) has the LF (3.179) and see how a representing

[53]It might be objected that the distributions involving *and* and *or* are just more complex than I have described them: When *and* or *or* occur 'positively', as for instance in (3.178.d-h), then distribution retains them; if they occur negatively, as in (3.178.a,b), then distribution involves a switch from *and* to *or* and from *or* to *and*. According to this more complicated distribution scheme, 'Frieda didn't submit an essay or give a presentation.' becomes 'Frieda didn't submit an essay and Frieda didn't give a presentation.' and in this case that is an equivalence-preserving transformation. But note that this still doesn't lead to a satisfactory equivalent in cases like (3.178.a) or (3.178.c). Moreover, if switches between *and* or *or* are needed to preserve equivalence when they occur in negative positions, then these distribution transformations could hardly qualify as cases of ellipsis. If by ellipsis reconstruction we understood what has been standardly understood by it, then it wouldn't allow for such switches back and forth between *and* and *or*. So for sentences like (3.178.a) and (3.178.b) ellipsis reconstruction would always give the wrong result.

DRS for the sentence can be computed from this LF.

(3.179)

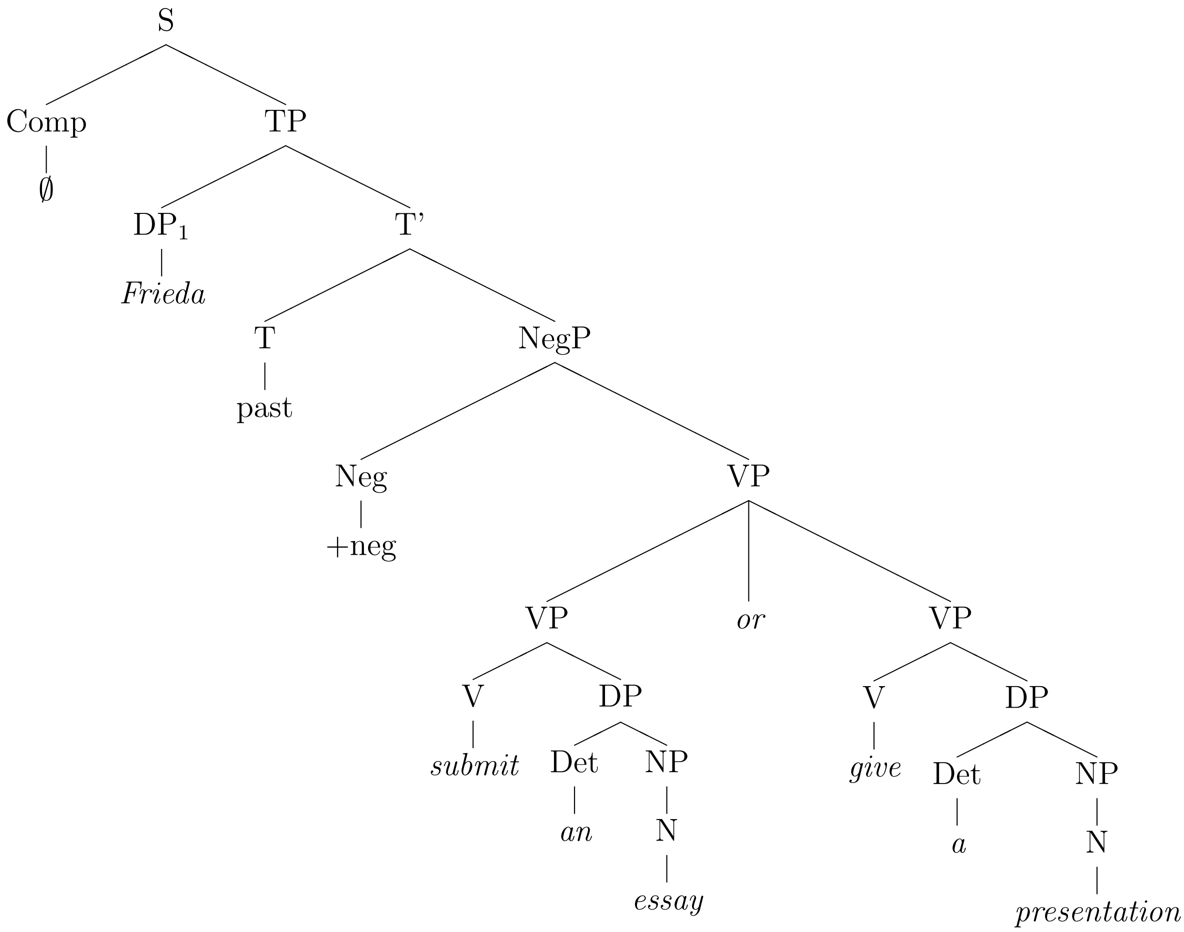

Let us assume that the semantic representations for the VP disjuncts have been constructed and that they are as in (3.180.a,b).

(3.180)a. $<e_{1,ref}, y| \begin{array}{|c|}\hline \\ \hline \text{essay'}(y) \\ e_1 : \text{submit'}(\underline{x}, y) \\ \hline \end{array} >$

b. $<e_{2,ref}, z| \begin{array}{|c|}\hline \\ \hline \text{presentation'}(z) \\ e_2 : \text{give'}(\underline{x}, z) \\ \hline \end{array} >$

The next step must form a disjunctive DRS Condition out of the representations (3.180.a) and (189.49.b). This might have been thought straightforward given the decisions made in the preceding section (the one on disjunction), except that (3.180.a) and (189.49.b) are not DRSs but representations that also involve stores. In forming disjunctions from such representations we follow the principle we adopted when combining representations of *if*-clauses and main clauses into conditional DRS Conditions: the stores are kept local. (The reasons for sticking to this policy for non-sentential disjunctions (and conjunctions) are like those that motivated dealing with stores in the case of conditionals, but I wont' say anything more about those here.) With this additional proviso the representation of the disjunctive VP *submit an essay or give a presentation* is as in (3.181).

$$(3.181) < e_{1,ref}, y| \begin{array}{|c|}\hline \\ \hline \text{essay'}(y) \\ e_1 : \text{submit'}(\underline{x}, y) \\ \hline \end{array} > \vee < e_{2,ref}, z| \begin{array}{|c|}\hline \\ \hline \text{presentation'}(z) \\ e_2 : \text{give'}(\underline{x}, z) \\ \hline \end{array} >$$

The presence of stores local to the disjuncts is typical for the representations of non-sentential disjunctions and constitutes a difference between non-sentential and sentence level disjunctions. Note that local stores pose a general issue of the construction algorithm, which so far I have failed to emphasize. When the point is reached for transferring drefs from local stores to DRS Universes, there is a problem to which DRS Universe each local store dref should go. For sentences with non-sentential disjunctions (and conjunctions) this can become a serious issue, since their preliminary representations may contain a considerable number of local stores – there is no upper bound to how many there ca be. So in general the question which drefs should go where can become quite complex and a special module will have to be added to the construction algorithm to deal with this question. It isn't possible to say more about this matter at this point, since for many drefs – all those that represent anaphoric DPs – the Universes to which they are to be transferred are determined by the resolution of the *identification presuppositions* of those DPs. So any discussion of this module will have to wait until Section 4, in which we will extend our formalism so that it can deal with presuppositions. Even then, however, some difficult questions will remain, among them those pertaining to drefs that represent indefinite DPs. In the meantime we will deal with the problem of store-emptying on an ad hoc basis.[54]

[54] There is also another complication that deserves to be pointed out. In (3.180.a), (189.49.b) and (3.181) the subject argument slots of the verbs *submit* and *give* have both been marked be the slot marker '$\underline{x}$'. This is intuitively right insofar as both slots will

Unfortunately (3.181) isn't quite what we want. The difficulty is that the phrase *submit an essay or give a presentation* is a VP, which ought to get a representation of the kind that we have been assuming in general for VPs. That is, its representation should be an eventuality description. In the case before us the need for this shows up at the very next step, in which the VP should be subjected to negation: The semantics of neg, as defined in Section 3.10.1, expects an eventuality description as input.

One way in which this desideratum can be satisfied is to see the event descriptions provided by the disjuncts in (3.181) as two ways in which an eventuality could be that satisfies the disjunction (3.181) as a whole. We can implement this idea by representing *submit an essay or give a presentation* as the description of an eventuality ev which can be either of the kind described by the first disjunct – i.e. be identical with an event satisfying that description – or of the kind described by the second disjunct and thus identical with some event answering that description. That is, the semantic representation of the disjunctive VP ought to be as in (3.182).

$$(3.182) < ev_{3,ref} | \begin{array}{|c|}\hline \\ \hline <y| \begin{array}{|c|}\hline e_1 \\ \hline ev_3 = e_1 \\ \text{essay'}(y) \\ e_1 : \text{submit'}(\underline{x}, y) \\ \hline \end{array} > \vee <z| \begin{array}{|c|}\hline e_2 \\ \hline ev_3 = e_2 \\ \text{presentation'}(z) \\ e_2 : \text{give'}(\underline{x}, z) \\ \hline \end{array} > \\ \hline \end{array} >$$

Since (3.182) is of the right form for inputs to the negation operator, this operator can be applied. But there is one complication here. Depending on whether the eventuality dref ev_3 represents an event or a state with the result of applying the operator to (3.182) is shown in (3.183).

---

have to be eventually filled by the dref for the sentence subject *Frieda*. But what is really behind this is that both slots will be co-indexed with the subject DP. We will assume that this co-indexation will be in place as soon as the semantic representations of the verbs have been inserted. And again we hold the parser responsible for this: Verification that the string in (3.178 .b) is well-formed and assigning it its syntactic structure (as in (3.179)). Part of that is recognizing the string *submit an essay or give a presentation* as a disjunctive VP with the verbs *submit* and *give* as the verbs of its two disjuncts, each of which has *Frieda* as its subject argument phrase. On the basis of this information the structure obtained from lexical insertion into the parse (3.179) can be enriched with co-indexation of the subject DP with both positions marked by '$\underline{x}$' in (3.181).

$$(3.183) < ev_{4,ref} | \begin{array}{|l|}\hline \\ \hline \neg \;\begin{array}{|c|}\hline ev_3 \\ \hline <y| \begin{array}{|c|}\hline e_1 \\ \hline ev_3 = e_1 \\ \text{essay'}(y) \\ e_1 : \text{submit'}(\underline{x}, y) \\ \hline \end{array} > \vee <z| \begin{array}{|c|}\hline e_2 \\ \hline ev_3 = e_2 \\ \text{presentation'}(z) \\ e_2 : \text{give'}(\underline{x}, z) \\ \hline \end{array} > \\ ev_3 \subseteq ev_4 \\ \hline \end{array} \\ \\ \hline \end{array} >$$

The remainder of the representation construction for (3.178.b) involves principles all of which are known. But there is one complication, which has to do with the application of the operator triggered by the tense feature past. Which of the two options for the application of the semantics of past should apply, that for events or that for states? In the present case the answer is intuitively straightforward. Since both e_1 and e_2 are events, ev_3 can only be an event; and then ev_4 will have to be an event too. So in view the representation in (3.183) can be rewritten as in (3.183.1), where the replacements of 'ev' by 'e' indicate that only events are involved.

$$(3.184) < e_{4,ref} | \begin{array}{|l|}\hline \\ \hline \neg \;\begin{array}{|c|}\hline e_3 \\ \hline <y| \begin{array}{|c|}\hline e_1 \\ \hline e_3 = e_1 \\ \text{essay'}(y) \\ e_1 : \text{submit'}(\underline{x}, y) \\ \hline \end{array} > \vee <z| \begin{array}{|c|}\hline e_2 \\ \hline e_3 = e_2 \\ \text{presentation'}(z) \\ e_2 : \text{give'}(\underline{x}, z) \\ \hline \end{array} > \\ e_3 \subseteq e_4 \\ \hline \end{array} \\ \\ \hline \end{array} >$$

The remaining steps of the construction are then instances of construction principles that are by now well-established. It is left to the reader to complete the construction.

As regards the aspectual status of the eventuality described by the disjunctive VP in (3.183) we were in luck, you might say, insofar as the two disjuncts have the same aspect status: both are event descriptions. But the matter won't always be as straightforward. In the following sentences one VP disjunct is a state and the other an event description.

(3.185)a. On Sunday Frieda will be in Amsterdam (already) or reach Amsterdam at noon.

b. On Sunday Frieda won't either be in Amsterdam (already) or reach Amsterdam at noon.

What is the aspectual status of an eventuality that can either be state of the kind described by one disjunct or an event described by the other disjunct? No coherent answer to this questions seems possible. Rather, the two options – state or event – have to be kept alive during further processing steps and if necessary cashed out separately. In (3.185.a) the point where this matters is when the future tense is applied to its input representation. In this particular case the input representation can be identified with the VP representation in (3.186), which we get by applying the same recipe that gave us (3.183).[55]

$$(3.186) < ev_{3,ref}, a | \begin{array}{|c|}\hline \quad \\ \hline \text{Amsterdam'}(a) \\ \begin{array}{|c|}\hline s_1 \\ \hline ev_3 = s_1 \\ s_1 : \text{in'}(\underline{x}, a) \\ \hline \end{array} \vee \begin{array}{|c|}\hline e_2 \quad t' \\ \hline ev_3 = e_2 \\ e_2\text{: reach'}(\underline{x}, a) \\ \text{noon'}(t') \quad e_2 \subseteq t' \\ \hline \end{array} \\ \hline \end{array} >$$

The problem is now how to apply the semantics of the future tense to this representation. And then the sea problem repeats itself for the temporal location by *on Sunday*. On the one hand, insofar as the eventuality dref ev_3 can represent a state of being in Amsterdam, it wants to be temporally

[55]The dref a representing 'Amsterdam' and the Condition 'Amsterdam'(a)' have been lifted form the representations of the disjuncts to the level of the DRS that contains the $\vee$-Condition. This is a hack, but it is one that will set itself out given the presupposition-based treatment of proper names in Section 4.3.

located as a state, which according to the principles adopted earlier would be that it is going on at the future time t introduced by the future tense, and as temporally including the time t'' representing the referent of the given use of *Sunday*. On the other hand, insofar as ev_3 can represent an event of reaching Amsterdam, it wants to be located as an event, viz. as temporally included within t and also within the time t''.

In essence the mechanism for constructing a representation for the sentence which resolves the apparent conflicts between these requirements. The point is that the inferential process that we have been invoking in earlier applications of fut and past to decide the choice between the two options those features present now has to be made dependent on which disjunct ev_3 is taken to instantiate. In the present instance this process takes, roughly speaking, the form of a 'proof by cases'. The resulting representation is shown in (3.187).

$$
(3.187)\quad
\begin{array}{|c|}
\hline
t \quad ev_3 \quad t'' \quad a \\
\hline
t \prec n \quad \text{Sunday'}(t'') \quad \text{Amsterdam'}(a) \\
\begin{array}{|c|}\hline s_1 \\ \hline ev_3 = s_1 \\ t \subseteq ev_3 \quad t'' \subseteq ev_3 \\ s_1 : \text{in'}(\underline{x}, a) \\ \hline\end{array}
\quad \vee \quad
\begin{array}{|c|}\hline e_2 \quad t' \\ \hline ev_3 = e_2 \\ ev_3 \subseteq t \quad ev_3 \subseteq t'' \\ \text{noon'}(t') \quad e_2 \subseteq t' \\ e_2\text{: reach'}(\underline{x}, a) \\ \hline\end{array} \\
\hline
\end{array}
$$

This DRS can be simplified to the one in (3.188).

(3.188)
$$\begin{array}{|c|} \hline t \quad t'' \quad a \\ \hline t \prec n \quad \text{Sunday'}(t'') \quad \text{Amsterdam'}(a) \\ \begin{array}{|c|} \hline s_1 \\ \hline ev_3 = s_1 \\ t \subseteq s_1 \quad t'' \subseteq s_1 \\ s_1\text{: in'}(\underline{x}, a) \\ \hline \end{array} \quad \vee \quad \begin{array}{|c|} \hline e_2 \quad t' \\ \hline e_2 \subseteq t \quad e_2 \subseteq t'' \\ \text{noon'}(t') \quad e_2 \subseteq t' \\ e_2\text{: reach'}(\underline{x}, a) \\ \hline \end{array} \\ \hline \end{array}$$

<u>Exercise</u> Show how (3.187) can be obtained by completing the DRS construction for (3.185.a), including a detailed account of how the application of fut and the temporal adverb rule for *on Sunday* give rise to the temporal Conditions in (3.187).

(3.185.b) differs from (3.185.a) in that here the aspectual ambivalence of the referential argument ev_3 of the disjunctive VP first arises as an issue in the application of neg. The new dref ev_4 introduced by the application of neg – compare (3.183) – is supposed to inherit its aspectual status from the referential argument (ev_3) of the input representation to which neg is applied. The simplest way to implement this idea is to add Conditions to the result of the application which state that ev_4 is an event/state if ev_3 is: 'EVENT(ev_3) *Rightarrow* EVENT(ev_4)' and 'STATE(ev_3) *Rightarrow* STATE(ev_4)'. These Conditions can then be used in the derivation of how tense features sand temporal locating adverbs are to be applied.

The general rule for forming the representations of VP disjunctions is given in (3.189).

(3.189) Suppose that VP is the disjunctive VP '(*either*) VP_1 or VP_2' and that VP_1 and VP_2 have the semantic representations $< ev_{1,ref}, \alpha_1..\alpha_n, | K_1 >$ and $< ev_{2,ref}, \beta_1..\beta_m, , | K_2 >$. Then the semantic representation of VP is:

$$< ev_{3,ref} \mid \begin{array}{|c|} \hline \\ \hline < \alpha_1..\alpha_n, | K_1' > \vee < \beta_1..\beta_m \mid K_2' > \\ \hline \end{array} >.$$

Here K_1' is the result of adding ev_1 to the Universe of K_1 and the Condition '$ev_3 = ev_1$' to K_1's Condition Set. Likewise for K_2'.

Although stated as a rule for constructing semantic representations for disjunctive VPs, (3.189) is applicable to many other non-sentential disjunctions as well. Any disjunction of semantic representations that function as descriptions can be represented this way. Among the grammatical categories for which semantic representations of disjunctions can be constructed according to (3.189) are: NPs, PPs, APs, AdvPs and Relative Clauses. For one further example, consider the disjunctive NP *Serbian or Croatian* in (3.178.i). (3.190.a,b) give the semantic representations of the nouns *Serbian* and *Coratian*. Applying the generalized version of (3.189) to these gets us the representation in (3.190.c) for the disjunctive NP. (3.190.d) gives the DRS for the sentence (3.178.i). (Note in particular that the semantic representation of the DP *a Serbian or Croatian* is formally indistinguishable from the semantic representation of its NP.)

(3.190)a. $$< x_{ref} \mid \begin{array}{|c|}\hline \\ \hline \text{croatian'}(x) \\ \hline \end{array} >$$

b. $$< y_{ref} \mid \begin{array}{|c|}\hline \\ \hline \text{serbian'}(y) \\ \hline \end{array} >$$

c. $$< z_{ref} \mid \begin{array}{|c|}\hline \\ \begin{array}{|c|}\hline x \\ \hline \text{croatian'}(x) \\ z = x \\ \hline \end{array} \vee \begin{array}{|c|}\hline y \\ \hline \text{serbian'}(y) \\ z = y \\ \hline \end{array} \\ \\ \hline \end{array} >$$

d.
$$\begin{array}{|c|}\hline t\quad e\quad z\quad v \\ \hline t \prec n \quad e \subseteq t \\ \text{“the first prize”}(v) \\ \\ \begin{array}{|c|}\hline x \\ \hline \text{croatian'}(x) \\ z = x \\ \hline \end{array} \vee \begin{array}{|c|}\hline y \\ \hline \text{serbian'}(y) \\ z = y \\ \hline \end{array} \\ \\ e : \text{win'}(z,v) \\ \hline \end{array}$$

Constructing semantic representations for disjunctions according to (3.189) also works for some DPs. An example is (3.178.h). As just noted in connection with (3.178.i), it is part of the construction algorithm developed in PART II that the semantic representations for the DPs look the same as the semantic representations for their NPs. (This is a distinctive property of indefinites (as opposed to other types of DPs), and a consequence of the role and status of the referential arguments of indefinite DPs and their NPs in our DRT-based set-up. The difference between NP- and DP-representations is in this case just how they will be used in further construction rule applications.) The upshot of this that the DRS for (3.178.i) is indistinguishable from that for (3.178.h). (If this may seem surprising, it certainly gibes with the intuition that (3.178.i) and (3.178.h) mean the same and seem hard to keep apart.

In general, however, (3.189) is not the right recipe for dealing with disjunctive DPs. Some dramatic examples of where (3.189) does not work are shown in (3.191).

(3.191)a. I will invite every boyfriend of Bill's or every girlfriend of Bill's.

b. I will not invite every boyfriend of Bill's or every girlfriend of Bill's.

c. I will invite every boyfriend of Bill's and every girlfriend of Bill's.

d. I will not invite every boyfriend of Bill's and every girlfriend of Bill's.

Disjunctions of quantifying DPs such as *every boyfriend of Bill's or every girlfriend of Bill's* is an altogether different story from disjunctions of expressions that function as predicates. The first and main reason for this is that the contributions that quantifying DPs make to sentence meaning take the form of operation on the representations of their sister nodes. The second reason is that for reasons discussed at some length in PART I we decided to treat quantifying DPs via Quantifying Raising, following in the essentials Heim & Kratzer.

The problems with a QR-based analysis of sentences like (3.191.a,b) are particularly plain. Consider for s start (3.191.a). For purely intuitive reasons Quantifier Raising out of a disjunction might seem an odd sort of operation. The main problem is: when the raised quantifying DP finally makes its contribution to the semantics of the sentence, its referential argument must be inserted into the right argument slot of some predicate in the nuclear scope of the Duplex Condition it introduces. What could that slot be in the case of (3.191.a)? There seems to be only one candidate, viz. the direct argument slot of *invite*. So when the raised DP is the left disjunct *every boyfriend of Bill's*, then the contribution made by this quantifier would presumably be the Duplex Condition in (3.192).

(3.192)
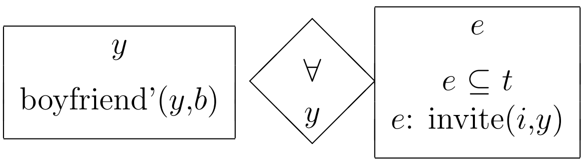

The other quantifying DP, *every girlfriend of Bill's*, also needs to be raised and will then also give rise to a Duplex Condition, with the same nuclear scope as in (3.192). Intuitively the representation of the sentence as a whole ought to be the disjunction of these two Duplex Conditions. This is because the direct object DP of (3.191.a) is a disjunctive DP. But how can the construction algorithm make use of this information?

Perhaps the most plausible story that can be told about this is one according to which the disjunctive direct object DP of (3.191.a) is raised to some higher adjunct position. I am not sure how high this position should be. In the present situation the most obvious adjunction site choices would be VP and T', but I have no more to say about this further question. In what follows right here I assume that the adjunction site is VP.

We assume, then, that the initial syntactic structure for (3.191.a) is that in (3.193).

(3.193)a.

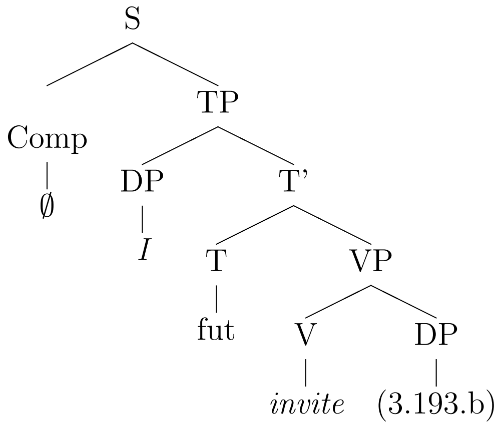

b.

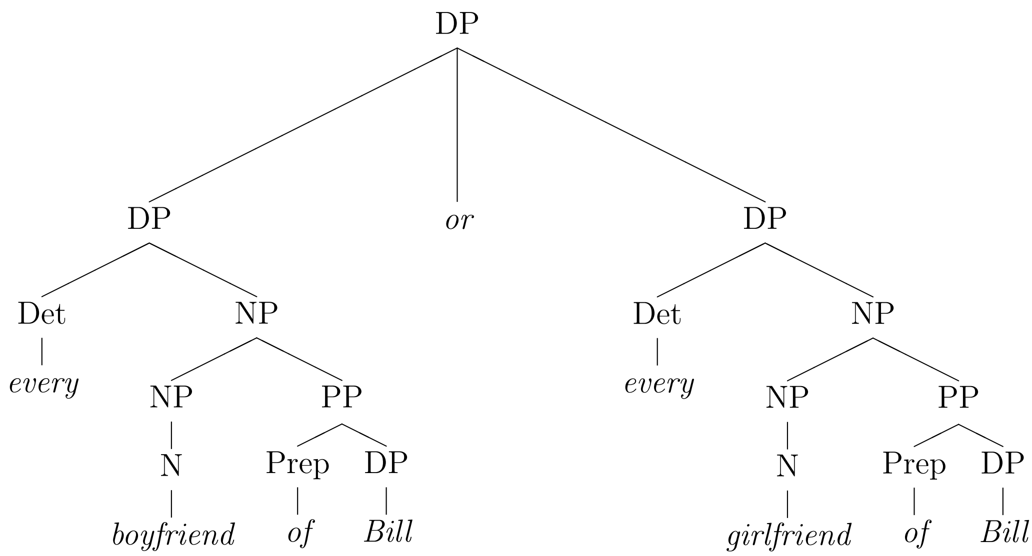

Insertion of the semantic representation for *invite* and raising of the direct object DP to VP adjunct transforms this structure into (3.194).

(3.194)

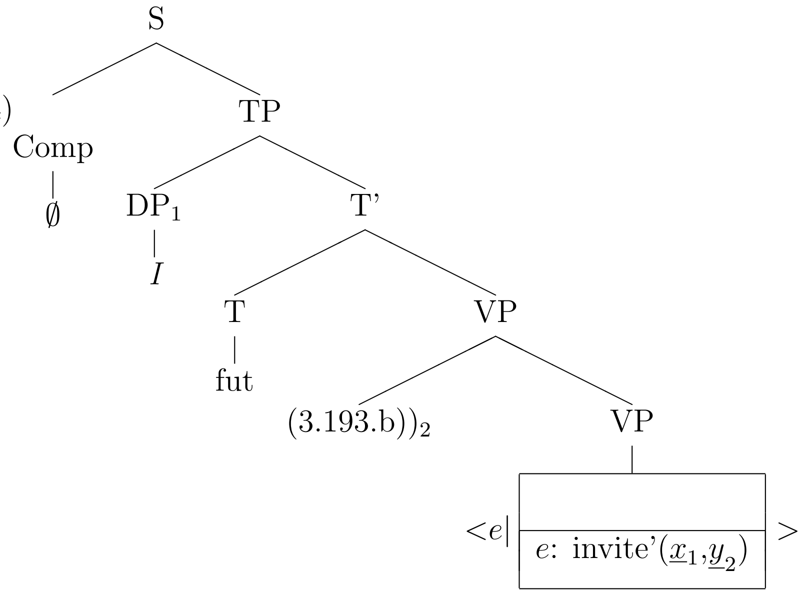

Distribution of the *or* of the raised DP now takes the following form: We make two copies VP1 and VP2 of the upper VP of (3.194). In VP1 the disjunctive DP is replaced by its first DP disjunct, yielding VP'1, and in VP2 it is replaced by the second disjunct, yielding VP'2. Then the two structures VP'1 and VP'2 are combined into a disjunction of category VP, which replaces the upper VP of (3.194). The result is shown in (3.195).

(3.195)a.

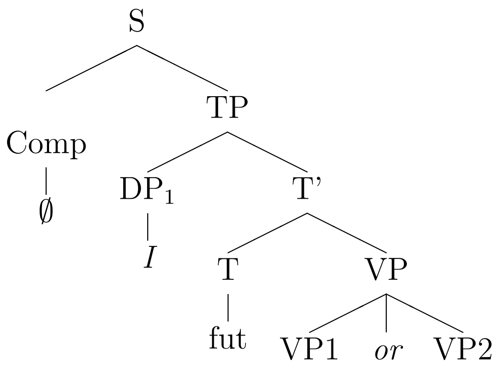

b. 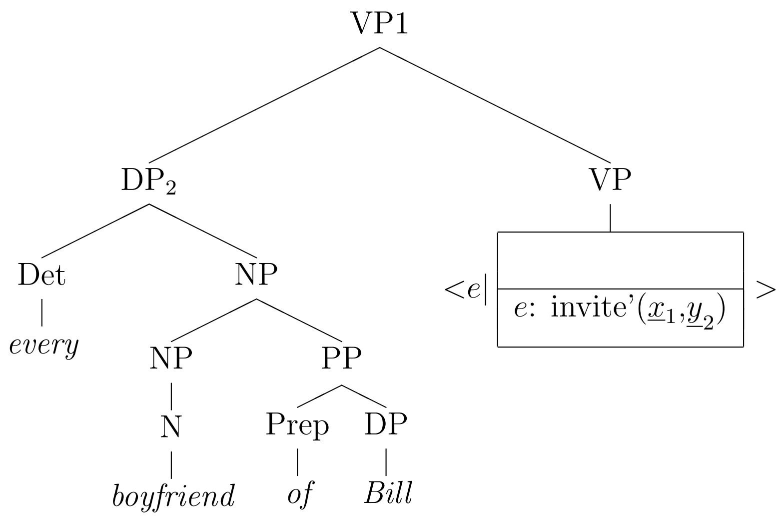

c. 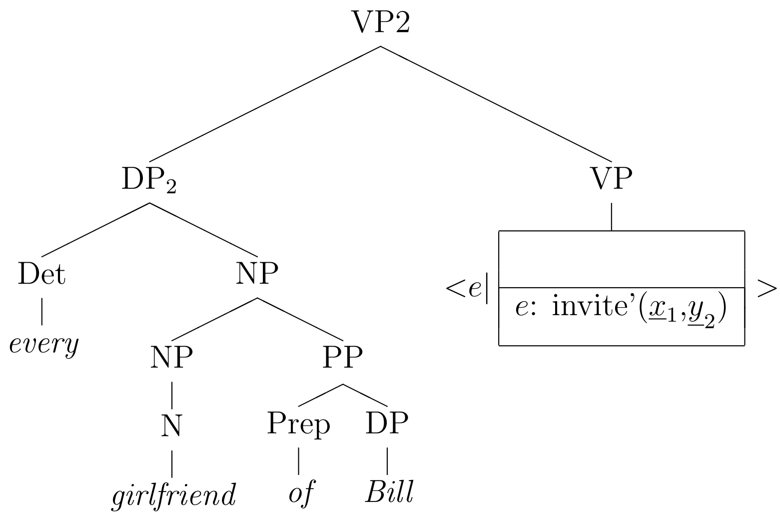

The next step can now be the construction of the semantic representations of the NP parts of the quantifying DPs of VP1 and VP2 and the introduction of Duplex Conditions to represent the quantifying power of these DPs. (The introduction of the Duplex Conditions proceeds according to the same principles as before, with the the representation of the adjunction site becoming the nuclear scope.) The representation of the NP of the first DP disjunct is given in (3.196.a) and the resulting representation of VP1in (3.196.b). The representation of VP2 is analogous.

(3.196)a. $<y, b \left| \begin{array}{|c|} \hline \\ \hline \text{Bill'}(b) \\ \text{boyfriend-of'}(y,b) \\ \hline \end{array} \right. >$

$$\text{b.}\quad <s'|s': \begin{array}{|c|}\hline y \\ <y,b\ | \begin{array}{|c|}\hline \\ \hline \text{Bill'}(b) \\ \text{boyfriend-of'}(y,b) \\ \hline \end{array} > \\ \hline \end{array} \ \left\langle \begin{array}{c} \forall \\ y \end{array} \right\rangle \begin{array}{|c|}\hline e \\ \hline e \subseteq s' \\ e\text{: invite'}(\underline{x}_1,y) \\ \hline \end{array} >$$

Applying the rule for forming ∨-Conditions leads for the upper VP of (3.195.a) the representation in (3.197). And from there familiar construction principles get us to the DRS for sentence (3.191.a) that is shown in (3.198).

$$(3.197) <s\ | \begin{array}{|c|}\hline s\ \ b\ \ b' \\ \hline \text{Bill'}(b) \quad \text{Bill'}(b') \\ \begin{array}{|c|}\hline s' \\ \hline s = s' \\ s'\text{: } \begin{array}{|c|}\hline y \\ \hline \text{boyfriend-of'}(y,b) \\ \hline \end{array} \left\langle \begin{array}{c} \forall \\ y \end{array} \right\rangle \begin{array}{|c|}\hline e \\ \hline e \subseteq s' \\ e\text{: invite'}(\underline{x},y) \\ \hline \end{array} \\ \hline \end{array} \\ \vee \\ \begin{array}{|c|}\hline s'' \\ \hline s = s'' \\ s''\text{: } \begin{array}{|c|}\hline z \\ \hline \text{girlfriend-of'}(z,b') \\ \hline \end{array} \left\langle \begin{array}{c} \forall \\ z \end{array} \right\rangle \begin{array}{|c|}\hline e \\ \hline e \subseteq s'' \\ e\text{: invite'}(\underline{x},z) \\ \hline \end{array} \\ \hline \end{array} \\ \hline \end{array} >$$

(3.198)

$$\boxed{\begin{array}{c} t \quad s \quad b \quad b' \\ \hline n < t \quad t \subseteq s \quad \text{Bill'}(b) \quad \text{Bill'}(b') \\ \boxed{\begin{array}{c} s' \\ \hline s = s' \\ s'\text{:} \ \boxed{\begin{array}{c} y \\ \hline \text{boyfriend-of'}(y,b) \end{array}} \ \left\langle \begin{array}{c} \forall \\ y \end{array} \right\rangle \ \boxed{\begin{array}{c} e \\ \hline e \subseteq s \\ e\text{: invite'}(\underline{x},y) \end{array}} \end{array}} \\ \vee \\ \boxed{\begin{array}{c} s'' \\ \hline s = s'' \\ s''\text{:} \ \boxed{\begin{array}{c} z \\ \hline \text{girlfriend-of'}(z,b') \end{array}} \ \left\langle \begin{array}{c} \forall \\ z \end{array} \right\rangle \ \boxed{\begin{array}{c} e \\ \hline e \subseteq s \\ e\text{: invite'}(\underline{x},z) \end{array}} \end{array}} \end{array}}$$

(To obtain the DRS for (3.191.a) in this final form we have resorted once more to ad hoc lifting of b, b', 'Bill'(b)' and 'Bill'(b')'; also the $\vee$-Condition of (3.197) has been merged with the main DRS.)

I have added the sentence (3.191.b) – ' I will not invite every boyfriend of Bill's or every girlfriend of Bill's' – to the list of examples discussed in this section from the desire to include a close replica of (3.178.a) (the sentence 'Frieda didn't submit an essay or give a presentation.'). But curiously (3.191.b) doesn't seem to be an acceptable sentence on either of the two interpretations that would come to mind on current theoretical grounds: that in which 'or- has narrow scope with respect to 'not' and that where the scope relation is the reverse. Presumably there is some violation of 'negative concord' here; but I am not sure exactly how this principle might be stated, given that (3.178.b) is not affected by concord problems.

However, while (3.191.b) is no good, sentence (3.191.d), in which *or* is replaced by *and* is acceptable, just as its unnegated counterpart (3.191.c). Presenting these examples of non-sentential conjunction as part of (3.191), is slightly running ahead of schedule, but not by much. Non-sentential conjunctions will be the topic of the next part of this discussion of disjunction and conjunction polysemy.

I have gone through a possible DRS construction for (3.191.a) with what might be perceived as an almost masochistic drive towards undoing pretty much all that I have endeavored to put into place up to this point. For the construction sketched is at variance with some of the most central construction principles we have adopted – principles for dealing with nominal quantification that have been with us almost from the beginning of these Notes, and which thus far survived the transition from PART I to PART II. In fact, the moral of this attempt to get at a reasonable blue-print for the construction of semantic representations with DP disjunctions involving quantifying DPs is that DP coordinations force us to rethink the entire range of assumptions we have made when trying to formulate rules for dealing with DP quantifiers. Such a rethinking would be a major undertaking and it is one that I will not engage in in these Notes. But nevertheless, this is to be a warning that some major rethinking is required for the reasons indicated should be kept in the backs of our minds, and a warning with a big exclamation sign attached to it.

<u>Exercise/Research Project.</u> Try to find a reformulation of the principles governing quantifying DPs which (a) can deal with DP coordinations involving quantifying DPs like the one in (3.191.a) and (b) which makes the same predictions about sentences involving quantifiers that have been considered up to now.

3.9.5.5 Non-sentential Conjunctions

Semantic representations of non-sentential conjunctions can be constructed along more or less the same lines as semantic representations of non-sentential disjunctions. And it is not uncommon that sentences with non-sentential conjunctions are easier to interpret than the corresponding sentences with *or*. (The sentence pair (**??**) and (**??**) is a good example.) The reason why I decided to start with non-sentential disjunctions nonetheless is that non-sentential conjunctions can be subject to a mode of interpretation for which there is no equivalent among the corresponding disjunctions (those expressions that can be obtained by replacing the *and* of non-sentential conjunction by *or*). The most prominent and common examples of this are conjunctions of DPs. A first piece of evidence of how DP conjunctions differ from DP disjunctions is given by the sentences in (3.199).

(3.199)a. She came to the party with a colleague or former student.

b. She came to the party with a colleague or a former student.

c. She came to the party with a colleague and former student.

d. She came to the party with a colleague and a former student.

The first two sentences of (3.199) illustrate a point already made: The disjunction of indefinite DPs that form the direct object of (3.199.a) and the direct object DP in (3.199.b), which is an indefinite with a disjunctive NP, seem to make no difference to the truth conditions of (3.199.a,b). But when *or* is replaced by *and* in these direct object DPs this is no longer so: The meanings of (3.199.c) and (3.199.d) are very different. (3.199.c) speaks of an individual who was both a colleague and a former student, (3.199.d) of two people, one of whom was a colleague and the other a former student.

One way in which this difference between (3.199.c) and (3.199.d) is often explained is to assume that when it connects DPs *and* stands for a fundamentally different operation than the one that it stands for when it is used to combine expressions of most other syntactic categories. NP is one of the latter categories. An NP of the form 'NP_1 and NP_2' expresses the property that a thing has iff it has both the property expressed by NP_1 and that expressed by NP_2. But a DP of the form 'DP_1 and DP_2' does not denote an individual that could be characterized both as the denotation of DP_1 and as the denotation of DP_2. What 'DP_1 and DP_2' denotes is the two-element set consisting of the denotations of DP_1 and DP_2; or, in our mereological ontology: 'DP_1 and DP_2' denotes the mereological sum of the denotations of DP_1 and DP_2. I will assume that this is always the case:

(3.200)The denotation of a conductive DP 'DP_1 and DP_2', where DP_1 and DP_2 are non-quantificational, is the mereological sum $\alpha \bigoplus \beta$ of the denotations α of DP_1 and β of DP_2.

Further evidence for this assumption is provided by conjunctive DPs in subject position. With conjunctive subject DPs verb-subject agreement always takes the form of plural morphology on the verb. In English this fact doesn't show up in all cases, but it does in the 3 person present tense, as shown in (3.201).

(3.201)a. A doctor and a nurse √are/*is in the room upstairs.

b. A doctor or a nurse *are/√is in the room upstairs.

c. Jim and James √are/*is in the room upstairs.

d. Jim or James *are/√is in the room upstairs.

Morphology makes clear that the subject of (3.201.a) i understood to be a set of two or more individuals (or a non-atomic mereological individual), whereas the subject of (3.201.b) is a single individual, of which it is left open whether it is a doctor or a nurse. A further illustration is provided by (3.201.c,d) where the constituent DPs are proper names. (3.201.d) expresses uncertainty whether the person that is said to be in the room upstairs is Jim or James. In contrast, (3.201.c) claims that both Jim and James are in the upstairs room. (*Jim and James* emphatically cannot be understood as denoting a single person who goes both by the name Jim and the name James. This is why the *is*-version of (3.201.c) is a case of plain ungrammaticality.)

Since plural have been declared terra incognito in PART II of the Notes, I won't have more to say about conjunctive DPs in this section.[56] But the question of non-atomic denotations also arises in connection with other grammatical categories. The conductive NP in (3.202.a) behaves semantically like that of (3.199.c); it expresses the property that is the conjunction of the property of being a colleague and that of being as former student. This is in keeping with the singular form*is*: the sentence is speaking of a single individual. (And in this respect we note once more the similarity with its disjunctive alternative in (3.202.b).)

(3.202)a. A colleague and former student is stalking her.

b. A colleague or former student √is/*are stalking her.

c. A colleague and former student are stalking her.

d. A colleague and a former student are stalking her.

[56]But note well, the problems that arise in connection with coordinations of DPs that we discussed in connection with the disjunctive DPs in (3.191) arise for DP conjunctions as well. The problems raised by the conjunctive DPs in (3.191.c,d) are not quite the same as those raised by the disjunctive DPs in (3.191.a,b), and they take on a different flavor in particular in the context of the claim made in (3.200), according to which conjunctions of non-quantificational DPs denote non-atomic individuals of a mereological ontology. Pursuing this matter seriously would get us far more deeply into various problems having to do with plurality than would be reasonable given that plurality has been set aside, and so we don't go in that direction here.

A curious case, which does not fit what I have been saying up to this point, is (3.202.c). This sentence ought to be no good given what has been said so far. it probably isn't, as opposed to (3.202.d) which is perfect and unambiguously speaks of the direct object being stalked by two people, one a colleague and the other a former student. If (3.202.c) is acceptable (and with this same meaning), then this may well be its close phonological resemblance with (3.202.d), which might encourage interpreters to think of the second *a* in (3.202.d) as 'slurred away'. (According to verdicts I have got from native speakers about analogous examples from German, where there is no such way of slurring the indefinite article and where the distinction between plural vs. singular marking on finite verbs is generally more prominent than in English, the counterpart of (3.202.c) is unequivocally ungrammatical.

Exercise: In (3.178.c), repeated below, we find a conjunction of two NPs. This conjunction is the NP component of the subject DP *His lawyer and man for shady deals*, a definite description which properly denotes iff there is a unique person (within the search space determined by the context) who is both a lawyer of the referent of *his* and that person's man for shady deals. Compare this sentence with the variants (3.203.a,b). What kinds of conjunctions do we find in (3.203.a) and (3.203.b)? What are the truth conditions of these two sentences? [Hint: phrases beginning with *the* can be used both as DPs and as NPs. (For instance, they are used as NPs when they occur as copula complements, as in 'He is the man for shady deals'.)]

(3.178.c) His lawyer and man for shady deals has just been arrested.

(3.203)a. His lawyer and the man for shady deals has just been arrested.

b. His lawyer and the man for shady deals have just been arrested.

[end Exercise]

The principle that a conjunctive NP denotes the conjunction of the properties denoted by its conjuncts equally applies to other grammatical categories that are used to express properties of individuals. (3.204.a) gives an example of a PP adjunct to NP and (3.204.c) for an AP that plays the role of copula complement. I have added the *or*-counterparts to these PP- and AP-conjunctions, to bring out once more the systematic semantic parallels between *and* and *or* for these grammatical categories. (3.204.a) describes the place where the subject will be waiting as having these two properties: It will be (i) in front of the church and (ii) opposite the Post Office. That gives a fair amount of information about the place, and a good deal more than (3.204.b), which only claims that the place has one of these properties and doesn't give the addressee much to go on. In a similar vein (3.204.c) says that the person waiting will wear a blouse that is both black and low-cut, a reasonably good clue as to what to look for. In contrast, (3.204.d) only says of the blouse the subject will be wearing that it will be either black or low-cut. That may leave the addressee in a quandary as to what he should focus on and seems an unpropitious starting point for picking his contact out.

(3.204)a. She will be waiting in a place in front of the church and opposite the Post Office.

b. She will be waiting in a place in front of the church or opposite the Post Office.

c. She will be wearing a blouse that will be black and low-cut.

d. She will be wearing a blouse that will be black or low-cut.

3.9.5.6 Conjunctions of Verbal Projections

The matter is different for conjunctions of verbal projections: conjunctions of expressions of category V, of category VP, and so on up the verbal projection line up to – at least – T'. Consider the following examples.

(3.205)a. Frieda submitted an essay and gave a presentation.

b. Frieda closed the shop and went home (as she normally does at 6.30).

c. Frieda closed the door and turned the key.

d. Frieda went to the cemetery and took the bus.

e. Frieda didn't submit an essay and give a presentation.

f. Frieda didn't close the shop and go home (as she normally does at 6.30).

g. Frieda didn't close the door and turn the key.

h. Frieda didn't go to the cemetery and take the bus.

Start with (3.205.a), once more to be thought of as said in the context of what different students in a class did towards meeting the class requirements. This sentence involves the conjunction of two VPs and it is naturally understood as describing two past events, one of Frieda submitting an essay and one of her giving a presentation. If this is right, then it seems to indicate that its VP conjunction describes a complex event that is the mereological sum of an event of submitting an essay and an event of giving a presentation. Given this interpretation of the VP conjunction we would expect that (3.205.c) expresses just the complement of these truth conditions, viz. that it is not the case that Frieda did two things, submit an essay and give a presentation. The sentence only records that she was not among those who did both. I believe that this is indeed the primary interpretation of the negation of (3.205.a) in (3.205.e) and perhaps it is the only one. We will return to this question below.

The next three sentences in (3.205) are different. The VP of (3.205.b) is naturally understood as describing a single complex event, that of Frieda closing up and going home. That is why the negation of this sentence, in (3.205.f), seems to just state that on the day that it speaks about there was no such complex event. This is of course also compatible with the possibility that Frieda did one of these two things on its own – that she closed the shop without going home or that she went home without closing the shop. So at this level there is no tangible difference between the way in which the truth conditions of (3.205.f) relate to those of (3.205.b) and the way in which the truth conditions of (3.205.e) relate to those of (3.205.a). In (3.205.c) the connection between the events described by the VP conjuncts is even more intimate. Closing and locking a door is a single complex event, with an internal causal-like structure in that the door has to be closed in order for turning the key to be possible, or at least to make the intended sense. Negating this sentence, as in (3.205.g), is naturally understood as the denial that such a complex event occurred. Questions about whether one of the part events may have occurred without the other aren't introduced as possible alternatives by this utterance, as they would have been by, say, an utterance of 'Frieda didn't closed the shop and didn't go home'.[57] Information structure is a topic set aside in these Notes and we won't further elaborate on this last remark.

(3.205.d) represents yet another case. Here the second VP conjunct – or T' conjunct; it doesn't matter for the point I want to make how we analyze this conjunction syntactically – describes the manner in which the going to the cemetery was executed. Here an analysis according to which the two conjuncts of the conjunctive VP describe different parts of the event described by their conjunction makes no sense; each conjunct describes an aspect of the same event which is not presented as having internal event structure. The

[57]This last sentence presupposes a certain part of the theory of *information structure*. More particularly, what is being implied here is that utterances bring certain a;alternatives into play – situations that the utterance excludes as *not* described. The alternatives theory of Rooth (see (Rooth 1985), (Rooth 1992) is the first decisive move in this direction.

negation (3.205.h) of this sentence is hard to interpret. Either one should negate that Frieda went to the cemetery at all, as in 'Frieda didn't go to the cemetery' or one should say that the manner of her going to the cemetery was not that of taking the bus, as in 'When Frieda went to the cemetery she didn't take the bus'.

To sum up this all too brief discussion of conjunctions of Verb projections: many of the events we describe in language are described or conceived as having internal 'event structure'. A description by a conjunction of verb projections can therefore often be understood as the description of a complex event $e_1 \bigoplus e_2$, with e_1 described by the first conjunct and e_2 described by the second conjunct. But this is only one of several possibilities. It is also possible for the two conjuncts to describe events that need not stand in any causal or other ontological relation to each other; and on the other hand the two conjuncts can be understood as providing complementary descriptions of a single event that cannot be decomposed into two components each of which is described by one of the conjuncts.

There thus appears to be a crucial difference between conjunctions in the verbal and in the nominal domain. Conjunctions of verbal descriptions can function as descriptions of simple, non-descomposable event, as descriptions of compound events of which each conjunct describes a component and as descriptions of pairs of mutually independent events, where each event is described by its own conjunct. Finer distinctions can be made between the different ways in which the described component events are connected or are jointly form a compound event and quite probably will have to bee made as part of a more refined semantics for conjunctions of verb projections. But in any case we find a spectrum of different possibilities here which we do not find in the nominal domain. There the grammatical categories are of two kinds: on the one hand the property-denoting categories like NP and adnominal PP, where conjunctions express the conjunctions of the properties denoted by they conjuncts and on the other the referential category DP, where a conjunction never denotes a single individual but rather the mereological sum of the denotations of the conjuncts.

<u>Exercise</u> A further complicating factor in the analysis of coordinated verb projections are certain aspects of syntax. (3.206.a) seems more or less synonymous with (3.205.e): Thee was something that Frieda did not do and that was: submit an essay and give a presentation. As noted in connection with (3.205.e), these sentences can be true even when Frieda did one of the two things of which the sentence denied that she did them together. But

(3.206.b) is different. this sentence only has the interpretation that Frieda did neither thing: she neither submitted an essay nor gave a presentation.

(3.206)a. Frieda failed to submit an essay and give a presentation.

b. Frieda failed to submit an essay and to give a presentation.

Find an explanation for why (3.206.b) and (3.206.b) differ in this way.
[end Exercise]

3.9.5.7 Back to Conjunctions of DPs

The fact that conjunctions of referential DPs denote the mereological sums of the denotations of their conjuncts is connected with a range of puzzles that sentences with such conjunctions present us with. The complexities of sentences with DP conjunctions are reminiscent of those we observed earlier for disjunctions of quantifying DPs. But the problems are not identical and having a look at some examples involving conjunctions of non-quanfiticational DPd will be instructive,for a reason that more will be said about towards the ned of this section about polymorphism.

We start with the following examples.

(3.207)a. Louise didn't interview the butler and the gardener. But she did interview the cook and the groom.

b. Louise didn't interview every chambermaid and every footman. But she did interview the cook and the groom.

c. Louise interviewed every chambermaid and every footman.

d. Louise interviewed the butler and the gardener.

The obvious interpretation of the second sentence in (3.207.a,b) is that the speaker interviewed both the cook and the groom. On the assumption that the conjunction *the cook and the groom* denotes the set consisting of the denotations of *the cook* and *the groom* and that the verb *interview* distributes over this set when the set occupies its direct object position, this is as expected. But what about the first sentences? The natural interpretation of the first sentence of (3.207.a) appears to be that Louise interviewed *neither* the butler *nor* the gardener (although the logically weaker interpretation according to which the speaker din't interview both is possible as well, as testified by 'True, I didn't interview the butler and the gardener. But I interviewed one of them'). How can the logically stronger reading be accounted for?

In the light of all that has been said so far there are two directions in which we may look for an explanation. The first has to do with our exploration of sentences involving disjunctions of quantifying DPs at the hand of example (3.191.b). The second builds on the principle that DP conjunctions denote mereological sums. As regards the first direction, let's begin by considering, as the first leg of a kind of bridge with the earlier discussion of (3.191.b), the sentence in (3.207.c). It would seem plausible that the mechanisms needed to deal with (3.191.b) should also apply to this sentence. If so, and if what we have been saying about (3.191.b) has been on the right track, it would seem reasonable to assume that operations similar to those we conjectured in connection with the disjunction of quantifying DPs in (3.191.b) are available also in connection with DP conjunctions like that of (3.207.c): (i) the conjunctive DP can be raised to an adjunct position to VP and (ii) distribution is then possible of the conjunct DPs over the adjunction site. In other words, the syntactic structure in (3.208) can be transformed into (3.209) and (3.209) then permits the distribution of *and* over the representation of the conjunctive DP's adjunction site so as to yield a VP conjunction of the form $[\ [\ \ ]_{VP1}\ \ and\ \ [\ \ ]_{VP2}\ \ ]_{VP}$. The first conjunct VP1 of this VP has the semantic representation given in (3.210). The conjunctive DRS Condition representing the upper VP of (3.209.a,b) is shown in (3.211).

There is a difference between forming the Conjunction Condition (the DRS Condition whose main operator is ;) and the $\vee$-Condition that was needed as part of the representation of (3.191.b). In the $\vee$-Condition the new state dref s was a kind of 'metavariable' ranging over the two possible quantification states s' and s'': s could take either s' or s'' as 'value'. In this way the implemented we have adopted in our syntax-semantics interface for temporal location by tense and temporal adverbs can be applied to the quantifications in the $\vee$-Condition by applying temporal location to s. For a Conjunctive DRS Condition like the one required in the representation of the upper VP of (3.209.a) such a 'metavariable' s is needed as well. But note that for such a conjunctive Condition s has to differently connected with the quantification state drefs s' and s'': s cannot be equal to the quantification state s' representing the first Duplex Condition and at the same time equal also to the quantification state s'' representing the second Duplex Condition. Instead, we get the right effect of tense and temporal adverbs on the quantification states represented by s' and s'' if we require that both s' and s'' are temporally included in s.

(3.208)a.

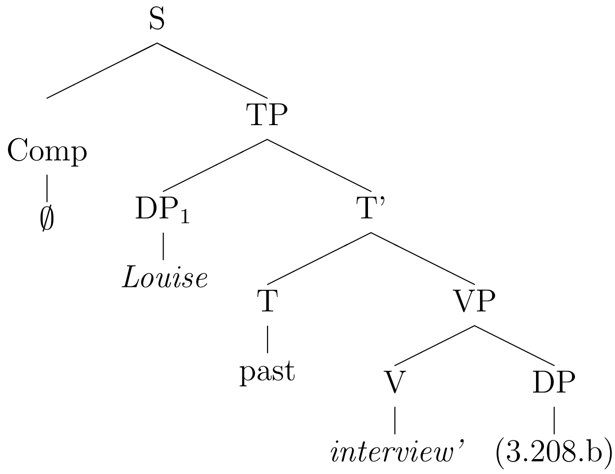
S
Comp
∅
TP
DP$_1$
Louise
T'
T
past
VP
V
interview'
DP
(3.208.b)

b.

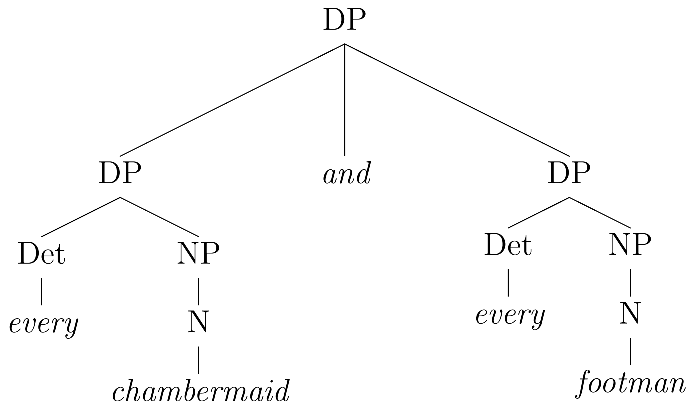
DP
DP
and
DP
Det
every
NP
N
chambermaid
Det
every
NP
N
footman

(3.209)

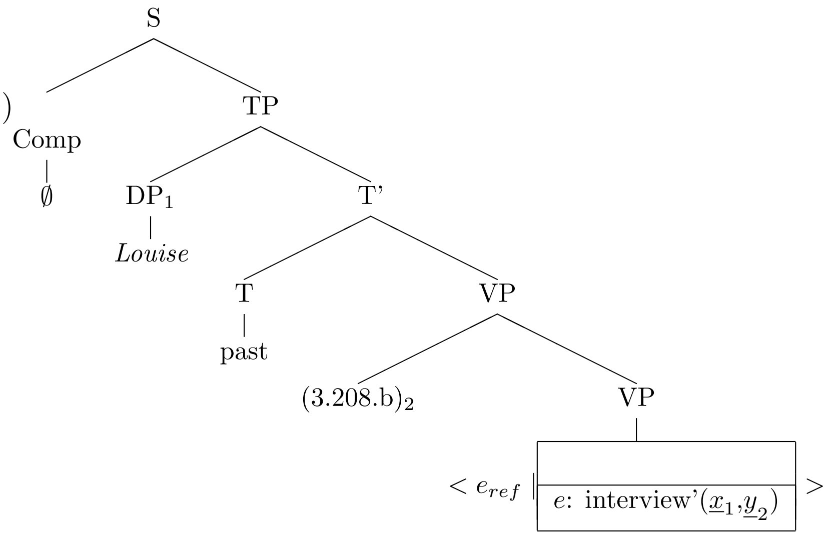
S
Comp
∅
TP
DP$_1$
Louise
T'
T
past
VP
(3.208.b)$_2$
VP
$< e_{ref}$ |
e: interview'($\underline{x}_1$,$\underline{y}_2$)
$>$

(3.210)$<s'_{ref}|\ s'$:

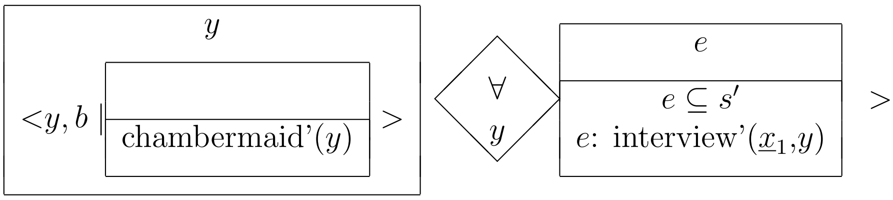

(3.211)$<s_{ref}\ |$

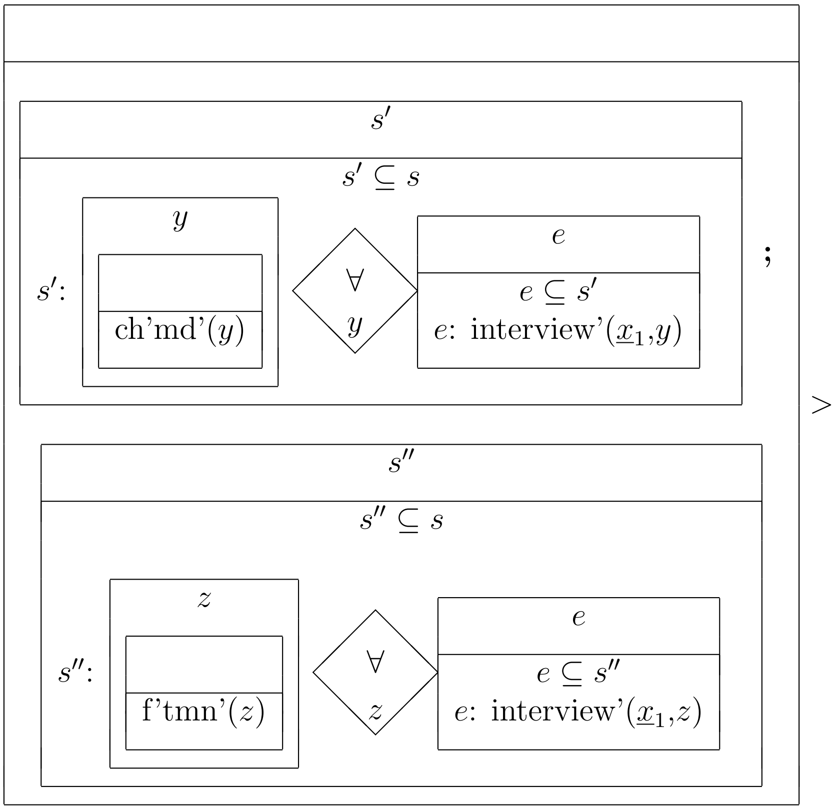

The DRS for (3.207.c).c) can now be completed in the familiar way. After completion of its construction the ;-Condition can be resolved in the sense that its constituent DRSs can be merged with the main DRS. (3.212) shows the result.

(3.212)

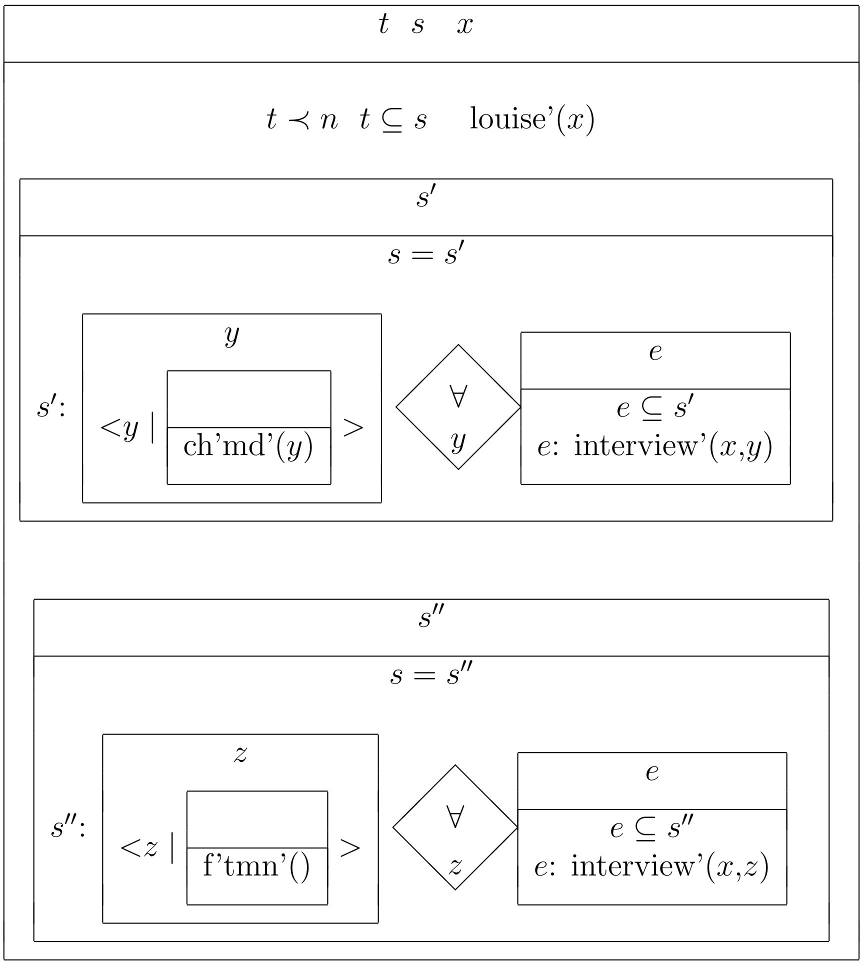

(3.207.c) differs from (3.207.a) on two counts, (i) (3.207.a) involves negation while (3.207.c) does not, and (ii) (3.207.c) has a conjunction of two quantifying DPs, whereas the conjunction in (3.207.a) is of two referential DPs. (3.207.b) which has negation, but also quantifying DPs. A semantic representation for (3.207.b) that captures the truth conditions that I described as the most prominent ones above can be constructed in almost exactly the same way as for (3.207.c), except that it must now be possible to adjoin the DP conjunction above the negation. At this point we have no clear precedent for this possibility. But we have already loosened up the constraints on waiting in this section, so I propose that we assume that such an adjunction is possible; and let us, more specifically, assume that the conjunctive DP of (3.207.b) can be raised to adjunction of NegP, as in (3.213).

(3.213)a.

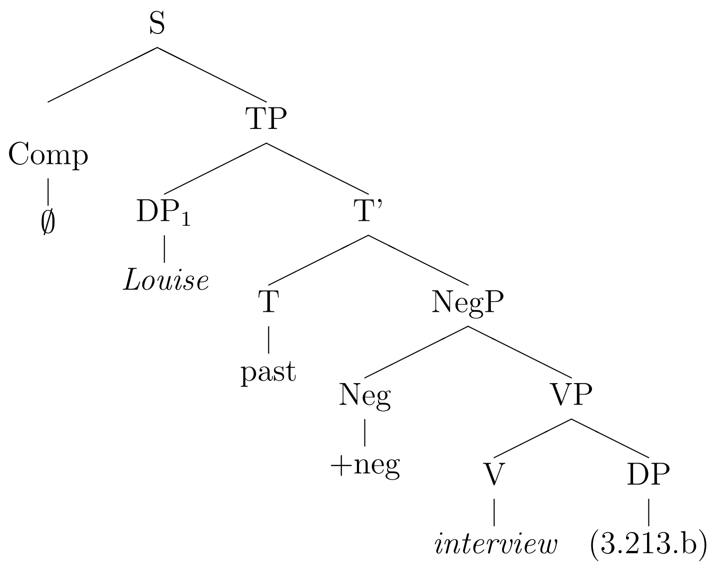

b.

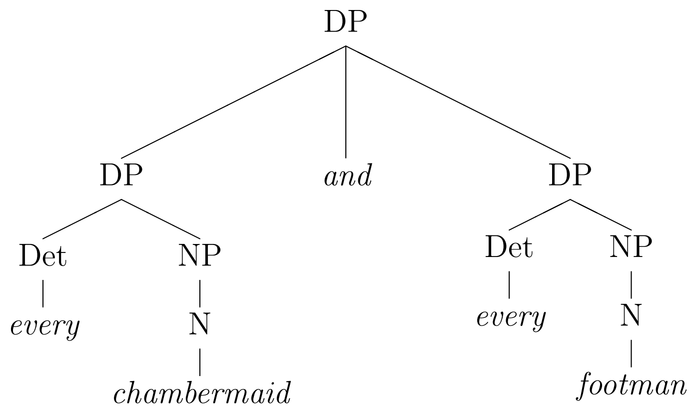

If we construct a DRS from (3.213) by applying the same rules that were invoked for the construction of (3.212), we obtain a representation that assigns the sentence (3.207.b) the truth conditions we have been aiming for.

(3.207.a) differs from (3.207.b) in having definite descriptions where (3.207.b) has universally quantifying DPs. But this difference need not be decisive. Suppose that the principle according to which DP coordinations can be raised to a position like NegP adjunct is independent of the form of the DP conjuncts, so that the direct object DP of (3.207.a) can be raised to the same position as that of (3.207.b). For (3.207.a) this, together with lexical insertion for the verb, gives us the syntactic structure:

(3.214)

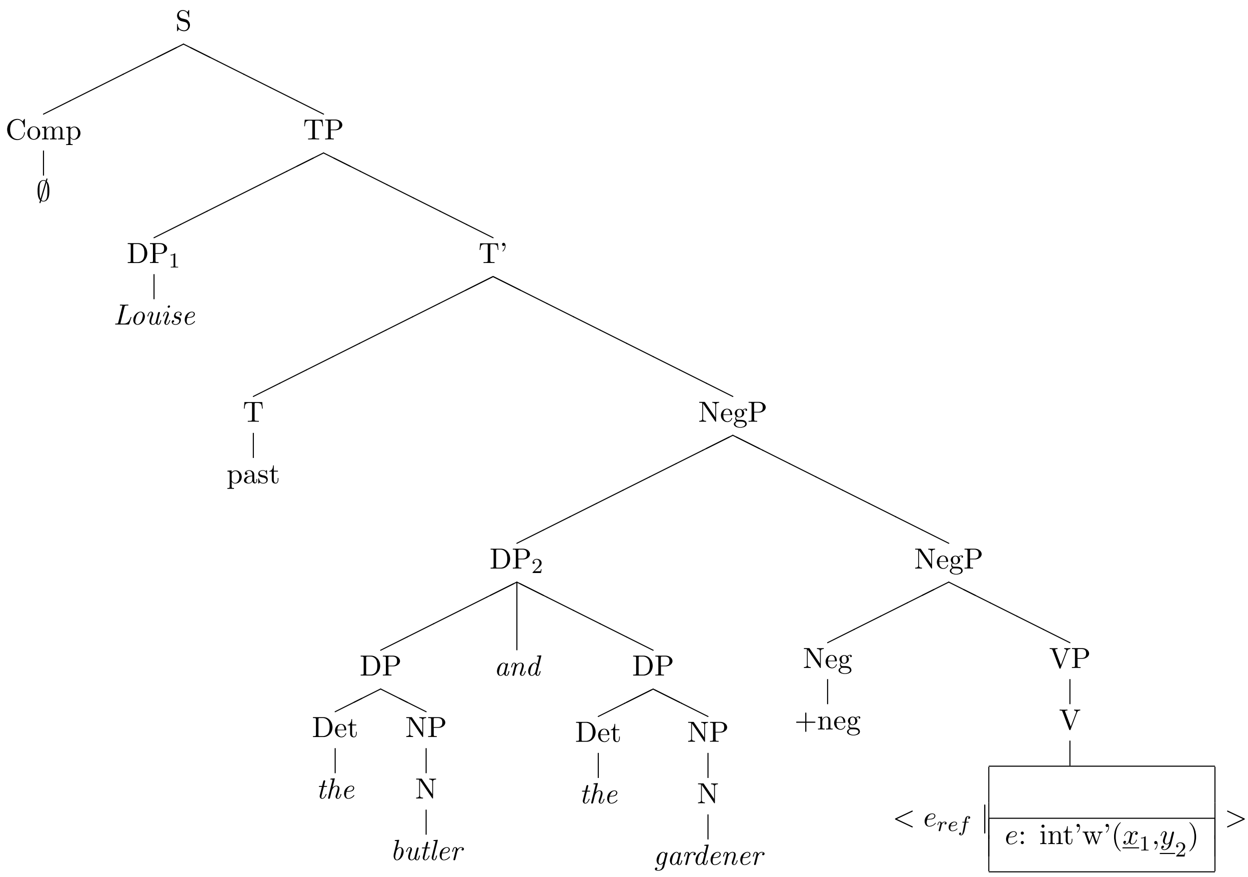

What is in this case the effect of distributing the conjuncts of the raised DP over its adjunction site? Imtuitively this should lead to the conjunction of two representations each of which is the result of inserting the referential argument of the relevant DP conjunct for $\underline{y}$ in the lexical predicate of the verb. We are not yet in a position to deal with this in proper detail because the proper treatment of definite descriptions has to wait till Section 4. But let us use the same proviso to which we have resorted more than once before, representing *the butler* by a dref y together with the 'Condition' "the-butler(y)" and *the gardener* by a dref z together with the 'Condition' "the-gardener(z)". Then the ;-Condition to which DP distribution in (3.214) gives rise is the representation in (3.215) for the upper NegP of (3.214).

$$(3.215)\quad <e_{ref},\text{y},\text{z}| \begin{array}{|c|}\hline \\ \hline \text{“the butler'}(y)\text{”} \quad \text{“the gardener'}(z)\text{”} \\ \begin{array}{|c|}\hline e' \\ \hline e' \subseteq e \\ e'\text{: interview'}(\underline{x}_1,y) \\ \hline \end{array} \; ; \; \begin{array}{|c|}\hline e' \\ \hline e' \subseteq e \\ e'\text{: interview'}(\underline{x}_1,z) \\ \hline \end{array} \\ \hline \end{array} >$$

Getting from the structure (3.214) with (3.215) as semantic representation for its upper NegP node to the DRS for the sentence is routine and left to the reader.

This was the first direction in which one may look for a justification of the 'not-interview the butler and not-interview the gardener' reading of the first sentence of (3.207.a). The second direction exploits the principle that conjunctions of referential DPs denote mereological sums. A central idea here has to do with the interpretation of verbs with arguments that are non-singleton sets, or, as in the remainder of this discussion, non-atomic individuals. English and related languages have small numbers of *plurale tantum*verbs, verbs which have argument positions that can only be filled by terms denoting non-atomic individuals. (English *gather* is an example of a verb that is *plurale tantum* with respect to its subject position.) But most verbs that can be used with argument terms that denote non-atomic individuals can also be used with terms denoting atomic individuals in the same slot. In such cases it is in principle always possible to interpret the combination of the verb with a term denoting a non-atomic individual as distributing over the atomic individuals of which that non-atomic individual is made up. Thus 'Louise interviewed the people taken into custody' can be interpreted as meaning that Louise interviewed each of those people (separately). Likewise (3.207.d), repeated below, can be interpreted as saying that Loyise interviewed the butler and that she interviewed the gardener.

(3.207.d) Louise interviewed the butler and the gardener.

In this sentence the verb *interview* occurs in combination with the non-atomic individual made up of the butler and the gardener and the distributive interpretation then leads to the reading just described.

This may work fine for an unnegated sentence like (3.207.d), but how can it get us to the interpretation we are after for the negated sentence (3.207.a)? Here is one possibility: Negated verbs can often be understood as if they were lexical verbs, with the same argument frames as the corresponding unnegated verbs. If an argument of such a ‘negation-verb’ is filled with a non-atomic individual, then it too may distribute over its components. In this way it is possible to interpret (3.207.a) as claiming that Louise din’t interview the butler and that she didn’t interview the gardener.

Formally we can construe this way of arriving at the reading for (3.207.a) that we are after as follows. The idea that the negation of *interview* can be understood as a kind of lexical operation, which turns the verb *interview* into the negated verb *not-interview*, can be captured by the following syntactic structure.

(3.216)

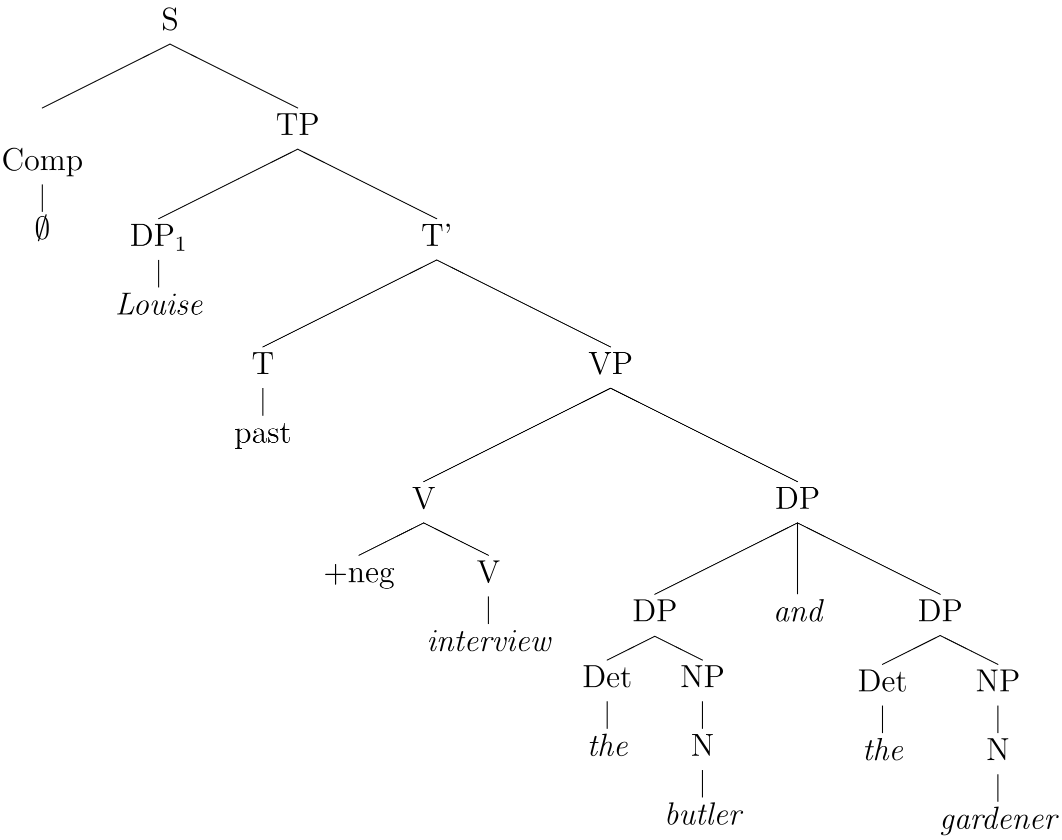

We now assume that the DP *the butler and the gardener* remains in situ and that its semantic representation is as in (3.217).

$$(3.217) < V_{ref}, y, z \mid \begin{array}{|c|}\hline \\ \hline \text{“the butler}(y)\text{”} \qquad \text{“the gardener}(z)\text{”} \\ \\ V = y \bigoplus z \\ \hline \end{array} >$$

Distributing the negated verb over the members of V is tantamount to forming the conjunction of the predication involving the negated verb and y and the predication involving the negated verb and z. That is, distribution gets us for the upper V node of (3.216) the representation in (3.218).

$$(3.218) < e_{ref}, e_1, e_2, V, y, z \mid \begin{array}{|c|}\hline \\ \hline \text{“the butler}(y)\text{”} \qquad \text{“the gardener}(z)\text{”} \\ e_1 \subseteq e \quad e_2 \subseteq e \quad V = y \bigoplus z \\ e_1\text{: ‘not-interview’}(\underline{x}_1,y) \\ e_1\text{: ‘not-interview’}(\underline{x}_1,z) \\ \hline \end{array} >$$

Using the mechanism we have introduced to represent negation in order to capture that the predication ‘not-interview’($\underline{x}_1$,y) really is the negation of the predication ‘e': ‘interview’($\underline{x}_1$,y)’ (and likewise with z in place of y) we can convert (3.218) into (3.219).

$$(3.219) < e_{ref}, e_1, e_2, V, y, z \mid \begin{array}{|c|}\hline \\ \hline \text{“the butler}(y)\text{”} \qquad \text{“the gardener}(z)\text{”} \\ e_1 \subseteq e \quad e_2 \subseteq e \quad V = y \bigoplus z \\ \\ \neg \begin{array}{|c|}\hline e'_1 \\ \hline e'_1 \subseteq e_1 \\ e'_1\text{: ‘interview’}(\underline{x}_1,y) \\ \hline \end{array} \\ \\ \text{‘}\neg \begin{array}{|c|}\hline e'_2 \\ \hline e'_2 \subseteq e_2 \\ e'_2\text{: ‘interview’}(\underline{x}_1,z) \\ \hline \end{array} \\ \\ \hline \end{array} >$$

With (3.219) as semantic representation for the VP of (3.216) the remaining steps needed to construct a DRS for the first sentence of (3.207.a) are once again familiar. The remainder of the construction is left for practice.

So much for possible ways of constructing a 'neither .. nor' interpretation for the first sentence of (3.207.a). As we noted, the sentence also has a 'not both' interpretation. Against the background of our syntax-semantics interface in the form in which in place when we started on our exploration of non-sentence level coordinations, this is the reading one would have expected, and it is curious (and probably important) that this reading isn't more prominent than it seems to be. One way to obtain this reading is to assume that the DP *the butler and the gardener* remains in situ and that distribution of the set {the butler, the gardener} is over the predicate 'interview' (and not over the quasi-verb 'not-interview').

This exploration of how (3.207.a) could have the 'neither nor' interpretation that seems to come so naturally to us suffers from the obvious shortcoming that I also drew attention to in our earlier discussion of quantifying DP disjunctions: What has been proposed in the course of trying to come up with a viable reconstruction of this interpretation has been highly speculative. What has been missing is a serious investigation of how the principles suggested work out when applied to a representative choice of other examples; and missing in particular has been any attempt to make sure that adding these principles isn't going to produce serious semantic over-generation (i.e. the possibility of constructing interpretation for sentences that they do not have).

In this regard the discussion of the sentences in (3.207) has been 'experimental' in the more pejorative sense of the term. But there is nevertheless an important moral to this tentative exploration: On the one hand the processing principles of our construction algorithm that were already in place when we started on our exploration of non-sentential disjunctions and conjunctions proved to insufficient to obtain the readings we wanted; on the other hand, when we were looking for the additional principles we needed we found that there was – at the level at which we have been carrying out our explorations – more than one way in which we could extend our system of construction rules so that it can deliver what we want.

This multiplicity of different options for reaching the intended interpretations may have an especially suspicious. 'So then, which of those options do you think is the right one?' may well have been more than one reader's reaction.

But the presupposition of this question, that there could be no more than one way in which a certain interpretation for a given reading can be obtained. There is no reason why different ways of construing certain syntactic structures semantically might not lead to truth conditionally equivalent interpretations. In fact, conjunctions and disjunctions of referential DPs may well be a case in point. Compare sentences (3.207.a) and (3.207.d) with the following sentences involving disjunctions of such DPs.

(3.220)a. Louise interviewed the butler or the gardener.

b. Louise didn't interview the butler or the gardener. But she did interview the cook and the groom.

(3.220.a) differs plainly from (3.207.d). It only has a disjunctive interpretation, according to which Louise did at least one of two things. (The possibility that she interviewed both the butler and the gardener isn't completely excluded, and can be cancelled, e.g by the addition 'and perhaps she even did both. But for reasons of implicature the exclusive interpretation is strongly preferred. For such 'Gricean implicature' effects see below.) This sets (3.220.a) with its disjunctive interpretation and (3.207.d) with its conjunctive interpretation clearly apart from each other. But for the negated first sentences of (3.220.b) and (3.207.a) this is not so. Like (3.207.a), (3.220.b) appears to have both a disjunctive and a conjunctive interpretation. The conjunctive interpretation (i.e. the 'neither nor' interpretation) is quite prominent, and in the light of what we have been saying that isn't surprising; this is what we get when we construct a semantic representation for (3.220.b) in which we apply the generalization of (3.189 for non-sentential disjunctions. For present purposes we can characterize this interpretation simply as 'in situ', meaning that the disjunctive DP is not raised and the disjunctive DP i interpreted as proposed earlier in this section (see (3.190.c)). But the disjunctive interpretation seems possible too, if perhaps somewhat less prominently and more dependently on a suitable context (such as for instance the context provided by the second sentence of (3.220.b)). Given what we have suggested about the interpretation of the sentences in (3.207) there are two ways in which this second reading could be obtained: (i) raising the coordinated DP to a position above Neg followed by $\vee$-distribution of the raised DP over its adjunction site; and (ii) reinterpreting the negated *interview* as a kind of lexical verb ('lexical negation').

(3.207.a) presents us with a somewhat different picture. The 'disjunctive' reading of (3.207.a) – or its 'not both' reading if you prefer – can according to what we have said be obtained in only one way, viz. by keeping the conjunctive DP in situ and then distributing the set {the butler, the gardener} over its predicate *interview*. (In this respect this reading for (3.207.a) corresponds to the 'neither nor' reading for (3.220.b).) The other reading for (3.207.a) – its 'neither nor' reading – can according to our speculations also be realized in one of two ways: (i) by raising the conjunctive DP to a position above the negation and applying *and*-distribution over the adjunction site; (ii) by keeping the DP in situ, but interpreting *not-interview* as a lexical verb.

The upshot of this is that according to the construction possibilities we have found reason to assume are available for some sentences (i) both (3.207.a) and (3.220.b) are ambiguous between a reading in which their coordination operator takes wide or narrow scope with respect to the negation; and (ii) that two of the possible readings can be obtained in more than one way.

3.9.5.8 Summing up; some further Methodological Implications

Here are some conclusions that can be drawn from these explorations of non-sentential disjunctions and conjunctions:

(1) The semantics of sentences with non-sentential occurrences of *and* and *or* turns up a host of problems that do not arise for their uses as sentence connectors. We have only fastened on a few of these and just dealing with these turned out to contain a range of challenges. And there is no reason that I can see that the issues that have been raised here come even near to giving a reasonably complete picture of the problems that are connected with other occurrences of non sentential coordinations.

But just the examples we have been trying to come to grips with have made it clear that for many fundamental syntax-semantics interface assumptions non-sentential coordinations present a serious challenge. We have seen is quite concretely with regard to the interface we have been developing in these Notes. But I believe this is equally true for other currently familiar syntax-semantics interface approaches.
(2) Contrary to a certain tendency in the current semantic literature, where a lot of attention is being paid to 'unexpected' interpretations of *or*, the perspective from which we have been looking at coordination here suggests that *and* is the more challenging of the two. This is because conjunctions of certain non-sentential categories allow for interpretations as mereological

sums and addition to the interpretations one gets when distributing local coordinations over other constituents in the sentences in which they occur.

(3) We have observed a notable difference between the semantics of conjoined categories in the nominal and in the verbal domain. The semantics of conjunctions of verb projections gives a quite complicated picture, we saw, with a many-stationed spectrum from completely independent eventuality described by the conjuncts via eventualities of which the connects describe different, and often differently connected parts to cases where the conjuncts can only be construed as different descriptions of one and the same event. And this range of options seems to apply for all syntactic projection categories of the verb – all those that in our syntax-semantics interface play the role of eventuality descriptions.

For nominal categories the picture is quite different. Here there appears to be a sharp division between categories that function as nominal predicates – these are nearly all the nominal syntactic categories – and the category DP, the category of expressions that contribute arguments to both verbal and nominal predicates. For the former conjunction is predicate conjunction, for DPs the semantics of conjunction depends on what the form of the conjunct DPs. A particular important case is that where the conjuncts are 'referential DPs', DPs that are naturally understood as referring to atomic or non-atomic individuals. For DP conjunctions of this latter kind the semantics is that of mereological sum.

Conjunctions of referential DPs play a part in the syntax-semantics interface of non-sentential coordination that is unlike that of any of the other forms of non-sentential coordination nadine that considerably complicates the general picture. But such conjunctions are extremely common and natural, so no serious theory of non-sentential coordination can afford to ignore them.

(4) An important concomitant feature of the semantics of referential DP conjunctions is their interaction with the predicates to which they can be arguments: Prominent among the possible interpretations of such predicate argument combinations is the distributive interpretation which was central to our discussion of conjunctive DPs like *the butler and the gardener*. Distribution is only one of a number of ways in which set-like arguments can semantically interact with predicates. There are also several other forms of interaction. All those are 'collective' in the sense that they cannot be reduced to interactions between the predicate and individual members of the set occupying an argument position of it. (Collective interpretations are

found mostly with verbs. I do not know of a good phenomenology of the different forms that collective interpretations can take. But this is an issue that belongs to the general topic of plurality and therefore not here, even if we have been forced into discussing some aspects of the theory of plurality in this present section.
(5) Another problem, we have seen, that arises primarily for coordinated DPs is the structural position form where they ale their semantic contribution. In the discussions above I didn't see how certain evidently available interpretations could be explained without assuming that such DPs can be subject to certain forms of DP raising. Already in PART I of the Noted we adopted Quantifier Raising as part of our syntax-semantics interface. But the reasons for doing so – having ultimately to do with the behavior of the so-called inversely linked quantifiers – are quite different from those that led us to adopt the option of DP raising in this section. These kinds of raising were always needed for coordinated DPs and it isn't even clear that they should be adopted also for non-coordinated DPs (although it would be odd, if such movements were not allowed in the absence of coordination).

it isn't clear at this point if the movements of coordinated DPs that we have relied on in some of our reconstructions should be seen as part of what happens at the level of syntax. Perhaps structural rearrangements of the kind illustrated bt the coordinated DP raisings we have made use of are also possible as part of semantic interpretation. but exactly what that should come to isn't clear to me. This way or that, the assumptions we have been making in this section about coordinated DP raising are among the most problematic from the point of view of the kind of syntax-semantics interface developed in these Notes, and I feel especially unhappy about this aspect of the proposals made in this section, and about leaving this aspect in the unsatisfactory state in which I am leaving it.

3.9.5.9 Non-sentential Coordination in Montague Grammar: Too good to be true

Montague Grammar offers an attractive account of non-sentential conjunction and disjunction. This account makes use of the notion of *type*. Among these types there are the 'purely extensional' ones. These are built from the basic types e (for 'entity') and t (fort 'truth value'). Complex types are built recursively from these basic types by forming ordered pairs: if α and β are types, then so is $<\alpha, \beta>$. Well-formed expressions of a natural language such as English are always of one such type, which is determined by their syntactic category. The semantics that comes with the types is an implementation of the general principle that the meaning of an expression of type $<\alpha, \beta>$ are functions from entities of type α to entities of type β (where those entities of type α or β will often will themselves be functions). Furthermore,expressions of the basic types e and t denote entities and truth values respectively, as implied above. This means in particulars that proper names are type e and sentences of type t. The meanings of expressions of type $<e,t>$ are functions from entities, so such expressions can be thought of as 1-place predicates of individuals, and so on. The polymorphism of *and* and *or* consists according to this account in that there are different versions of them for all types that 'end in t'. A type that 'ends in t' is one that is either t itself or that leads to t by successively taking away first members. So t, $<t,t>$, $<e,t>$, $<<e,t>,t>$, $<<e,t>,<<e,t>,t>>$ are all types that end in t. The versions of conjunction and disjunction for type t are the familiar sentence connectives & and $\vee$ of classical logic whose semantics is given by the standard truth tables. The versions of & and $\vee$ for complex types $<\alpha,\beta>$ are defined in terms of the versions for α and β. For instance, $\&_{<\alpha,\beta>}$ is the function which maps two entities $\mathcal{F}$ and $\mathcal{G}$ of type $<\alpha,\beta>$ to the entity $\mathcal{F}\ \&_{<\alpha,\beta>}\ \mathcal{G}$ that is defined by the condition that for any Q of type α, $(\mathcal{F}\ \&_{<\alpha,\beta>}\ \mathcal{G})(Q) = \mathcal{F}(Q)\ \&_{\beta}\ \mathcal{G}(Q)$. More specifically, when $\alpha = e$ and $\beta = t$: $\&_{<e,t>}$ is the function which maps any two entities P and Q of type $<e,t>$ to the function $P\ \&_{<e,t>}\ Q$ that maps any type e entity d to the truth value given by $P(d)\ \&_t\ Q(d)$, where '$\&_t$' is another notation for the sentence connective &.

This account of the polymorphism of *and* and *or* goes a fairly long way to account of the data, which in view of its elegance and simplicity is a remarkable fact. But as our exploration above have rvealed, it cannot go far enough all on its own. The interpretations that the account assigns to particular occurrences of*and* and *or* are fixed once their logical type has been determined,and it is one of the standard assumptions of MG that these are determined by the syntactic categories of the expressions that the given occurrences combines as coordinates. So there is no room for ambiguity in the account itself. Therefore, if a sentence is ambiguous and this ambiguity has to do with the contribution made by *and* or *or*, then the ambiguity must either be an ambiguity of the input to the account – i.e. it must be a matter of syntactic ambiguity – or the ambiguity must be the result of further interpretational processing of the computed semantic content. As far as syntactic ambiguity is concerned the MG approach might in principle not be worse off than the proposal made above, which also postulate syntactic ambiguities as part of explaining why sentences with non-sentential conjunctions or disjunctions can be ambiguous and how. But I do not know how a combination of a theory of the syntactic ambiguity of the kind of sentences we have been looking at might be combined with the MG account in a way that will make the right predictions about the kinds of sentences we have been looking at, but haven't looked at the matter in sufficient depth to feel confident about this.

A second strategy for dealing with apparent ambiguity is to locate its source at the level of post-semantic processing. The general method is indelibly connected with the name of Grice. Grice was the first to see that complex data about interpretation of sentences in use which superficially seemed to contradict the logical account of of words like *or* and *and*should not be taken to contradict a logo-based approach to semantics, but instead can often be explained by an elegant combination of such a semantics and an account of how the use of certain sentences with a given semantic content can convey additional content because the speaker reveals through their use information without which she could not have used the sentence legitimately in the given context. One type of application of this general strategy is to sentences with occurrences of *or*. Of some such sentences it can be argued that their legitimate use in a given context or for a certain purpose entails that more must be true than just their semantic content. One type of example, which has preoccupied the semantic community for decades are 'permission sentences' like (3.221).

(3.221)You may take an apple or a pear.

When (3.221) is used as permission granting utterance – to extend a permission to some addressee over whom the speaker has the relevant authority – then the sentence has the force of extending both the permission to take an apple and the permission to take a pear. For a theory that takes the semantics of this sentence to be that it is permitted to the addressee to make true the proposition that he takes an apple or a pear this may well seem a surprising result, for there are all sorts of ways in which this proposition can be made true and nothing tells us for instance that it may be made true by taking an apple (as opposed to taking a pear). A fully satisfactory solution to this problem seems to have been eluding the community and from what I know about the literature that is still the case, after about half a century of worrying and proposals.

Nevertheless I don't think there can be any doubt that a large part of the solution has to do with the pragmatics of sentences like (3.221). Important for the effect just described is whether the sentence is used to issue a permission or to report to someone else what one believes to be permitted to him. On the other hand, however, there is also the curious fact that (3.221), with its disjunctive DP, seems the most natural form to use when one's purpose is to issue a permission. Alternative forms, like (3.222.a) and (3.222.c), can also be used for the purpose of issuing permissions. But (3.222.c) seems quite unnatural when used in this capacity and may be slightly unnatural when used as permission report. ((3.222.b) seems to be only marginally well-formed, for reasons that I do dare to pronounce on.)

(3.222)a. You may take an apple or take a pear.

b. You may take an apple or may take a pear.

c. You may take an apple or you may take a pear.

Sentences like those in (3.221) and (3.222) may not differ much in their content and use. But here too a full account should be able to deal in detail with the syntax-semantics interface for the different non-sentential and sentential disjunctions they contain. I am stressing this point for two reasons which could think of as each other's mirror image. On the one hand work on the pragmatics of *or* should be pay close attention on the syntactic categories of the disjunctions involved; but by the token accounts of the syntax and syntax-semantics interface for such sentences must also keep firmly in mind that there are aspects of the sermonic properties of *and* and *or* that need a pragmatic explanation and about which syntax and semantics as considered in this section (and in these Notes generally) cannot have anything to say.

3.9.5.10 Biconditionals

The biconditional is the last of the different connectives that in presentations of the propositional or predicate calculus are commonly included among the primitive sentence connectives . As the term 'biconditional' suggests, a biconditional $A \leftrightarrow B$ is true if and only the conditionals $A \rightarrow B$ and $B \rightarrow A$ are true. When the biconditional is not included among the primitives, it is as the conjunction of these two conditionals that it is typically defined.

DRT is an awkward representation formalism for the biconditional. Since it has a representational format for the conditional, in the form of the conditional DRS Condition $K \Rightarrow K'$, where K and K' are DRSs representing the antecedent and the consequent of the represented conditional, one might have thought that the biconditional could be defined as the conjunction of two such Conditions, as in (3.223).

(3.223)
$$\begin{array}{|c|}\hline \quad \\ \hline K \Rightarrow K' \\ K' \Rightarrow K \\ \hline\end{array}$$

But unfortunately this won't do in general. The reason is that discourse referents in the Universe of K or K' have a different semantic function depending on whether their DRS occurs as antecedent of a conditional Condition or as its consequent. For instance,

$$\begin{array}{|c|}\hline \quad \\ \hline
\begin{array}{|c|}\hline x \\ \hline P(x) \\ \hline\end{array} \Rightarrow \begin{array}{|c|}\hline y \\ \hline Q(y) \\ \hline\end{array} \\[1em]
\begin{array}{|c|}\hline y \\ \hline Q(y) \\ \hline\end{array} \Rightarrow \begin{array}{|c|}\hline x \\ \hline P(x) \\ \hline\end{array} \\ \hline\end{array}$$

says that for every x such that $P(x)$ there exists a y such that $Q(y)$ and that for every y such that $Q(y)$ there exists an x such that $P(x)$: x is acts as a universally quantified variable and y as an existentially quantified one in the first Condition and y as universally and x as existentially quantified in the second.

To make sure that the drefs in the Universes of K and K' make the same logical contributions in each of the two Conditions we have to shield them from this source of variance in quantificational force by embedding K and K' more deeply, as for instance in (3.224).

(3.224)
$$\begin{array}{|c|}\hline \quad \\ \hline
\begin{array}{|c|}\hline \quad \\ \hline \begin{array}{|c|}\hline \quad \\ \hline \emptyset \Rightarrow K \\ \hline\end{array} \Rightarrow \begin{array}{|c|}\hline \quad \\ \hline \emptyset \Rightarrow K' \\ \hline\end{array} \\ \hline\end{array}
\quad
\begin{array}{|c|}\hline \quad \\ \hline \begin{array}{|c|}\hline \quad \\ \hline \emptyset \Rightarrow K' \\ \hline\end{array} \Rightarrow \begin{array}{|c|}\hline \quad \\ \hline \emptyset \Rightarrow K \\ \hline\end{array} \\ \hline\end{array}
\\ \hline\end{array}$$

(Here $\emptyset$ denotes the empty DRS $\begin{array}{|c|}\hline \quad \\ \hline \quad \\ \hline \end{array}$.)

In (3.224) all drefs occurring in U_K or $\mathrm{U}_{K'}$ behave as existential quantifiers. so both K and K' make the same semantic contribution whether they occur on the left of the main $\Rightarrow$ or on the right.

However, most definitions require more than a biconditional with the satisfaction conditions of (3.224). Definitions typically are universal quantifications of such structures. Consider for instance the definition of humans as featherless bipeds. (A rather curious definition, which trades on our good sense to ignore plucked chickens. But for logical and linguistic purposes that doesn't matter, and it is nicely (if atypically) simple.) This definition wants to say that it is true of anything whatever that it is human if and only if it is a featherless biped. If we want to express this in a DRS language using (3.224), then K and K' will be 'improper' DRSs, which contain a dref in their conditions that is 'free' in them by not belonging to any of their DRS Universes. This DRS must be bound outside, and in such a way that it universally quantifies over the structure in (3.224). The simplest way to achieve this effect is to use a Duplex Condition in which the two conditional Conditions from(3.224) form the Condition Set of the nuclear scope and the restrictor is essentially empty, as in (3.225).

(3.225)

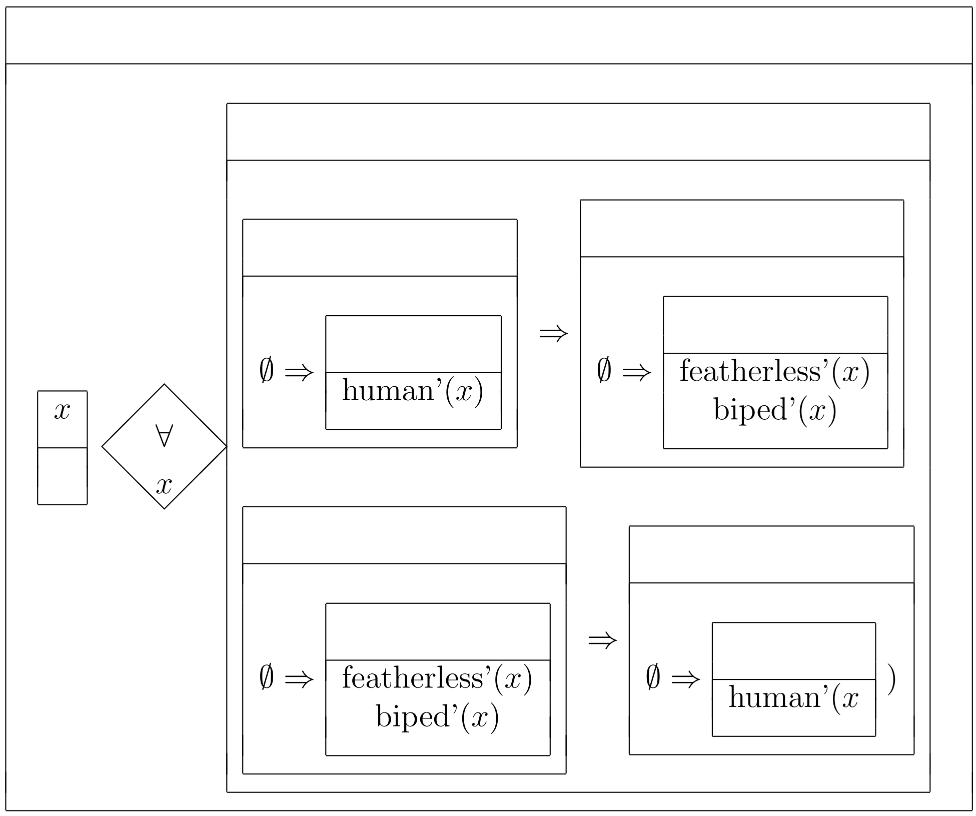

This looks dreadfully cumbersome, but once the template, instantiated by (3.225) and shown below in (3.226), has been defined, users can employ it while closing their eyes to the seemingly redundant formal complexities. If we want to, we can give a structured name to the template, e.g. as ‘Def(K,K',$<\alpha_1,..,\alpha_n>$), where K and K' and $<\alpha_1,..,\alpha_n>$ (with n $\geq$ 1) are as shown in (3.226). (To use standard terminology, K is the *definiendum* in this definition and K' the *definiens*.)

(3.226)

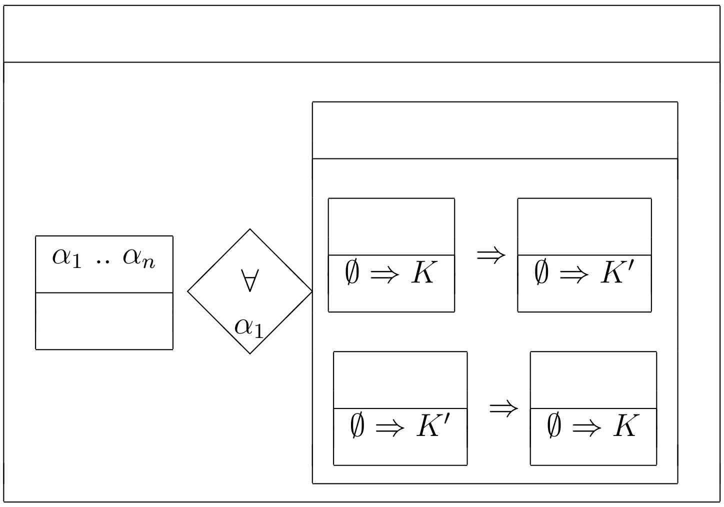

Definitions in Natural Language

So much for how definitions can be stated within a DRS language. But more interesting for the linguist is how definitions are expressed in natural languages. In English there are various ways of doing this, and the form chosen will depend to some extent on the details of the actual definition. We won't go into a systematic investigation of this, but consider just one example. (3.227) is a natural way of formulating the definition according to which humans are all featherless bipeds and are all the featherless bipeds.

(3.227) Something is human if and only if it is a featherless biped.

There are several aspects to this sentence that merit a comment. First, there is the expression *if and only if*. In mathematics, logic, philosophy and linguistics, and presumably in other branches of science as well, *if and only if* has become a kind of idiom for expressing equivalences, including those that are essential to explicit definitions. But whether or not it is to be called an idiom, it is an expression whose meaning is largely determined by general compositional principles.

However, its compositional meaning is not one that can be taken for granted on the basis of the syntax-semantics interface that we have developed up to this point. In the following discussion we won't put all the principles in place that would be needed to carry out a detailed step-be-step construction of the semantic representation of this sentence; but I will give pointers to the principles that would be needed for such a construction.

1. *if and only if*

The first point is that the *and* of *if and only if* is a case of non-sentential conjunction, as discussed in the section on Polymorphism. This is one of those cases where the result of the conjunction – the conjunction *if and only if* – has a meaning that comes close enough to the biconditional of formal logic to serve as its natural language stand-in: 'A if and only if B' is true if either A and B are both true or else are both false.

The second point concerns the role of *only*. Intuitively, 'A only if B' is taken to express the converse of what is expressed by 'A if B'. But how can it mean that? How can *only* make this particular difference? The answer to this question is contained in the analyses of *only* provided by Alternatives Semantics ((Rooth 1985), (Beaver & Clark 2008)). Accounts of *only* according to Alternatives Semantics vary in their details, but a common core is that an occurrence of *only* in a sentence S determines an *associated focus*. The focus associated with an occurrence of *only* is some syntactic constituent of S. And what that comes to is roughly this. When a syntactic constituent Γ of S is the associated focus of *only* in S and if S' is the sentence obtained by eliminating *only* from S, then for S to be true S' itself must be true and any sentence S" obtained from S' by substituting some 'alternative' constituent for Γ in S' must be false. What this amounts to in detail depends crucially on what the set of Alternatives to Γ are. More will be said about this below.

Which constituent of a sentence S is the associated focus of an occurrence of *only* in S is sometimes fully determined by the syntactic position that *only* occupies within S, and in other cases it is determined by that position together with the prosodic contour with which S is uttered. The context in which S is used (e.g. when S is produced as answer to a preceding *wh*-question) often plays a part as well. But there are also cases where it is not indeterminate which constituent is the associated focus of *only*.

Here are a few examples. The first is from Rooth ((Rooth 1985)), the other two are close variants.

(3.228)a. Mary only introduced Bill to Sue.

b. Mary introduced only Bill to Sue.

c. Only Mary introduced Bill to Sue.

In (3.228.b) and (3.228.c) the focus associated with *only* is determined by word order alone. The focus is the DP immediately following *only*, i.e. the proper name *Bill* in (3.228.b) and name *Mary* in (3.228.c). In (3.228.a) word order isn't enough to determine the focus, which could be either (i) *Bill* or (ii) *Sue* or (iii) the Verb constituent *introduced* or (iv) the VP *introduced Bill to Sue* or, less prominently, the subject *Mary*. When (3.228.a) is spoken, then prosody will in many cases suffice to disambiguate between these possibilities. Focal stress on *Bill* will select it as the associated focus; and focal stress on *introduced* singles out the Verb. But if stress on *Sue* is to determine *Sue* as the focus associated with *only*, it will have to be recognizably different from default stress on *Sue*. In the practice of spoken language such a distinction is often hard to detect. Furthermore, when the speaker intends the focus to be the verb phrase *introduced Bill to Sue*, then this too may lead to focal stress on *Sue*, so even when there is a perceptible extra stress on *Sue*, this last ambiguity – between the focus being *Sue* and it being *introduced Bill to Sue* – won't be resolvable by prosody alone.[58]

The associated focus Γ of an occurrence of *only* in a sentence S determines in its turn an *Alternatives Set* AS(Γ). In some cases this set can be identified as a set of entities, one of which is the denotation of Γ. The other entities in the set must then be entities 'of the same kind' as this denotation, but exactly what 'of the same kind' can amount to has been a matter of debate. For one thing the answer to this question depends on what sort of expression Γ is. (For instance, what could AS(Γ) be when Γ is *white* in the sentence 'There are only whíte balls in the bag', or when Γ is *tennis* in 'There are only ténnis balls in the bag'?) The simplest and clearest examples of what AS(Γ) can be (and the cases on which much of the discussion of focus association in the literature has concentrated) are those in which the associated focus is a definite DP, and more particularly where it is a proper name. *Bill* in (3.228.b), *Mary* in (3.228.c) and any of *Bill*, *Sue* or *Mary* in (3.228.a) are examples of this. In these cases the Alternatives Set can be taken to be a set of individuals. The denotation of Γ – the person Bill, Mary or Sue in the examples considered – will be one of the members of the set AS(Γ), but what else should go into this set? All individuals in the entire world? That doesn't sound very plausible. It is natural to think for instance of (3.228.b) as a statement that is made about some particular event – a party, say, or a

[58]There is also entirely different interpretation of *only* in (3.228.c), in which it plays a very different part from the one that is relevant to the present discussion. This interpretation of the sentence can be phrased as 'It is only that Mary introduced Bill to Sue', or as something like 'But we shouldn't forget that Mary introduced Bill to Sue'. When *only* is used this way, it doesn't have any associated focus. Please, set this use of *only* aside for the remainder of the present discussion.

reception. In that case the more natural choice for AS(Γ) would be the set of people present at that event. (This choice renders (3.228.b) equivalent to 'Bill was the only person at the event that Mary introduced to Sue', and that seems about right.) In general the intuitively plausible choices for AS(Γ), in those cases where Γ is a DP, is subject to this and other kinds of contextual restrictions. There is an ineradicable element of vagueness in the context dependence of AS(Γ), but for the most part the vagueness is harmless insofar as it is immaterial to the truth conditions of the sentence exactly which set AS(Γ) is taken to be. For instance, the message conveyed by (3.228.b) is intuitively to the effect that there was no relevant introduction at the event in which Mary was the introducer and Sue the beneficiary, apart from the one in which Bill was the introducee; and for that message to come across many different choices for AS(Γ) will do.

To sum up the discussion of *only* up to now, we have encountered two sources of ambiguity for sentences with *only*: (i) What is the focus Γ associated with *only*? (ii) Once Γ has been determined, what is the set AS(Γ)? But note well that these are quite different kinds of ambiguity. The first, concerning the choice of Γ, is discrete and the set of choices – when there is a choice at all – tends to be quite small (just a few constituents of the sentence in which *only* occurs). The second ambiguity tends to be open-ended, is nearly always there and mostly there are no clear limits to the different ways it can be resolved (i.e. to the different choices for AS(Γ) that can be made).

In cases where the associated focus is an individual-denoting expression the focus community has by now secured a reasonably good grip on what Alternatives Sets can and cannot be. But when the associated focus is an expression of some other syntactic category, matters are different and on the whole much murkier. The associated foci *white* and *tennis* noted above are telling examples. For another example, consider the interpretation of (3.228.a) according to which the associated focus of *only* is the verb *introduce*. What could the Alternatives Set AS(Γ) be in this case? Intuitively it should contain, besides the relation expressed by *introduce* itself, other relations that consist in one person establishing a relation between two other persons (in the sense in which Mary in (3.228.a) establishes a relation between Bill and Sue by introducing Bill to Sue). Candidates for such alternative members of AS(Γ) might be the relation of Mary praising Bill in front of Sue, or of her recommending Bill to Sue as a possible date. But exactly which relations should belong to AS(Γ) seems hard to decide and, it appears, more so than when the associated focus is a DP like *Bill*. All that (3.228.a) appears to be saying (given that *introduce* is taken to be the associated focus) is that 'what

Mary did was to introduce Bill to Sue; and she did no more than that to the two of them'. What that 'more' remains largely implicit.

The case that triggered this discussion of the semantics of *only* is one where the associated focus of *only* is an *if*-clause. To understand what the force of *only* could be when its associated focus is of this kind, we first have to reflect on the function of *if*-clauses in sentences without *only* such as, for instance, the conditional 'if x is a featherless biped, then x is human'. *if*-clauses qualify the claims made by the main clauses to which they are attached: the claim is made provisionally on the assumption that the *if*-clause holds. But what does that mean? There is wide agreement that in general it means more than just that either the claim made by the *if*-clause is false or the claim made by the main clause is true. Rather, the *if*-clause should provide a basis for concluding that the main clause must be true, and that guarantee should hold across a set of different epistemically possible situations. The way this is commonly put in more formal terms is that it is true for each of a range of situations, or possible worlds, or times, that if the *if*-clause is true in any of them, then so will be the main clause. This assumption gives rise to a family of different accounts of the truth conditions of *if*-clause-main clause sentences. One dimension along which these differ from each other has to do with whether their semantics is articulated in terms of situations, possible worlds or times, and on how these sets are determined by form and content of *if*-clause and man clause, as well as, often, the context in which the utterance is produced. For present purposes I will assume that the set is a set W of possible worlds, i.e. that we are talking about the truth of *if*-clause and main clause at different worlds belonging to W.

It is in this sense, then, that we will analyze *if*-clause-main clause combinations '*if* A, B' in what follows: Both A sand B are assumed to have a truth value in each of the worlds in W. Let $[[A]]$ be the set $\{w \in W$: A is true in $w\}$. We will refer to this set as the 'proposition relative to W expressed by A'. Likewise, $[[B]]$ - the 'proposition relative to W expressed by B' – is the set $w \in W$: B is true in $w\}$. We specify the truth conditions of *if* A, B by stipulation that *if* A, B is true (in any world w') iff $[[A]]$ $subseteq$ $[[B]]$.[59]

[59] This is one of the few places in these Notes where we make a small foray into the realm of the non-extensional. (Earlier forays were made in the discussion of intensional operators like PROG in Section 3.5.2 and of non-actual futures in Section 3.7.1.) The intensional truth conditions for conditionals adopted above constitute a considerable simplification in comparison with accounts – there are quite a few of those – that can be found in the current literature. For the purposes of ht present discussion our simple account will do fine.

The conditional to which our analysis of *only if* is to be applied as part of our discussion of (3.227) is of the form ‘if x is a featherless biped, then x is human’. This is a conditional in which x occurs as a free dref. That causes slight complications, so let us before turning to it first consider an example in which this particular difficulty does not arise.

(3.229) If John is blue-eyed and Mary is blue-eyed, Billy is blue-eyed.

We can represent this conditional as a DRS with a conditional Condition in which the conditional operator is now interpreted intensionally, in accordance with the propositional inclusion relation between antecedent and consequent which we adopted above. The DRS is given in (3.230). The conditional operator $\Rightarrow_{int}$ in (3.230) has been given the subscript $_{int}$ to distinguish it from the extensional conditional operator $\Rightarrow$ that we have been using up to now. To strip down the discussion to its essentials reference to time has been eliminated in (3.230). For instance we use ‘blue-eyed’(j)’ to represent the information that John is blue-eyed, instead of a condition like ‘s: blue-eyed’(j)’.

$$(3.230)\quad \begin{array}{|c|}\hline j\quad m\quad b \\ \hline \text{John'}(j)\quad \text{Mary'}(m)\quad \text{Billie'}(b) \\ \begin{array}{|c|}\hline \\ \hline \text{blue-eyed'}(j) \\ \text{blue-eyed'}(m) \\ \hline\end{array} \Rightarrow_{int} \begin{array}{|c|}\hline \\ \hline \text{blue-eyed'}(b) \\ \hline\end{array} \\ \hline\end{array}$$

Presupposed by the use of $\Rightarrow_{int}$ is that th DRSs that this operator combines can be evaluated for truth or falsity at different worlds of W. Consider for instance the antecedent DRS K_{ant} of the conditional Condition in (3.230). This DRS will be true in a world $w \in W$ if the individuals John and Mary represented by j and m both have the property ‘blue-eyed’ ’ in w. In accordance with the notation adopted above we denote this set of worlds as $[[K_{ant}]]$; likewise the corresponding set $[[K_{con}]]$, for the consequent DRS K_{con} of the conditional Condition in (3.230), is denoted as K_{con}. (3.230) is true, then, (in any word w') iff $[[K_{ant}]] \subseteq [[K_{con}]]$.

From this way of thinking about *if*-clause-main clause sentences it is only a short step to thinking of the main clause as representing a claim that could be made conditional upon any one from a set of possible assumptions, and where any antecedent chosen from this set provides one of the possible conditional claims with the main clause as consequent. This conception is the key, I want to suggest, to what *only* contributes when it combines with the *if*-clause of an *if*-clause-main clause combination: *only* selects the *if*-clause as its associated focus – this is one of the constructions where the syntax alone determines which constituent the associated focus is – and, given this choice of focus, its contribution then is that if the *if*-clause of the *if*-clause-main clause combination is replaced by another member of the Alternatives Set, then the resulting conditional claim is false: the alternative antecedent does not guarantee the truth of the consequent.

But what does 'guaranteeing the truth of the consequent' come to precisely? That depends on what is in the set AS(Γ). And as far as that is concerned, the semantics of *only* when it combines with the *if* of an *if*-clause appears to be subject to a certain constraint, which does not apply to uses of *only* of the kind found in (3.228). The constraint is that if in any of the relevant worlds – for us: any world in W – in the *if*-clause of an *if*-clause-main clause sentence is not true, then the main clause is not true in that world either.[60]

How can we secure this effect of *only* in *if*-clause-main clause sentences – that failure of the *if*-clause dental failure of the main clause – by imposing suitable constraints on the Alternatives Set? I can see more than one way to do this. The one I will adopt here may in the end be not be conceptually optimal. but it is easy to state and explain, and since it is one that for our for present purposes does well enough, it is the one I adopt. The constraint is to the effect that if any world w in W is not covered by the focus proposition

[60] We also find this constraint in other uses of *only* in which its focus specifies a condition under which the remainder of the sentence is true. The effect of adding *only* to such a sentence in which it select the constituent in question as associated focus is to produce the claim that the main clause is not true when the condition imposed by the focused does not hold. An example of this that does not involve an overt *if*-clause is shown in (3.231).

(3.231)a. This bolt can be loosened with a special ratchet.

b. This bolt can only be loosened with a special ratchet.

The focus associated with in (3.231.b) is the PP *with a special ratchet*. In (3.231.a) this PP sates a condition under which it is possible to loosen the bolt. The effect of adding *only*, as in (3.231.b), is to turn the sentence into the claim that is not possible to loosen the bolt *without* the ratchet.

(i.e. the if the *if*-clause is not true in w), then the proposition that the world is w, i.e. the singleton set $\{w\}$, is a member of the Alternatives Set. This constraint is stated as point 2. in (3.232), which summarizes the effects of *only*when it syntactically modifies the *if* of an *if*-clause.

(3.232)Let S be a sentence of the form 'Only if A, B', where A and B are clauses.

1. The focus associated with *only* is the proposition $[[A]]$.

2. For each world w from W that does not belong to $[[A]]$, $\{w\}$ is a member of the Alternative Set AS($[[A]]$).

3. 'Only if A, B' is true (in any world) only when B' does not hold on any condition in AS($[[A]]$) that is distinct from $[[A]]$. That is, for no such proposition C in AS($[[A]]$) distinct from $[[A]]$ do we have: $C \subseteq [[B]]$.

It is easy to see that (3.232) entails that when w is a world in which the subordinate clause A of an *only if*-sentence 'only if A, B' is not true, then B is not true in w. For let w be such a world. Then $\{w\}$ is a member of AS($[[A]]$) which is distinct from $[[A]]$. So it is not the case that $\{w\} \subseteq B$.

But that means that B isn't true in w.[61,62]

So much for (3.227). As noted, the analysis of the *only if* part of our definition in (3.227) is more complex, since the conditional that is central to the analysis of both the *if* part and the *only if* part has a free dref. That is, in the notation we have been using the central conditional DRS condition has the form shown in (3.233).

(3.233) $$\begin{array}{|c|}\hline \\ \hline \text{featherless'}(x) \\ \text{biped'}(x) \\ \hline \end{array} \Rightarrow_{int} \begin{array}{|c|}\hline \\ \hline \text{human'}(x) \\ \hline \end{array}$$

As a matter of fact, however, this difference with what we have seen in discussing (3.229) isn't all that difficult to deal with. So long as the dref x isn't

[61]There is also a further conventional aspect of *only if*. In our account of the truth conditions of the sentences in (3.228) we assumed not only that the sentences were false for any Alternatives (that is: any elements of the Alternative Set distinct from the denotation of the associated focus), but also that the sentence is true for that denotation itself. (For instance, (3.228.b) was analyzed as true iff Bill is *the only* individual that Mary introduced to Sue.) But *only if*-sentences are usually understood as stating no more than that the main clause is false for any of the Alternatives. This seems to be a real difference between *if* as part of *only if* and the *only* that we find elsewhere (including in the sentences in (3.228)). (I suspect that this difference is the result of a special conventionalization within the mathematical sciences, where there is a clear need for a simple form for a conditional whose logical force is that of 'If B, then A', but where one wants to treat 'A' as the topic, e.g. because it expresses the proposition in which one is currently interested and of which one is exploring under which conditions it is true. If 'B' expresses a necessary condition for the truth of 'A', what would be a good form to express this. One option is 'If not B, not A'. But a need may nevertheless have been felt for a simple expression in which 'A' occurs unnegated and that looks like and at the same time contrasts with 'If B, A' in a manner analogous to the way in which 'necessary condition for A' resembles and contrasts with 'sufficient condition for A'. Given such a need, 'Only if B, A' would have been a natural candidate: All that is required for it to play that part is to ignore the '*if* part' of an *only if*-sentence, even though a corresponding meaning constituent is usually taken as part of what *only* expresses in other sentence constructions.) Of course, when *only if* is part of *if and only if*, this special property of *only if* sentences is neutralized by the explicit presence of the 'If A,B' conjunct of the *if and only it* construction.)

[62]Note well that the conditions in (3.232) do not determine AS($[[A]]$) completely. They only impose positive constraints, to the effect that certain propositions must belong to it, but no negative constraints, to the effect that certain propositions must not belong to it. it is obvious however that there must be negative constraints as well. If the Alternatives Set that is assumed for the *only if* part of an *if and only if*-sentence is too big, then the analysis of the *if and only if*-sentence may actually become self-contradictory because for some work w the *if* part part requires that the main clause be true in$w$, while *only if* part requires that the main clause be not true in w.

assigned a value, neither the antecedent nor the consequent DRS in (3.233) determines a proposition (i.e. set of possible worlds in which the DRS is true). But they both do determine a proposition once an individual has been assigned as value to x. Suppose that some (arbitrarily chosen) individual d has been assigned to x and consider the propositions $[[K_{ant}]]_d$ and $[[K_{con}]]_d$ that are expressed by antecedent and consequent DRS K_{ant} and K_{con} of (3.233) given this assignment. Then in view of what we have committed ourselves to in the discussion of (3.227), Condition (3.233) is true (in any world w') on this assignment of d to x iff $[[K_{ant}]]_d \subseteq [[K_{con}]]_d$. What does it mean for the corresponding *only if* claim to hold? Well that depends once again on what the right Alternatives Set is in this case. The crucial part of this is determined by the conditions laid down in (3.232): for each world w from W that does not belong to $[[K_{ant}]]_d$ the singleton set $\{w\}$ belongs to the Alternatives Set. With these assumptions it follows just as before that the result of applying *only* to (3.233) with the value of x fixed to be d is that $[[K_{con}]]_d$ will be true only in worlds in which $[[K_{ant}]]_d$ is true.

Summing up this discussion of how the contribution of *only* to *only if* can be seen as a special case of what *only* does in general:

(i) We can understand what *only if* has in common with other uses of *only* by seeing the *if*-clause of an *if*-clause-main clause as a condition under which the main clause can be true and which in this capacity can be seen as competing with other possible conditions on the truth of the main clause. Implementing this conception in such a way that it returns intuitively correct truth conditions requires making some special assumptions about the Alternatives Sets that are involved in the semantics of *only if*-sentences.

(ii) To deal with the starting point for this discussion, the definition in (3.227), the analysis referred to under (i) has to be allowed also for *if*-clause-main clause structures which contain free variables. So far we have seen only one part of how this works. The remainder will become clear the comments about the semantics of *something* and *it* in (3.227) to which we turn next.

2. *something* and *it*

The second aspect of our *if*-clause-main clause definition (3.227) that deserves closer attention are the roles played by the indefinite *something* and the pronoun *it*. When we first discussed pronouns in PART I, starting in Section 1.5, it was in order to motivate DRT as a replacement for Montague Grammar. Both in donkey discourses and in donkey sentences, it was argued there, the anaphoric relations between pronouns and indefinite antecedents show that we need a regime for compositional semantics that MG cannot provide. In this argument donkey sentences played a particularly important part: one of the crucial observations was that a pronoun in the consequent of a conditional, or in the nuclear scope of a universal quantification, can be anaphoric to an indefinite in the conditional's antecedent, or in the restrictor of the quantifier, but not the other way round.

All this seems incompatible with what we find in (3.227). Here it is the pronoun *it* that occurs in the *if*-clause – strictly speaking it is an *if and only if*-clause, but for the present point that makes no difference – and its apparent anaphoric antecedent *something* occurs in the main clause. Does this show that the initial arguments for DRT brought up in PART I were based on false premises,and thus, it might be argued, should be dumped on the scrapheap of failed proposals)? No. (If I believed that, these Notes would not have been written.) But it is true that the facts about pronominal anaphora are a good deal more complex than the discussions in PART I have so far revealed. A proper discussion of some of the additional complexities will be given in Section 4 (more specifically, in Section 4.3.3). There it will be argued, following observations that were made in Section 3.7.2 on the Present Tense, that sentences often admit generic readings, and that this is especially common for sentences in the Simple Present tense. Moreover, generic interpretations of Present Tense sentences impose no special restrictions on the time periods over which they are supposed to hold. Arguably the Present tense contributes the information that n belongs to this period, but that imposes no constraints on the length of the period; in particular this constraint is compatible with the possibility that the period includes all of time, which as a rule is how definitions are intended. In any case, this is of no direct relevance right here, since the temporal aspects of (3.227) have already been set aside.

Many generic interpretations take the form of generic quantifications. In our DRT-based framework these generic quantifications are represented in the form of Duplex Conditions, with a generic quantifier GEN filling the central diamond. One difficult question is what exactly is responsible for generic quantification, i.e. which constituents of the syntactic structures that serve as input to DRS construction trigger the introduction of generic Duplex Conditions. This is a question I will not attempt to deal with here. In fact, I will do no more here than give some hints as to how a DRS for (3.227) can be constructed. But the primary aim is to present the DRS that such a construction should yield and to give some indication of which parts of this DRS correspond to which constituents of the sentence. Important is in particular that the DRS Conditions corresponding to the *if* part and gthe *only if* part of (3.227) belong to the Condition Set of the nuclear scope DRS of the generic quantification. (3.234) is a first, partly schematic presentation of the DRS that ought to be outcome of the representation construction for (3.227).

(3.234)

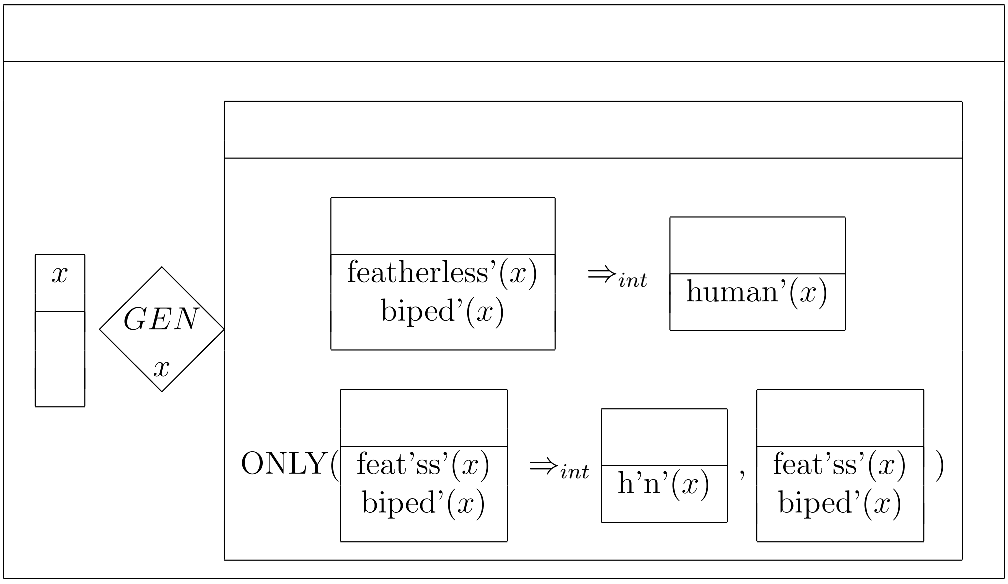

The part of (3.234) that needs explanation is the second Condition in the nuclear scope of the Duplex Condition. This is a provisional representation of the effect of applying *only* to the $\Rightarrow_{int}$ Condition above it. ONLY is represented as a 2-place predicate which takes as arguments (i) the representation of the sentence part that is modified by *only*, which in this case is the DRS for 'if x is a featherless biped, x is human', and (ii) the antecedent of this conditional, which the syntax identifies as the focus associated with *only*. To state what this amounts to semantically, we have to keep in mind that the semantic evaluation of (3.234) will, by the time it reaches the ONLY Condition, already have assigned some individual d to the dref x. Under this assignment both the antecedent and the consequent DRS of the $\Rightarrow_{int}$-Condition determine the propositions $[[K_{ant}]]_d$ and $[[K_{con}]]_d$ (see the discussion of (3.233)). It is this pair of propositions that determines the truth conditions of the ONLY-Condition for the Assignment of d to x in the way described above. The upshot of this is that for any assignment of an individual d to x the truth value (in any world w') of both Conditions in the nuclear scope of the GEN quantification are determined. The truth conditions of (3.234) as a whole further depend on the semantics of GEN, but that is a matter which cannot be discussed here.

As argued earlier, the truth conditions of the ONLY Condition entail those of the 'inverse conditional' (3.235)

(3.235) $$\begin{array}{|c|}\hline \\ \hline \text{human'}(x) \\ \hline\end{array} \Rightarrow_{int} \begin{array}{|c|}\hline \\ \hline \text{featherless'}(x) \\ \text{biped'}(x) \\ \hline\end{array}$$

The converse entailment, from (3.235) to the ONLY Condition, can also been shown to hold, given some additional but quite plausible assumptions. (These are of little interest in their own right and I will not trouble to maker them explicit.) On the assumption that the (3.235) and the ONLY Condition stand in a relation of mutual entailment,we can replace the latter by the former in (3.235) . The resulting representation for our definition (3.227) is shown in (3.236).

(3.236) $$\begin{array}{|c|}\hline \\ \hline \begin{array}{|c|}\hline x \\ \hline \\ \\ \hline\end{array} \; \left\langle \begin{matrix} GEN \\ x \end{matrix} \right\rangle \; \begin{array}{|c|}\hline \\ \hline \begin{array}{|c|}\hline \\ \hline \text{featherless'}(x) \\ \text{biped'}(x) \\ \hline\end{array} \Rightarrow_{int} \begin{array}{|c|}\hline \\ \hline \text{human'}(x) \\ \hline\end{array} \\ \\ \begin{array}{|c|}\hline \\ \hline \text{human'}(x) \\ \hline\end{array} \Rightarrow_{int} \begin{array}{|c|}\hline \\ \hline \text{featherless'}(x) \\ \text{biped'}(x) \\ \hline\end{array} \\ \hline\end{array} \\ \hline\end{array}$$

It is worth comparing this representation with the schematic representation (3.226) of the Biconditional. (3.236) isn't exactly an instance of this schema. But it comes close. First, the special 'shielding' provisions of (3.226), which provide for an extra embedding layer of the DRSs for definiens and definiendum are unnecessary for the definiens and definiendum of (3.227). For these DRSs K and K' (3.226) can be simplified to (3.237).

(3.237)

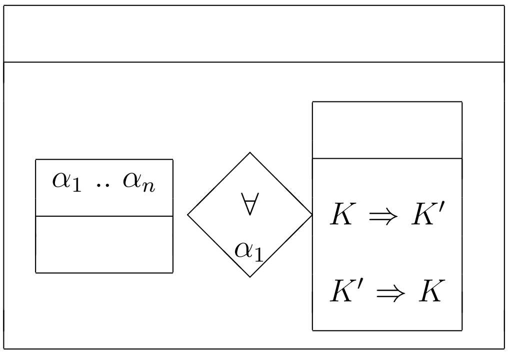

Two differences still remain between between instance and schema. The first is that (3.236) has the intensional operator $\Rightarrow_{int}$ where (3.237) has the extensional $\Rightarrow$. This is an important difference for many purposes, but arguably it is not in the case of definitions, since these are generally meant to hold non-contingently. The second difference is that (3.236) has the generic quantifier GEN whereas (3.237) has the universal quantifier $\forall$. This too is arguably not an important difference. GEN can take on different meanings depending on the context in which it is used. $\forall$ is among the special meanings it can take on in different contexts and definitions provide one of the contexts in which that is usually the case.

Quite a few details of this story still need filling in, but I leave this treatment of (3.227) in this unfinished state.

3. *featherless biped*

I conclude with a comment on the NP *featherless biped.* From a syntactic point of view this is a case of AP adjunction to an NP, of the kind discussed at length in Section 3.9.3. But there the adjuncts were relative clauses, whereas in the present case the adjunct is an Adjective Phrase. For the AP-NP combination *featherless biped* the earlier NP adjunction story seems adequate, since 'x is a featherless biped' is intuitively equivalent to 'x is a biped and x is featherless' (or as 'x is a biped that is featherless'). In view of this the reader might wonder why the semantics of the constituent *featherless biped* should be mentioned at all in this catalogue of novel problems connected with (3.227). The reason why I mention it in spite of the fact that predicate conjunction is an unproblematic analysis of the AP-NP combination in this case, is that in general this analysis of AP-NP combinations is wrong more often than it is right. Some of the classical examples for which such an analysis is not right are *small elephant*, *red hair*, *skillful cobbler*, *former senator*, *fake gun.* In all these examples the interpretation of the adjective is influenced by the noun that follows it, though in different ways, depending on what kind of adjective the adjunct is. (Compare for instance from this point of view *small* with *skillful* or *former*.) For this reason it has been proposed that prenominal adjectives should be analyzed as 'modifiers' of nouns, in the sense that the meaning of the adjective is a function that maps the meanings of the nouns with which it can combine to 'complex noun meanings'. For instance, the adjective *small* will be treated as denoting a function that maps the property denoted by *elephant* to that of the 'complex noun' *small elephant*, the property denoted by the noun *knife* to that of the 'complex noun' *small knife* and so on. Only for some adjectives the functions they denote will correspond to a property A such that the result of applying the function to any noun N will be equivalent to the conjunction of N and A. *featherless* is arguably one of those.

This is as much as I want to say here about the English wording in (3.227) of the definition according to which humans are featherless bipeds. The best that can be said about our discussion of this sentence, which has focussed exclusively on all that is still missing from the syntax-semantics interface we have thus far developed towards an adequate treatment of it, is that like our discussion of non-sentential conjunctions and disjunctions in the last section, this has been a sobering experience: So much was and still is missing from the syntax-semantics interface system we have been engaged in putting together.

Returning to the more specific question of the role of biconditionals in definitions: In the above we have looked at just one of a range of different ways in which the definition of humans as featherless bipeds can be expressed in English. As an (open-ended) exercise the reader may try to think of some other ways in which this definition, and also other definitions, can be stated in English and reflect on what is still missing from a detailed treatment of the syntax and semantics of those other formulations. (Assuming that something would still be missing; but the overwhelming likelihood is there will be.)

<u>Exercise</u>: Find as many other English formulations of the definition according to which humans are featherless bipeds besides(3.227) and reflect on what needs to be added to the syntax-semantics interface as it has been developed so far in PART II of these Notes to make an adequate treatment of those formulations possible.

3.10 More on Temporal Adverbs

This section gives a first and somewhat eclectic overview of some of the different kinds of temporal locating adverbials that are found in English. So far we have encountered only one such adverbial, the PP *at 18.00*. In one way that example is as good an illustration as any of the role that temporal locating adverbials play: that of providing times which include described events or at which described states obtain. But temporal locating adverbials differ from each other in one or both of two ways: (a) the way in which the location times they contribute are determined, and (b) the way they relate the eventuality they serve to locate to the location time they select. Because of the different options that exist for both (a) and (b) an exhaustive classification and description of all the different types of locating adverbs in a language like English is an elaborate undertaking and there can be no question of undertaking it here in earnest. Yet, even if we can do no more than skim the surface of this terrain, that will be useful, for one thing because it will give us a sense of how different types of temporal adverbials function and, just as important, because having ways to deal with a reasonably diverse range of locating adverbs will broaden the repertoire from which we can choose naturally sounding sample sentences when we want to illustrate various other phenomena.

English temporal locating adverbials can be classified in various ways and there may be no optimal classification. The classification with which we proceed here is based on a primary division into two main classes, the class of *calendar-based* adverbs and the remainder. The second of these two classes, however, is not homogeneous and will have to be subdivided further.

Some examples of adverbials involving calendar-based terms are given in (3.238). Such adverbials come in two varieties, *absolute* (or *complete*) calendar-based adverbs and *relative* (or *incomplete* or *context-dependent*) calendar-based adverbs. (The terms that I will be using most often in what follows are 'complete' and 'incomplete'.) The phrases in (3.238.a) are examples of absolute calendar-based adverbs. They determine a particular portion of actual time without the need for support of information from the context. There is only one calendar day fitting the description 'fifth of October 1973', only one month fitting the description 'June of 1605', only one summer fitting the description 'summer of 1066', only one moment fitting the description '10.05 p.m. January 1st, 1905', only one year fitting the description '527 BC' (or '1609'; I assume that when the indication 'BC' is absent, then there is a tacit agreement that the time referred to is AD.)

(3.238)a. on the fifth of October 1973; on Christmas day, 1966; in June of 1605; in the summer of 1066; at 10.05 p.m. January 1st, 1905, in 527 b.c.

b. on the fifth of October; on the fifth; on Christmas day; in June; in the summer; at 10.05 p.m, on Wednesday.

The phrases in (3.238.b) differ from those in (3.238.a) in that contextual information is needed to determine the times they refer to. To interpret *on the fifth of October* you need to know which year is being targeted, and likewise for *on Christmas day*; to interpret *on the fifth* you need to know the year and the month; to interpret *in June* you need to know the year; to interpret *at 10.05 p.m.* you need to know the year, the month and the day of that month; and to interpret *on Wednesday* you would need to know something like the year and the month, but even that won't be enough. (What you really need is some period of about a week – a period within which there is only one Wednesday – and then the indicated time will be the unique Wednesday within that period.) A prominent strategy for interpreting a temporal adverb like *on Wednesday* is to zero in on a weekday which is not a Wednesday, and then to interpret *on Wednesday* as referring to the first Wednesday after that day or the last Wednesday before it. But other options for interpreting this phrase exist as well, especially when the sentence in which it occurs is part of a larger discourse.

Note that the phrases in (3.238) all have the form of Prepositional Phrases – they all begin with one of the prepositions *on*, *in* or *at*. The times involved in the interpretations of the adverbials mentioned in our informal explanation of the difference between absolute and relative calendar-based adverbials just now were the denotations of the DPs governed by these prepositions – in this regard they all are like the adverb *at 18.00* that was part of the sentences (3.20) which we discussed in detail in Section 3.5.1. The prepositions *on* and *in* have a very similar function to that of *at*: all three say that the eventuality the adverbial is used to locate coincides with – that is: is included in, includes or overlaps with – the time that is denoted by the DP that the preposition governs. Details follow below.

on, *in* and *at* are not the only prepositions to be found in temporal locating adverbials. There are also temporal locating adverbs with other prepositions, and some of those are calendar-based as well (in the sense that they involve prepositional complements that select their denotations by reference to our calendar). A few such calendar-based adverbials are given in (3.239), in which the prepositions express other relations between the time of the locating adverb and its locandum (the eventuality that is the referential argument of the sister node).

(3.239) before 10 o'clock, after the first of January, between 14.00 and 14.30

Some temporal locating adverbs have the overt form of a DP (rather than a PP with a covert preposition). Examples are *today*, *yesterday*, *last week* and so on, as we find them in sentences like (3.240).

(3.240) He came today/yesterday/last week.

We will treat such adverbs also as PPs, which contain a tacit preposition the semantics of which is like that of *at*, *on*, *in* in that it comes down to some suitable form of temporal coincidence.

The formation of the complete and incomplete calendar terms that are the DPs of locating adverbs like those in (3.238) and (3.239) follows some of the general rules of DP formation in English, but also shows some idiosyncrasies that are specific to this particular subsystem of the grammar of English nominal expressions. We don't go into the details of this subsystem here. But it should be clear that an articulation of the compositional syntax of such terms is needed as the basis for a correct compositional account of their semantics.

3.10.1 Temporal Measure Phrases

As the expression implies, the semantics of calendar-based terms depends on our calendar system – the system which we use to divide time up into years, months, days, hours, minutes, seconds (and when desirable, as in astronomy or microphysics, into much bigger or much smaller units, like that of a light year or of a nanosecond). Each of these nits partitioned time – its full extent, from the distant past to the distant future – into intervals each of which has the length of the given unit. In this way each unit imposes a *metric* on time, a way of assessing the sizes of its different parts, by counting how many adjacent intervals of unit length are included in each given part. (Of course this can give no more than an approximation of the size of the part, and the larger the unit, the coarser the approximation.) But this isn't a serious problem, given that different temporal units are consistent in the following (intuitively obvious) sense: Whenever two given parts of time contain n copies of unit U1 and the first contains m copies of unit U2, then the second will also contain approximately m copies of unit U2; and this approximation will improve when we move to smaller units. Another, somewhat more abstract way to express this is that there is a fixed ratio r between U1 and U2, such that if any part of time includes n copies of U2, then it will contain r × n copies of U2. For the temporal units that we use these ratios are familiar to all of us: we all now that 60 is the ratio of seconds to minutes and that of minutes to hours, that there are 24 hours to a full day, and a little more than 365 days to a year. This knowledge helps us to represent sizes of temporal intervals tin ways that are easy to remember and process, by choosing the units that we find convenient. We could represent a period of five days as one of 5 × 24 × 60 × 60 seconds, but for most purposes this would introduce information into representation that would just be in our way.

Describing the size of a temporal interval in terms of a temporal unit requires (a) a name for the chosen unit – the words *year*, *moth*, *week*, *day*, *hour*, *minute*, *second*; there are more, but these will do – and (b) a repertoire of names for numbers. Minimally this repertoire should make available standard names for the natural numbers, such as *one*, *two*,.., *ten*, .., *twenty one*,.., *three hundred and sixty five*,... *1*, *w*,.., *10*,.., *21*,..., *365*,... . These unit names and number names enable us to form an infinite of *temporal measure phrases*: *five days*, *twenty five minutes*, *17 hours*, *300 years*, and so on. But there are of course other temporal measure phrases as well, such as *four and a half hours*, *several days*, *at most three weeks* and so on. (It isn't hard to give a context free grammar for English temporal measure phrases that includes all these examples. We leave this as an exercise for anyone disposed to try their hand at this.) We will often refer to temporal measure phrases with the abbreviation 'TMP'.

One syntactic environment for measure phrases is that in which they occur as modifiers of adjectives. Thus we find *three hours later*, *half an hour too late*, *five minutes early*. Combinations of a measure phrase and the comparative adjectival forms *later* and *earlier* can be combined with comparative phrases or clauses beginning with *than* to form complex adverbials with a temporal locating function. (3.248) gives a couple of examples.

(3.241)a. John arrived five minutes later than Mary did.

b. Mary arrived five minutes earlier than Mary did.

In (3.248.a) the phrase *five minutes later than Mary* temporally locates the event of John's arrival which the sentence describes, as haven occurred five minutes after Mary's arrival; likewise for (3.248.b). Like other modified and unmodified comparatives, the phrases *five minutes later* and *five minutes earlier* can also occur on their own, as in (3.242). In such cases context is needed to interpret the phrase, i.e. to answer the question 'Five minutes later/earlier than what?' This exemplifies one of the many ways in which locating adverbs can depend on context for their interpretation.

(3.242)a. John arrived five minutes later.

b. Mary arrived five minutes earlier.

The comparative phrases *earlier than* and *later than* can be replaced without change of meaning by the words *before* and *after*. *before* and *after* are standardly classified as prepositions (when they occur as in (3.243.a)) or as sentence operators (when they occur as in (3.243.b)).

(3.243)a. before/after the party/ half an hour before/after the party.

b. before/after Mary arrived at the party/ half an hour before/after Mary arrived the party.

We note in passing that these parallels between *earlier/later than* and *before/later* suggest on the one hand that *before* and *after* are also comparatives of sorts (a claim for which there is a good deal of diachronic support); this explains why *before* and *after* can take measure phrases as modifiers, as well as that they can occur without complements, as in 'I have heard that before', or 'We will deal with that later'. On the other hand the parallels also throw a certain light on comparative phrases with *earlier* and *later* and on comparative phrases more generally. both from a syntactic and a semantic point of view *earlier than* and *later than* behave much like temporal prepositions which can combine with DPs into temporal Prepositional Phrases, as in *earlier/later than the party.*

In the constructions with TMPs we have looked at above the semantic function of the TMP was to measure the temporal distance between the eventuality described by the clause to which the TMP containing adverb belongs and some other eventuality, which is either denoted by some constituent of the adverbial or else must be reconstructed from the context. There are also other types of temporal adverbs in which TMPs can occur as constituents. These are PPs beginning with *for* or with *(with)in*, such as *for an hour* or *(with)in an hour*. In the theory of Tense and Aspect temporal *for*- and *in*-PPs have a long history, going back to ((Vendler 1967)). Vendler used these PPs as tests for determine the aspectual status of verbs and some of their projections: when a verb can be combined with a *for*-PP, that shows, he observed, that it is *non-telic* (either a state verb or an *activity* verb); if a verb can be combined with an *in*-PP, that is an indication that it is *telic* (either a so-called *accomplishment* verb or a so-called *achievement* verb). Thus the combinations in (3.244.a) and (3.244.c) are fine, while those in (3.244.b) and (3.244.d) are ungrammatical.

(3.244)a. John ran for an hour/John was angry for an hour.

b. John opened the safe for an hour/John wrote a letter for an hour.

c. John opened the safe in an hour/John wrote a letter in an hour.

d. John ran in an hour/John was angry in an hour.

A Vendler himself observed his '*for an hour*/*in an hour* test doesn't always lead to unequivocal answers. But it zeros in on one important distinction among the eventualities that verbs and their projections describe. When there is no natural intrinsic termination to an eventuality – the eventuality could have come to an end earlier or could have gone on for longer – then a *for*-PP is felicitous as a way of saying for how long the eventuality went on. When the described eventuality is an event with a culmination (when the safe is finally open, or the letter finished) then this such an addition is not felicitous. The presupposition carried by *for*-PPs that the eventuality described by the event description they modify could have gone for less or more time isn't satisfied in those cases. *in*-PPs carry an opposite presupposition: that the eventuality have an intrinsic culmination; of such eventualities they then say that it took no more than the amount of time denoted by their measure phrase for the culmination to be reached.

What matters in the context of our discussion of TMPs here is that the TMPs of both *for*- and *in*-PPs the TMP has to do with the duration of the described eventuality itself, and not with the length of the temporal distance between it and some other eventuality. A *for*-PP tells us for how long the described eventuality went on. There is an implication that this is exactly how long the eventuality lasted, but this doesn't appear to be a strict entailment. For instance, (3.245.a) does not seem a contradiction.

(3.245)a. John was miserable for two weeks. And even at the end of that period his mood didn't lift. His misery just didn't want to go away.

b. He wrote the letter in half an hour. In fact, it only took him 20 minutes.

in-PPs differ in that they provide outer bounds for the duration of the described event. Here too there is a usually strong implication that the event took just as long as indicated by the *in*-PP. But here too the implication can be overruled, as demonstrated by (3.245.b).

The function, we just said, of *for*- and *in*-PP s is to provide information about how long an eventuality lasted. They do nothing to locate the described eventuality along the time axis. This is true when the DP governed by *for* or *in* is a TPM of the sort considered so far. But there also *for*- and *in*-PPs which do both things– indicate the duration of the eventuality and locate it in time – at once. Examples are those in (3.246)

(3.246)a. For the first two weeks in January 2003 John was miserable.

b. He wrote his latest novel in the first three months of this year.

That the PPs of (3.246.a) and (3.246.c) provide particular locations (as well as telling us something about their duration) is evident.
That DPs like *the first two weeks* are not TMPs is indicated by the fact that they cannot be used to modify *earlier* and *later*. *the first two weeks earlier* is plainly ungrammatical. (*the first two weeks after John arrived in Paris* is grammatical, but its syntactic analysis is a quite different one. Here the subordinate clause *after John arrived in Paris* is, just like *in January 2003*, a modifier of the NP *week*; *the first two weeks* is a modifier of *after*, in the way that *two weeks* is a modifier of *after* in *two weeks after John arrived in Paris.*)

For now this is the end of the interlude about temporal measure phrases and we return to the topic off calendar-based terms. We conclude the present subsection with a thought about the status of calendar-based terms in the sentences that contain them.

English speakers interpret the calendar-related terms of their language accruing to *their* calendar. That calendar is a product and mark of our culture. Other cultures have developed and adopted other calendars, which assign different denotations to some of the calendar-based terms we use (in particular to names that take the form of number terms or years, like *1916* and so on) and in some cases they have adopted different terms altogether. But even though the interpretation of calendar-based terms rests on a set of conventions that are subject to cultural variation, a variation that testifies to their contingency, the contributions that calendar terms make to the sentences in which they occur have for the most part a kind of 'logical' feel to them. For instance, the sentence (3.247) comes across as a tautology

(3.247) If Uncle Jim died in May 1906 and Aunt Mathilda died in August 1905, then Aunt Mathilda died before Uncle Jim.

It is important to draw a sharp distinction between the contingency that consists in adopting one particular calendar system in preference to others, some of which are found in cultures different from our own, and the status of the semantics of the calendar terms of our language that is fixed in terms of this system, once that system has become an integral part of the network of our cultural presuppositions. A sentence like (3.247) can be identified as true simply because of its grammatical structure and the meanings of its terms, including its calendar terms. In virtue of what these calendar terms denote the times that they do denote – that their denotations are determined by a calendar system that was adopted by our society in the course of time – doesn't make it any less true that this sentence can be assessed as true just on the basis of those semantic conventions, of which our calendar and the terms in our language that we use to refer to it are also part. Given those conventions the sentence could not have been false.

3.10.2 Incomplete Calendar Terms, Incomplete Temporal Comparatives and Temporal Indexicals

Incomplete calendar terms, such as *at five o'clock*, *on the seventh*, *in March*, depend for their interpretation on the context in which they are used. To know what month a speaker is referring to who uses the adverb *in March* the context has to provide us with a calendar year (so that we know she is referring to March of 2015 and not to March of 2014 or March of 2016 or March of any other year). Likewise to interpret *on the seventh* the context has to provide with some particular month (of some particular year), and to interpret *at five o'clock* the context has to provide us with a particular day (of some particular month, of some particular year). The general principle governing these 'completions' an be summarized as follows: The calendar-related temporal unit terms form a kind of hierarchy: *hour* - *day* - *month* - *year* - (*BC*/*AD*). Call this the 'Calendar Term Hierarchy'. A *coherent* calendar term is one that covers some sub-interval of this hierarchy. *at 5 o'clock, January 2nd* is well-formed in this sense, but *at 5 o'clock, January* is not, since it is missing the specification of a day within the intended January. Let τ be a well-formed incomplete calendar term and let γ be the upper element of the sub-interval of the Calendar Term Hierarchy it covers. Then interpretation of τ requires identification of some particular element of the extension of the term from the Calendar Term Hierarchy that immediately follows γ. For example, *in March* covers the one point subinterval $\{month\}$.

So its interpretation requires the identification of some particular element of the extension of the next term from the Calendar Term Hierarchy, that is of the term *year*. In other words, what needs to be identified is some particular year. The denotation of the term *(in) Martch* is then the unique month of that name in the year identified. Likewise, interpretation of *on the seventh* requires identification of some particular calendar day, whereas *at 5 o'clock, January 2nd* (which covers the subinterval {*hour*,*day*,*month*}) requires again the identification of some particular year.

Temporal adverbials involving the words *earlier*, *later*, *before* or *after* can also be incomplete. For instance, like complete expressions of these types that we find in (3.248), and which are repeated below as (**??**.a,b), we also have the incomplete adverbials in (**??**.c,d).[63]

(3.248)a. John arrived five minutes later than Mary did.

b. Mary arrived five minutes earlier than Mary did.

c. John arrived five minutes later.

d. Mary arrived five minutes earlier.

five minutes later and *five minutes earlier* are also expressions that depend on context for their interpretation. 'Later/earlier than what?' are the questions that the context must help us answer to interpret these expressions. These adverbials too it is natural to think of as incomplete. (In this regard they are like other incomplete comparative expressions, as in 'John is taller.' or 'John is two inches taller.' which provoke the questions 'Taller than who?' or 'Two inches taller than who?').
Incomplete calendar terms and incomplete temporal comparatives have in common that context is needed to correctly interpret their tokens. But that is where the similarities end. We have just seen that the kind of contextual information needed to interpret an incomplete calendar term is of a very specific kind. For adverbs like *five minutes later* the relevant interpretation principles are quite different much. They are much like those that govern the recovery of the missing complements of incomplete comparative terms generally – those principles that are also involved in the interpretation of incomplete comparatives like *taller* or *two inches taller* in 'John is (two inches) taller.'. The entity that needs to be recovered to interpret an incomplete

[63]Expressions like *five minutes after* sound somewhat marked, but native speakers tell me that they are acceptable. *five minutes after that* and perhaps *five minutes afterwards* sound more natural. *five minutes before* seems more or less interchangeable with *five minutes before*. From now on we only consider the *earlier* and *later* variants.

temporal comparative can be anything that is either a time or an entity that determines a location in time, i.e. an event.

There is laos a third type of incomplete temporal adverb. Recall the examples of (3.246), which are repeated here as (3.249.a,b).

(3.249)a. For the first two weeks in January 2003 John was miserable.

b. He wrote his latest novel in the first three months of this year.

c. For the first two weeks John was miserable.

d. In the first three months he wrote another novel.

The adverbials *for the first two weeks* and *in the first three months* in (3.246.c) and (3.246.d) are incomplete counterparts to *for the first two weeks in January 2003* in (3.246.a) and *in the first three months of this year* in (3.246.b). The incompleteness in these cases is an instance of so-called 'incomplete definite descriptions', as these are often referred to. The DPs *the first two weeks* and *the first three months* are incomplete descriptions in the sense that their descriptive contents – the NPs *first two weeks* and *first three months* do not identify a unique referent; their extensions, to put this in other terms, are not singleton sets. There are lots of two week periods that can qualify as 'the first two weeks'. To get to a unique referent we need more information, which narrows the extension of *first two weeks* down to a singleton. But as in the case of incomplete comparatives the missing information can take all sorts of forms, so long as it can be understood as as a predicate of weeks. For instance, the missing information could be 'in January 2003' (as in (3.246.a)), but it could also be something like 'of John's illness'.

Besides the three types of context-dependent temporal adverbs mentioned above there is yet another category of temporal adverbs that depend for their denotations on context. These are the so-called *temporal indexicals.* Indexicality is a p;phenomenon that we find both within the temporal domain and outside it. The clearest and least controversial examples of English indexical terms are non-temporal words. viz. the first and second person singular pronouns *I* and *you.* *I* always refers to the speaker or author of the utterance in which it occurs and singular *you* always to the addressee, in those cases where the utterance has a single addressee. There are a few exceptions to this, such as direct quotation, but on the whole these are easy to identify and set aside.[64] In the temporal domain the word that has often

[64]Here is a typical example of the exceptions.

been cited as a kind of paradigm is the word *now*, which has been described as always referring to the time at which the utterance containing it is made (again with the executions mentioned in connection with *I* and *you*). But in actual fact *now* is not a particularly good example. It can be used, and often is used, to refer to times other than the utterance time (usually times in the past of it) in contexts that do not fall under the exceptions for *I* and *you* and in which *I* and *you* therefore unequivocally refer to speaker and addressee. Two examples of this are shown in (3.251).

(3.251)a. Alan was sitting by the fire, trying to relax. The past week had been very stressful. But **now** all that was behind him and he could look forward with some pleasure to the days ahead of him.

b. He told me that he had got back quite depressed from his trip to outer Mongolia, with all its disappointments and dashed hopes, but that he was **now** ready to make a fresh start.

c. ?? I went to the butcher this morning and *now* bought a rack of lamb.

In (3.251.a) *now* refers to the time when Alan was sitting by his fire, and in (3.251.b) it refers to the time at which the subject of the sentence was telling the speaker about the aftermath of his trip to Mongolia. In both examples what justifies the use of *now* to refer to a time in the past of the utterance time of (3.251.a) or (3.251.b), respectively is that the content of the clause containing *now* is presented from the protagonist (Alan, or the referent of *he*) at the time off the eventuality that the clause describes. Compare these two with the decidedly odd appearance of *now* in (3.251.c). *now* is so odd here because there is nothing that indicates the presence of, or shift to, the kind of shifted perspective that is easy to accept in (3.251.a,b), without which a past shifted interpretation of *now* isn't possible.

Arguably better examples of temporal indexicals are the adverbs *today*, *yesterday* and *tomorrow*. These too are not perfect, as they too have non-indexical uses. For instance, *today* can be used to refer to some other day

---

(3.250)Fred said: "I have been treated so badly." So I told him: "Why are you always complaining?"

The second occurrence of *I* in (3.250) refers to the speaker of (3.250), as the rule for *I* tells us. But the first occurrence, within the direct quotations, does not refer to the speaker but to the Fred she is referring to. Likewise, *you* doesn't refer to the addressee of (3.250) but also to Fred. Not all exceptions to the mentioned interpretation rules for *I* and *you* are as easily identifiable as in tho example. But by and large the class of exceptions is well-defined and the definition transparent and straightforwardly applicable.

other than the one on which the utterance takes place, and again without this involving direct quotation. But for *today*, *yesterday* and *tomorrow* such uses appear to be more restricted than they are for *now*. [65]

now, *today*, *yesterday* and *tomorrow* are not the only temporal indexicals. Yet other examples are DPs beginning with the words *next* and *last* and in which these are followed by a calendar unit term, as in *next week*, *last month* and so on. *next week* is normally interpreted as referring to the week immediately following that which contains the utterance time, *last month* as referring to the month immediately preceding the month in which the utterance takes place. Likewise *last Wednesday* refers to the last Wednesday preceding the utterance time and *next April* to the first month of April in the future of the time when the phrase is uttered.[66]

Of the indexical adverbs *now*, *today*, *yesterday*, *tomorrow* and the bare noun phrases consisting of *next* or *last* and a calendar unit term *now* is the only one that is indexical in the strict sense that (when used as an indexical) it refers to the utterance time itself. The others refer to other times, when they are used as indexicals, but those time are related to the utterance time in some systematic, obvious way. For instance, *today*, we already noted, refers to the day containing the utterance time, *tomorrow* the day following that day, *last month* the month immediately preceding the calendar month containing the utterance time, and likewise for the others. For each of these expressions there is a simple formula which relates the denotations of their indexically used tokens to the corresponding utterance times. When a temporal indexical expression is interpreted as used indexically, then we say that the utterance time is the *origin of computation* of the interpretation of the expression: Indexical interpretations of temporal indexicals are those which use the utterance time as origin of computation for their reference.

We noted in passing that temporal indexicals also have non-indexical uses (and not only in direct quotation contexts) but that the conditions under which this is possible vary between them, with *today*, *yesterday*, *tomorrow* subject to stricter restrictions than *now*. Noun phrases beginning with *next*

[65] For a discussion in relation to corresponding temporal indexicals of French see ((Kamp & Rohrer 1983*b*)), [Schlenker ??].

[66] Note that the absence of a determiner in these phrases is crucial to their indexical behavior. For example, *next week* is quite different from *the next week*. 'I'll call you next week' is clearly a promise to call in the week after the one in which on says this. 'I'll call you the next week' is odd, and if it means anything, it is in a context in which *the next week* refers to some were other than the immediately following one.

or *last* also can be used non-indexically. They too can do this only under certain conditions. But it seems that speakers differ on the question what these conditions are; some speakers appear to be more liberal than others on this point. One of the challenges that temporal indexicals present is to identify the conditions under which the different types can be used non-indexically, and to explain what may be responsible for these differences.
As we have presented the interpretation principles of indexical uses of temporal indexicals, these expressions seem to be quite different from the three types of incomplete temporal adverbials described in the first part of this section. Indexical interpretations involve the utterance time in an essential way. Nothing that we said about the incomplete adverbials suggests that the utterance time plays a similar role in their semantics. This impression is essentially correct. But to see that and in which sense this is true this some more discussion is needed. Consider once more the incomplete calendar terms. Let's take *in June* for an example. Suppose this expression is used by me now, in December 2015. There are various ways in which the indeterminacy of the phrase could be resolved. The referent could be June of 2015 or it could be June of 2016. It could also be, in the right sort of context, the month of June of any other calendar year. (It is enough for that if the year in question has been mentioned in the preceding sentence, as in 'They moved into provisional quarters in the spring of 1939. In June they were forced to leave again and move on.') But while such alternative interpretations, in which the referent is neither the last month of June before the utterance time nor the first after it, are possible in the right contexts, without such a context they are hardly possible. The utterance time related interpretations (last such time before or first such time after) are clear defaults. In the absence of contextual information to the contrary any other interpretation is practically impossible to get.

There is a striking difference as far as this is concerned with incomplete temporal comparatives. It is very hard to interpret *an hour earlier* as referring to the time one hour before the utterance time, and even more so to interpret *an hour later* as referring to the time one hour after the time of utterance. Instead of being default interpretations, these are hard or impossible even when the context indicates that an utterance time related interpretation is intended, as in 'I don't have time now. Why didn't you ask me an hour earlier?' or 'I don't have time now. But I will deal with it an hour later.' The natural thing to say in the first example would have been *an hour ago*, and in the second *in an hour* or *an hour from now*. Such examples show that incomplete temporal indexicals are *anti-indexical*: they resist interpretations that use the utterance time as origin of computation even in contexts where everything else points in the direction of such an interpretation.

This is equally true of temporal adverbials like *for the first two hours*. 'I won't have time for the first two hours.' doesn't mean the same thing as 'I won't have time for the next two hours.': *the first two hours* can now be understood as 'the two hour period that starts now'. *for the last two hours* is different in this respect. 'I have been wasting my time for the last two hours.' is naturally understood as talking about the two hour interval that ended at the utterance time. *last*, as it occurs in phrases like *for the last two hours*, is *not* anti-indexical.

The picture that emerges from this discussion may seem a rather checkered one. There is a good deal more to be said about the possibilities and requirements of utterance time related interpretations of temporal adverbs. That would add more diversity to the picture. But further bewilderment is not what we need at this point. The purpose of this section is to give an overview of the various ways in which the interception of temporal locating adverbs can depend on context. For this purpose the different cases we have discussed suffice.

The central contrast that has emerged so far is that between temporal adverbials for which an utterance time related interpretation is the default, and those which resist such an interpretation. But there is also, I insinuated earlier, an important difference that subdivides the first group: Temporal indexicals, I suggested, differ from incomplete calendar terms, even if for both, as has subsequently transpired, utterance time related interpretations are a kind of default. The remainder of this section is devoted to the question what the difference between temporal indexicals and other expressions for which utterance time related interpretation is the default precisely is. What follows in the newt couple of paragraphs is perhaps not so easy to digest on a first pass. But it is of central importance for most of what we will say about temporal reference, and in particular about the interpretation of tense, in Section 4.

The missing part of an incomplete calendar term can be any entity of the right calendar related granularity so long as the context makes this entity available. The reason why the interpretations of such terms are so often utterance time related is that utterances are among their own context-determining factors. In particular the utterance time is always available as a potential temporal anchor (or origin of computation, in the terminology introduced above). So, unless there is some other context to compete with one established by the utterance, it is the utterance time that interpretation will exploit.

The interpretations of temporal indexicals resemble those of incomplete calendar terms only superficially. Contexts in which a temporal indexical gets an interpretation that is not utterance time related have to special. It is not just that they need to make some other time available as origins of computation. The context must induce a *change of perspective*. (Compare the comments to (3.251).) Perspectival shifts often involve a shift in time, and sometimes they appear to be just that. The temporal shift is always from the utterance time to some other time (usually, though not invariably, a time in the past of the utterance time). We refer to the time to which the perspective is shifted as the *Temporal Perspective Time*, or *Temporal Perspective Point* (TPpt). We also use this term in cases where no perspectival shift takes place. In those cases the TPpt is the utterance time itself. Thus, in any context of temporal interpretation there is a TPpt. This point is the utterance time if and only if no perspectival shift has taken place. the temporal indexical *now* always refers to the TPpt. When there is no perspectival shift, *now* refers to the utterance time; when there has been a shift, then the referent of *now* will be some other time.

By the standards of formal semantics of natural language the notion of temporal perspectival shift has a long history. Its importance was first recognized by Reichenbach in 1947 ((Reichenbach 1947)). Reichenbach used the term *reference Time* where we are using *Temporal Perspective Point* here. (The reasons for the change in terminology will be explained in Section 4.) His point of departure was the observation that the Past Perfect is often used to locate a described eventuality as situated in the past of a time that is itself in the past of the utterance time. An example is the past perfect of one of first examples in Section 3, (3.1), 'John proved the theorem in twenty lines. Mary had proved it in ten lines.'. The past perfect *had proved* in the second sentence locates Mary's proof in the past of the past time at which John is said to have proved the theorem. Reichenbach went on to analyze all tense forms of English as involving a Reference Time. For some tenses the Reference Time coincides with the Speech Time (our utterance time), while for other tenses Reference Time and Speech Time differ. (The past perfect is the paradigm example of this second possibility.) We will follow Reichenbach in the role he assigns to the Reference Time in spirit in that we will assume that temporal interpretation always involves an assumption about Perspectival Shift: either there isn't any shift – this is Reichenbach's first case, in which Reference Time and Speech Time coincide – or there is a Perspectival Shift, to some particular TPpt different from the utterance time; this second possibility corresponds to the cases where Reichenbach takes Reference Time and Speech Time to be distinct.

Note well: this cannot be the whole story about perspectival shift. I said earlier that different indexical expressions come with different requirements on contexts that allow them to get non-indexical interpretations. For instance, stricter constraints are operative in relation to *today*, *yesterday*, *tomorrow* than in relation to *now*. And a similar question should also be raised in relation to *now* and a terns form like the past perfect. But those subtleties will play no further part in these Notes.

Here is a summary of what this subsection has been about. First, we noted three types of temporal adverbials whose interpretation must rely on context - incomplete calendar terms, incomplete temporal comparatives, and incomplete temporal descriptions. We noted that the principles governing how contextual information can be exploited in the interpretation of these three types differ. We then introduced a new type of temporal adverb, the temporal indexicals, which also depend on context for their interpretation, but which do this according to principles in which the utterance time is central in a way that it isn't for the incomplete expressions already discussed.

After this we observed that the non-indexical incomplete temporal adverbials are split between the 'anti-indexicals' for which utterance time related interpretations are impossible or marginal, and the others, which are not subject to such an anti-indexical constraint. These latter ones not only permit utterance time related interpretations, but for them such interpretations apparently have a kind of default status.

Having utterance time related interpretations as a default now seems to be a common feature between these expressions and the temporal indexicals. The final part of the section was a first attempt at accounting what sets the temporal indexicals apart, as expression that can only get interpretations that are not utterance time related when they are used in contexts that involve Perspectival Shift. Perspectival Shift will be central to much of what is still to come.

3.10.3 More about *before* and *after*

In the discussion of the anaphoricity of locating adverbs of the forms *(10 minutes) earlier/later/before/after* I more than once said that the interpretations of such adverbs requires finding in the given context that can be 'either a time or an eventuality'. Well, what is it? Is it always a time, always an eventuality, or sometimes one and sometimes another? There may be no simple answer to this question, for the ultimate effect will typically be the same; even when the antecedent is an eventuality, the temporal relation expressed by *(10 minutes) earlier/later/before/after* will to the time at which that eventuality is located. However, the following consideration would seem to be relevant. We already observed that the words *before* and *after* can play both the part of a preposition and that of a conjunction; in the first case the complement is a DP, in the second it is a finite or gerundial clause. Let us first consider some examples in which *before* and *after* function as prepositions.

(3.252)

a. before the 1st of July 2018

b. after 10 o'clock

c. before the party/the eruption

d. 15 years after World War II

e. before Mary's illness

(3.252.a,b) show that the DP complement of *before* or *after* can be a term that denotes a time, (3.252.c,d,e) that it can be term denoting an eventuality. That overtly present complements of *before* and *after* can be both times and events doesn't prove that the anaphoric antecedents of expressions in which the complement of *before* or *after* is missing can also be either times or eventualities, but it certainly strongly suggests this.

The reason why the choice between reconstructing the missing argument to the temporal predicates *before* and *after*will normally make no difference to the resulting truth conditions is that when the reconstruction leads to an eventuality ev, then the temporal relation will hold between the location time t that is the referential argument of the *before* or *after*-PP and ev iff it holds between t and the duration t' of ev (i.e. $t' = \mathrm{dur}(ev)$). A difference in truth conditions could arise only if the choice in reconstruction of the missing argument was between t' and th time t'' that served as location time for ev in the semantic representation of the clause in which ev was introduced. For instance, when ev is an event dref. e, then this semantic representation will contain the condition '$e \subseteq t''$'. this condition is compatible with further information , for instance that t extends beyond e in the direction of the past. If that were so (i.e. t did extend beyond e in the direction of the past), and if the temporal adverbial we are dealing with were a *before*-phrase, then the 'before'-condition '$T < t''$', which would result when the missing argument of the *before*-phrase is reconstructed as t'', would be stronger than the condition '$t < t'$' (or, equivalently, '$t <$ dur(ev))' that results when the argument is reconstructed as ev.

Whether such cases can arise depends on the further questions of linguistic analysis as well as on details of how the construction algorithm is spelled out and implemented. not all these details have been sorted out; but in any case to address them properly will be possible only when the mechanisms are in place that guide context-based reconstruction ad that will have to wait until Section 4.

In the discussion above we have spoken of *before* and *after* as selecting for times or eventualities. But is 'eventuality' quite the right term here? Two of the examples in (3.252) in which the complement DP denotes an eventuality, (3.252.c) and (3.252.d), denote events and only the DP *Mary's illness* of the last example, (3.252.e), might be thought to denote a state. However, there are reasons for doubting that even this DP denotes a state rather than an event. Note the the eventuality denoted by *Mary's illness* is understood as lasting the entire period of time during which Mary was ill; by 'illness', one might say, is meant the entire 'bout of illness' that Mary went through, from falling via being ill to full recovery. Intuitively such an eventuality seems to qualify as an event, even if the adjectival predicate from which the noun *illness* is derived – the adjective *ill* – is naturally understood as describing a state, as it does for instance when it is part of the copular construction*be ill*. Other cases in which the complement DP of *before* or *after* may seem to refer to a state also invite such an alternative interpretation according to which it refers to an event. I will therefore assume that the prepositions *before* and *after* select for either times or events; in those cases where the DP complement looks like it might denote a state it is 'coerced' into an interpretation according to which it denotes a corresponding event; the duration of this event is some maximal period of time during which the (apparently) state describing predicate expressed by the descriptive content of the DP is satisfied. (This kind of coercion is sometimes referred to as the so-called *maximization coercion of state descriptions into event descriptions*. Maximization coercion also plays a part elsewhere in tense and aspect semantics.)

We saw that *before* and *after* not only accept DPs as complements. They can also function as 'conjunctions', taking finite or gerundial clauses as complements. Here we will limit attention to finite complement clauses. Examples of subordinate clauses are those in (3.253).

(3.253)

a. before Fred got the job

b. after Mary joined the army

c. before Mary was ill

d. after Mary was ill

The complement clauses in (3.253.a,b) are unequivocally event clauses. Their semantic contributions to sentences which contain them as temporal subordinate clauses are clear. For instance, in (3.254.a) the contribution made by *before Fred got the job* is that the locating time of the state of Susan being happy precedes the event of Fred getting the job. Analogously, the contribution of *after Mary joined the army* to (3.254.b) is that the time at which Mary is said to have been happy was later than the time when she joined the army.

(3.254)

a. Susan was happy before Fred got the job.

b. After Mary joined the army, she was happy.

The examples in (3.254.c,d) illustrate a twist to the story about *after* and *before* that I have so far been telling. Above I stated that the eventualities that *before* and *after* select for are events and that this triggers maximization coercion. But that isn't quite right is shown by an observations which goes back quite a long time ((Anscombe 1964)). The observation was that there appears to be a certain asymmetry between the conjunctions *before* and *after* which shows up in sentences in which both the main clause and the *before/after*-clause are stative. Two examples of such sentences are given in (3.255).

(3.255)

a. Fred was depressed before he had mononucleosis.

b. Fred was depressed after he had mononucleosis.

The difference between these two sentences is that (3.255.b) is ambiguous in a way that (3.255.a) is not. (3.255.a) only has a reading according to which the time at which Fred was depressed preceded the beginning of his mononucleosis. But (3.255.b) appears to have two interpretations, one according to which Fred's depression followed his mononucleosis – his depression started after his mononucleosis was over – and one according to which he became depressed after his mononucleosis started. This second interpretation is logically weaker than the first; it follows from the first interpretation in that becoming depressed after some period of time t entails becoming depressed after the beginning of t. But on the other hand there are situations in which the sentence is true on its second reading but false on the first one – for instance one in which Fred was depressed during his mononucleosis. (For some the second, weaker reading of (3.255.b) is apparently not all that easy to get. But nevertheless, cases where speakers confirm that the second reading is possible for them are well attested.)

We can explain this asymmetry by assuming (i) that when the complement of *before* or *after* is an apparently stative clause there is 'state-to-event coercion' just as there is when the complement expression is a DP; and (ii) that there isn't just one kind of 'state to event coercion', but two: not only the maximization coercion but also one that we will *inchoative state to event coercion*. Inchoative state-to-event coercion is the reinterpretation of a description of a state s into the description of an event$e$ that is the beginning of s. Let us assume that the state-to-event coercions that *before* and *after* can trigger are of either one of these two types. Then the asymmetry between *before* or *after* can be seen to fall into place. For *before*-phrases the choice between the two coercions makes no difference as far as temporal location is concerned: both maximization coercion and inchoative coercion produce the effect of locating the eventuality described by the 'main' clause before the time when the stative condition expressed by the *before*-clause started to hold. But for *after*-phrases the two coercion options produce different effects. Maximization coercion leads to location of the main clause eventuality after the time when the stative condition expressed by the *after*-clause stopped holding; inchoative coercion locates the eventuality at some time after the beginning of the period during which this condition condition is satisfied.

Note that the explanation of the last few paragraphs rests on the fact that the coercion options are maximization maximization coercion and inchoative and inchoative coercion and that there is no 'result coercion': to the state that immediately follows the holding of the stative condition expressed by the subordinate clause. It seems to be a general feature of state-to-event coercion that this isn't one of the options. I have no good explanation for why this should be so.

I also noted that inchoative coercion of stative *before*- and *after*-clauses is marginal for many speakers. Moreover, its accessibility seems to vary as a function of what the subordinate clause says. For instance, when 'had mononucleosis' in (3.255.b) is replaced by 'was ill', I do not seem to be able to get the inchoative coercion reading. I have no idea why there should be such variation and even less of an idea what the relevant factors could be.

To sum up this section: *before* and *after* can occur in various syntactic configurations, (a) with or without a modifying Temporal Measure Phrase and (b) with or without a complement DP or clause. The occurrences with overt complements show that they select for either times or events and it is reasonable in the light of this that when *before* and *after* occur without an overt complement, the missing argument that has to be retrieved from contact may be either a time or an event as well.

3.10.4 Temporal Prepositions, Temporal Conjunctions and Temporal Comparatives: a bewildering landscape littered with idiosyncrasies

Nearly all that we have said about the semantics of *before* and *after* also applies to *earlier* and *later*.These words can also occur with or without TMPs and they too can occur with or without complements and these complements can either involve DPs or clauses. The only difference is that the complements of *earlier* and *later* must begin with *than* and that the complement clauses are always finite clauses, and not gerundives. Why *earlier* and *later* require, like all other English comparatives complement phrases or clauses in which the presence of *than* is obligatory is a question for which there may be no illuminating answer and in any case this is a matter that falls outside the scope of the present discussion.

That apart from the obligatory presence of *than* the behavior of *earlier* and *later* appears to be quasi-indistinguishable for that of *before* and *after* is remarkable. Perhaps there are some differences after all that have escaped my attention. Readers are invited to see if they can find some.

Another pair of temporal expressions that show considerable similarities to *before* and *after* are *since* and *until*. (*since*, as opposed to *until*, also has a non-temporal use, but this use will not be discussed here.) *since* and *until* are words with intriguing properties and each has given rise to a literature of its own. *since* has had a good deal of attention in the context of the study of the English perfect, a tense form with a number of puzzling features, some of which show up with particular clarity when they appear in combination with *since*-phrases or -clauses. *until* presents other challenges.

The pairs (*before*, *after*) and (*since*, *until*) in that (a) all four permit both DP complements and clausal complements and (b) that the two members of each pair are, at least in first approximation, temporal mirror images of each other as far as their semantics is concerned. But there are also important differences between (*before*, *after*) and (*since*, *until*). First, the interpretation of a *since*- or *until*- phrase or-clause always requires the recovery of a time from context. *since*- and *until*- phrases and -clauses denote intervals of time. One end of such an interval is given by the complement clause or DP of *since* or *until* (or has to be recovered from context when there is no complement, just as for *before* and *after*). But the other end always has to be recovered from context.[67] In fact, this kind of retrieval follows the pattern we alluded to in our discussion of 'indexical' adverbs like *now*: the retrieved time has to be the Temporal Perspective Time (TPpt), and this requirement presupposes that the constraints imposed on TPpt selection that are imposed by the *since*/*until*-prhase or -clause on the one hand and the remained of the sentence in which this protease or clause are embedded are consistent. For instance, when the sentences in (3.257) are uttered at some time in 2017, then this consistency requirement is satisfied for (3.257.b) and (3.257.c) but not for (3.257.a) and (3.257.d): The requirement imposed by the present perfect of the (main) clause that the TPpt coincide with the utterance time renders the *since*-phrase of (3.257.a) semantically incoherent, since the starting time of the interval it is supposed to denote (viz. 2027) is later than its termination time (which is the TPpt and thus, in this case, some time in 2017). (3.257.d) presents the same kind of problem: for the *until*- phrase of this

[67] An example in which the occurrence of *since* rehires context-based retrievals of both ends of the denoted intervals would be that in (3.256)

(3.256)

He hasn't been here since.

In this case the end of the denoted interval is the TPpt, which in this case must coincide with the utterance time because of the present perfect. The beginning of the *since*-interval must be retrieved from the context in some other way. Here, since I haven't provided any context, this retrieval is impossible.

sentence the TPpt is the starting time of the interval the phrase is supposed to denote. And this time, somewhere within 2017 is later than the final time of the interval, 2007. Different utterance times will of course produce different effects. When uttered in 2037 (3.257.a,b) are coherent while (3.257.c,d) are not. When the utterance time is, say, 1997, it is the latter two that are coherent and the former that are incoherent. The reader may find it also a useful exercise to see what happens to the coherence of the sentences in (3.257) at different utterances times when the present perfects in (3.257.a,b) are replaced by past perfects and the simple future tense form in (3.257.c,d) are replaced by the corresponding future of the past (i.e. 'would live').

(3.257)

a. Fred has lived here since 2027.

b. Fred has lived here since 2007.

c. Fred will live here until 2027.

d. Fred will live here until 2007.

Another difference between *since* and *until* on the one hand and *before*and *after* on the other is that the former do not allow for modification by Temporal Measure Phrases. That is of course not surprising, as the interpretation of *since*- and *until*-phrases and -clauses requires the independent determinations of the beginning and end points of the denoted periods. These two points fix the interval of which they are the two endpoints and therewith also determine its length. So TMP modification cannot serve the purpose of helping to fix the length of the interval; that presumably accounts for its ungrammaticality. Recall in this connection that the TMP modification of a *before*- or *after*-phrase/clause provides additional information about the temporal location of the main clause eventuality.[68] For the reasons just given, TMP modification of *since*- and *until*-phrase or -clause couldn't cannot do this. But although these are good reasons why TMP modification of *before*- and *after*-phrases and -clauses is disallowed, this difference with *before*and

[68]It is possible to add a TMP to a *since*- or *until*-phrase/clause as an appositive, as in (3.258

(3.258)

Fred had lived there since 1965, for more than twenty years.

But that is in keeping with the general function of appositives: to provide additional information about the referent of the phrase to which they are apposed, but where that phrase suffices by itself to identify its referent.

after is nonetheless one that deserves explicit mention.

There are a number of temporal prepositions that cannot be used as conjunctions: *at*, *on*, *in*, *during* – the list is not complete; it could for instance be extended with compound prepositions like *in the course of* or *throughout*; but here I will limit myself to the four prepositions listed. Of these four, *at*, *on*, *in* select exclusiveely for complements that denote times, whereas *during* also allows for eventuality-denoting complements. For each of the four prepositions there is a corresponding conjunction, which does with its complement clause the same thing that the preposition does with its complement DP. For *at* this is the conjunction *when*, for *on* and *on* it is either *when* or *while* and for *during* it is just *while*. These correspondences are rough ones, but I will make no attempt here to say more about how rough or good they are. I will say no more about these corresponding conjunctions except for an observation below about *when*.

First, however, a general and quite superficial observation: all four prepositions *at*, *on*, *in*, *during* and the two corresponding conjunctions *when* and *while* express some kind of simultaneity, or some kind of temporal overlap, between their denotation and the eventualities described by the clauses in which thee y are embedded. What more, if anything, can be said about what kind of simultaneity or temporal overlap is involved depends on further factors, such as what kind of eventuality – events or states – the embedding clause describes and the temporal extent of the denotation of the PP or subordinate clause. Note in this connection that *at*, *on* and *in* differ in their selection restrictions. We say *at five*, *on Wednesday* and *in March*, *in the fall*, *in 2017*, *in this decade*,*in the Middle Ages* and so on. These restrictions – *at* with clock times, *on* with days and *in* with DPs whose heads correlate with larger calendar units – suggest that —em at wants short intervals, which can be conceived as temporal points, *in* longer intervals which are naturally understood as longer intervals of time, for which it is natural that they properly include the durations of events that occur within them, while the denotations of *on*-PPs are somewhere in between. Whether these differences are an indication that *at*, *on* and *in* express different simultaneity or overlap relations is a question that I will not try to address. it is well, however, to keep in mind when engaging in such reflections, that the relevant notions of a 'long', 'short', 'punctual' or 'extended' amount of time must be understood at a conceptual level. The time denoted by the complement DP of *in* might be by every day standards quite short. But this will be fine so long as the time can be understood as extended in the context of the particular discourse. For instance, it can be perfectly appropriate to say 'In that fraction of a second several things happened in close succession ...'. That fractions of seconds are normally considered short period of time is no ground for thinking such sentence would be bad.

during differs from *at*, *on* and *in* in that it 'selects for durative complements'. That sound rather like a tautology. But what is meant is this: The complement of *during* must be understood as a period *within* which something happens or *throughout* which something is the case. In this regard *during* is closer to *in* than it is to *on* or *at*. But *during* differs from all of *at*, *on* and *in* in that it allows for DPs that denote eventualities, and not only for DPs that denote times (and in fact, eventuality denoting DPs seems to be on the whole preferred).

The conjunctions *when* and *while*, it was surmised above, can do the same things with clauses that *at*, *on* and *in* and it 'selects for durative complements'. That sound rather like a tautology. But what is meant is this: The complement of *during* can do with time or eventuality denoting DPs. *while* corresponds most directly to *during* in that it too confers durativity upon the denotation of its (clausal) complement. And since the complements of *while* are always clauses, this means that these complement clauses must be always be interpretable as descriptions of states. *while* is also like *during* in that it expresses proper temporal inclusion of the event described by the clause embedding the *while*-clause within the state described by the complement clause of *while*, in case the embedding clause is an event description.

when is a rather different cattle of fish. It selects for event describing rather than state describing complement clauses. Furthermore, like all the other prepositions and conjunctions discussed in this section *when* expresses some kind of simultaneity. But in the case of *when*, this 'simultaneity' has been argued to be something more than, or different from, a purely temporal relation. Suppose that both the complement to *when* and the embedding clause describe events – let these be e and e', respectively. Then the relation between e and e' is often understood to be one of (close) temporal succession. An example is the following sentence:

(3.259) When he told her she was a fool, she put down the receiver.

However, this isn't always the way we understand the relate between the 'main clause event' e' and the event e contributed by the *when*-clause. The relation is often one of actual temporal overlap and it has even been argued that there are cases where e' precedes e.

The general story about the semantics of *when* appears to be something like this. The topic of a sentence with a *when*-clause adjunct is typically taken to be some event complex, or 'episode', of which main clause and *when*-clause each single out an event constituent. And these two events are then understood as standing in some kind of causal relation, of which the temporal relation is a consequence ((Moens & Steedman 1988), (Webber 1988)). In fact, there are many similarities between sentences consisting of a main clause and a *when*-clause in which the two clauses stand in causal or rhetorical relations that we also find between successive sentences in narrative discourse. The semantics and pragmatics of *when*-clauses thus reaches into a domain – that of the pragmatics of discourse relations and discourse coherence, which falls outside the scope of these Notes. Some – but only some – these issues will be touched upon in Section 4. But even those do not belong here, in this section of the Notes.

Evidently the exploration of temporal locating adverbials that we have engaged in this section (and here I mean all of Section 3.11) hasn't been very systematic, and it hasn't been anything but exhaustive. It has been guided by the aim to give an impression of the range of different mechanisms that can play a role in the interpretation of the various locating adverbials that can be found in a language like English. As indicated here and there in the course of our exploration, a formal account of most these mechanisms will be possible only in a framework that handle presuppositions. That will be put in place only in Section 4 and we will return to some of those mechanisms there.

But a more formal look at the semantics of some of the expressions we have discussed will also bring the apparently idiosyncratic features our discussion encountered into sharper focus. I am fairly convinced that many of these apparent idiosyncrasies are real ones. If there is any kind of explanation for them, then it will be one in diachronic developments are likely to play an ineliminible part.

3.10.5 Some sample DRS Constructions

This section presents DRS constructions for a number of sentences with various kinds of temporal locating adverbials. Up to the lower TP node the structure of the sentences will be simple and familiar and we will take the semantic representation construction up to that point for granted. The focus will be on the construction of the representations for the temporal adverbials and the ways in which those are combined with the representations of the TP nodes that we assume to be their adjunction sites.

We start by having another look at the one sentence with a temporal adverb for which we have so far shown a DRS construction. The sentence, (3.20), is repeated below as (3.260.a). The result of the representation construction for this structure up to the lower TP, originally shown in (3.39), is repeated as (3.260,b).

(3.260)a. At 18.00 Frieda closed the shop.

b.

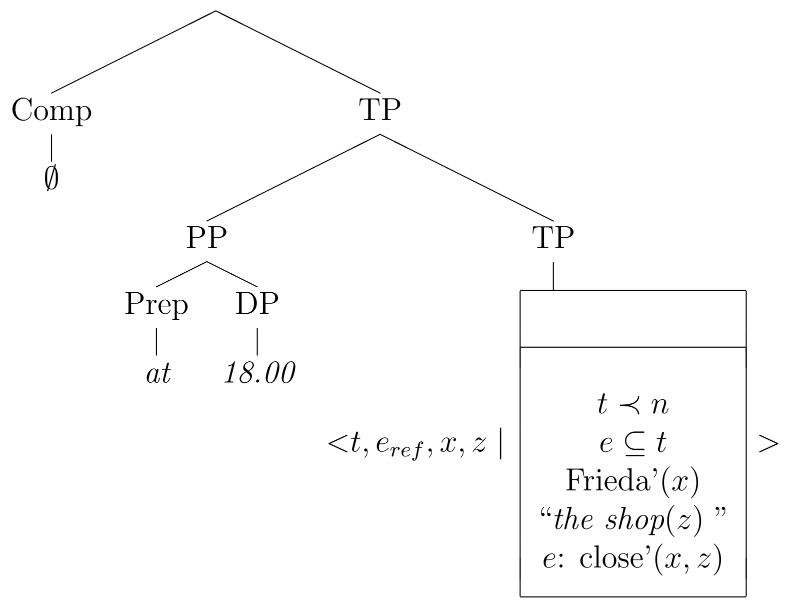

The temporal adverb of this sentence is a PP with a DP complement that is an incomplete calendar term. In the absence of a special context which offers alternatives to the utterance time as basis for the computation of the DP denotation, the utterance time will serve as a default. Our discussion of incomplete calendar terms so far wasn't fully explicit, however, on what taking a particular time as 'basis for the computation of the DP denotation' precisely comes to. Here is what was missing from that discussion: Considered on their own, incomplete calendar terms function as predicates of parts of the time axis. For instance, *18.00* is the predicate that is true at all and only those times that, according to the way in which we are keeping time and are expressing this in our language, can be described as '18.00'.

When an incomplete calendar term occurs as DP of a PP that is used as temporal adverb, then the term is to be interpreted as referring to one of the times in its extension, just as a definite description like *the table* can be used to refer to some particular table – that is, to some particular member of the extension of the predicate 'table'. In either case – *18.00* or *the table* – context is needed to zero in on the intended member of the extension. When an incomplete calendar term is interpreted as getting its referent via the utterance context, this is assumed to work as follows: the referent is either (a) the last instance of the predicate before n or (b) the first instance of the predicate after n. In simple sentences like (3.260.a) the choice between these two possibilities is decided by tense. For instance, in (3.260.a), whose tense is past, the only consistent choice is the last time before n; so that is how in this instance the phrase is interpreted according to the utterance time-based strategy. (Note well that the ambiguity between (a) and (b) cannot always be resolved in this simple way. For example, in the sentence 'He said he would be here at 18.00' the interpretation according to which *18.00* denotes the last time before n and the one according to which it denotes the first time after n are both possible. But then of course occurrences of '18.00' in such sentences may refer to many other times as well.)

This interpretation strategy can be applied to an incomplete calendar term so long as it is possible to determine the predicate P in such a way that the term denotes the last P before n or the first P after n. But that is not a real problem. Determining P from the form of the calendar term is always possible, and for the most part it is straightforward. In fact, for many incomplete calendar terms the determination of P proceeds directly from their form. Examples of temporal locating adverbials with such terms are *on Wednesday*, *on the 22nd*, *in May*. *Wednesday* is the name of a day of the week and *May* the name of a month of the year. Both are predicates that are applicable to large numbers of distinct segments of the time axis. The predicate *Wednesday* has an extension consisting of all Wednesdays – all calendar days that are Wednesdays according to our calendar – and the extension of the predicate *May* consists of all calendar months that our calendar identifies as 'months of May'. Here *Wednesday* and *May* are the respective Ps. The matter is much the same for a term like *the 22-nd* when it occurs as part of the PP *on the 22-nd*. Here *22-nd* is a predicate the extension of which consists of all days that according to our calendar are the 22-nd day of some month or other. And lastly, terms like *18.00* carry their P on their sleeve as well. *18.00* is true of all and only those *instants of time* when it is 18.00 on the clock. (There is a special problem in this last case about the 'segments' of the time axis that form the extensions of the predicate. What is an 'instant'? Are instants durationless, in the way the individual real numbers are, if the common assumption is made that time has the structure of the real number line? Or are the instants referred to by calendar terms like *18.00* genuine segments, with non-zero durations? But if the second of these two possibilities is true – if the denotations of *18.00* have non-zero durations – then what is the size of those durations? And are they all the same size? We set these questions aside here, but will turn to them briefly in the final subsection of Section 3.11.)

There is, we have just seen, a sense in which incomplete calendar terms are predicates. But from a syntactic perspective they do not function like predicates. They function like entity denoting terms that provide arguments to predicates represented by other expressions, such as the preposition like*at* in our example. In fact, incomplete calendar terms have many of the characteristics of a familiar, though semantically controversial and much debated category of DPs, that of the proper names. First, incomplete calendar terms are DPs without overt determiners, and, second, they share with typical proper names (such as *Susan*, *John Smith*, *London*, *Springfield*, *Rue Voltaire*) the property of being ambiguous: What entity is being referred to by a given occurrence of any of these expressions depends on the context. However, at least this second similarity is a rather superficial one. The ways in which these 'ambiguities' are resolved in the case of names like *Susan* or *Rue Voltaire* and the way it is resolved for incomplete calendar terms are very different. The question who a speaker or author is referring to when using *Susan* or *Rue Voltaire* is in last analysis a matter of her *referential intentions*. She must 'have some particular person or street in mind' when using the name, and at least when the intended individual or street has the name she is using for it, that then is the referent of her use of the name on this occasion. Also, a name like *Susan* or *Rue Voltaire*, as an expression of the language, doesn't have an *intension*, an independently determined principle that fixes what qualifies as a 'Susan' or a 'Rue Voltaire'. At best, names come with an associated sort. When we hear someone say 'Rue Voltaire', we ave good reason to expect that what she is referring to is a street, and when we hear her say 'Susan', we will normally expect that she is referring to a woman. But even these expectations are easily overruled. 'Susan' is as good a name for a horse, a boat, a motorcycle, a dog or any other kind of pet as it is for a woman. (Perhaps it is not so good for a male dog or a stallion, but times are changing.)

For incomplete calendar terms this isn't how their references are determined. When *18.00*, for instance, is used as a calendar term, it can refer, and only refer, to some time that is a '18.00 time', a kind of time of which there are countlessly many. And when the expression refers 'properly', it will refer to the unique time of this kind that can be found in some given period of the magnitude of a calendar day. The task of the context will be to provide such a period, in which there is exactly one *18.00* time. As far as this is concerned, incomplete calendar terms are more like another category of DPs, that of the definite descriptions. More precisely, they are like those definite descriptions that depend for their reference on the context in which they are used. Such a definite description will refer, in the context in which it is used, if and only if the context determines a 'search space' within which there is a unique satisfier of its descriptive content. For instance, *the fridge*, *the stove*, *the sink* etc. will properly refer in a context that selects some particular kitchen as search space, in which there is a unique fridge, a unique stove, a unique sink.

In the light of these various considerations incomplete calendar terms appear as a kind of hybrid between context-dependent descriptions on the one hand and proper names with multiple possible referents on the other. We might, in line with this, cast a new term for them: 'description-based names'. But more important than the name is the syntactic and semantic treatment that such expressions should be given. This is the treatment that we will adopt: (i) We assume that the syntactic structure of an incomplete calendar term is a DP with an empty determiner and in which the words that make up the incomplete calendar term are all constituents of the NP complement to this empty determiner (see the DP of the temporal adverb in (3.260,b).) Such an analysis treats incomplete calendar terms – syntactically as well as semantically – as predicates whose extensions with calendar-based extensions, as indicated above for *18.00.* (ii) The full semantics for such DPs must then articulate how the contexts in which they are used select some particular time satisfying this predicate. The utterance time based strategy is one way in which this can be done. But as hinted more than once, it is just one such strategy.

In our final treatment of the denotations of incomplete calendar terms this demand on context – that it provides a period within which there is just one time of the kind described by the term – will be handled as presuppositions (just as we will handle the unique satisfaction requirement that comes with definite descriptions more generally). But once again that treatment will have to wait till Section 4. In the DRS construction shown below we will gloss over this aspect of the semantics of *18.00.* This means that we will have to improvise, by assigning compete and incomplete calendar terms the semantic representations that they would be assigned by those contextual strategies.

(3.261)

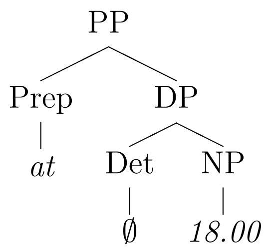

In general the NPs of such calendar term DPs are more complex of course than the one in (3.261.b). On the one hand we have DPs like *the 24-th of August 79 A.D.* in which the NP *24-th of August 79 A.D.* has a complex structure (a structure that our representations can and will display). On the other hand, name-like expressions like *18.00* also have an internal structure. The subdivision of a full 24 hour day into 24 hours,starting at midnight, or 0.00, and ending at the next midnight, or 24.00, allows us to refer concisely to certain times within that 24 hour period by certain terms for numbers. which can be systematically projected onto any 24 hour period and thereby pick out certain parts of it. The conciseness of a term like 18.00 is made possible by the general conventions of our notation for numbers, the decimal system which exploits the arabic position principle according to which the position of a digit in a finite sequence indicates the power of 10 with which it must be multiplied to determine its additive contribution to the number denoted by the sequence. Notation systems like this one have their own syntactic and semantic rules. In a fully explicit treatment these rules should be specified, as the syntax-semantics interface for this submodule of of the grammar of the language as a whole. For the submodule that is relevant here spelling out those rules isn't a particularly difficult task. But it is a task that involves quite a bit of work, and work from which little could be learned that would be directly relevant to our present concerns. So we will treat clock terms like *18.00* as unstructured NPs, as in (3.261.a).

Note by the way that this complication – the internal structure of the expressions that make up calendar term NPs – arises for clock time terms but not for for calendar terms of coarser granularity levels. For instance, *Wednesday* in *on Wednesday* and *March* in *in March* are semantic primitives, which can't be analyzed in terms of smaller linguistically relevant parts. Terms of the granularity level of years, such as 2017 or 79 AD do of course present the same kind of internal structure issues as clock time terms.

Since we are not analyzing the NP *18.00* into its structural constituents here, we cannot do better than treat it as an atomic predicate. So the result of 'lexical insertion' for *18.00* yields the NP representation in (3.262), where the dref t' is chosen as referential argument introduced by the predicate expression *18.00*.

(3.262)

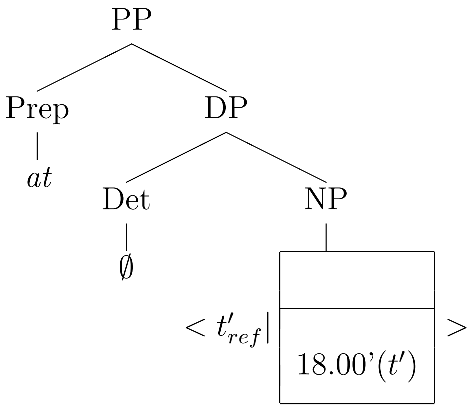

The step from the NP representation in (3.262) to the representation of the DP containing the represented NP involves the context=dependent resolution that we cannot properly deal with at this stage of the development of our theory, at which the treatment of presuppositions has not yet been included. All we can do is give a provisional account of this step.

As noted earlier, the resolution of incomplete calendar terms can take different forms, one of which is what we called the utterance time based resolution strategy. Principle (3.263) states in more precise terms what we were hinting at at that earlier time. What (3.263) says is that the denotation of the DP *18.00* is either the time of this description that is nearest to n in the past of n or the one that is nearest to n in the future of n. These two possibilities are $\overset{!}{\vee}$-alternatives in the sense that interpretations in which the principle is applied will always make a choice between them. In all cases we will consider in these Notes, the choice will be determined by tense (In the example under discussion, in which *at 18.00* is part of a simple past tense clause it is obviously the first $\overset{!}{\vee}$-disjunct of (3.263) that must be chosen.)

(3.263) (application of utterance time-based strategy in the interpretation of an incomplete calendar term DP τ)

Let $<t'_{ref}, ... \mid K>$ be the representation of the NP that is sister to the Det node of τ. Then the representation of the DP node obtained via the utterance time-based strategy is as given to the right of $\leadsto$ in the following transition schema:

$$<t'_{ref}\ldots \mid K> \;\leadsto$$

$$<t'_{ref}\ldots \mid \begin{array}{|c|}\hline \\ \hline K \cup \begin{array}{|c|}\hline \\ \hline t' \prec n \\ \neg \begin{array}{|c|}\hline t'' \\ \hline t' \prec t'' \prec n \\ P(t'')\;\; [K]t''/t' \\ \hline \end{array} \\ \hline \end{array} \\ \overset{!}{\vee} \\ K \cup \begin{array}{|c|}\hline \\ \hline n \prec t' \\ \neg \begin{array}{|c|}\hline t'' \\ \hline n \prec t'' \prec t' \\ P(t'')\;\; [K]t''/t' \\ \hline \end{array} \\ \hline \end{array} \\ \hline \end{array} >$$

Application of (3.263) to (3.262) yields (3.264).

(3.264)

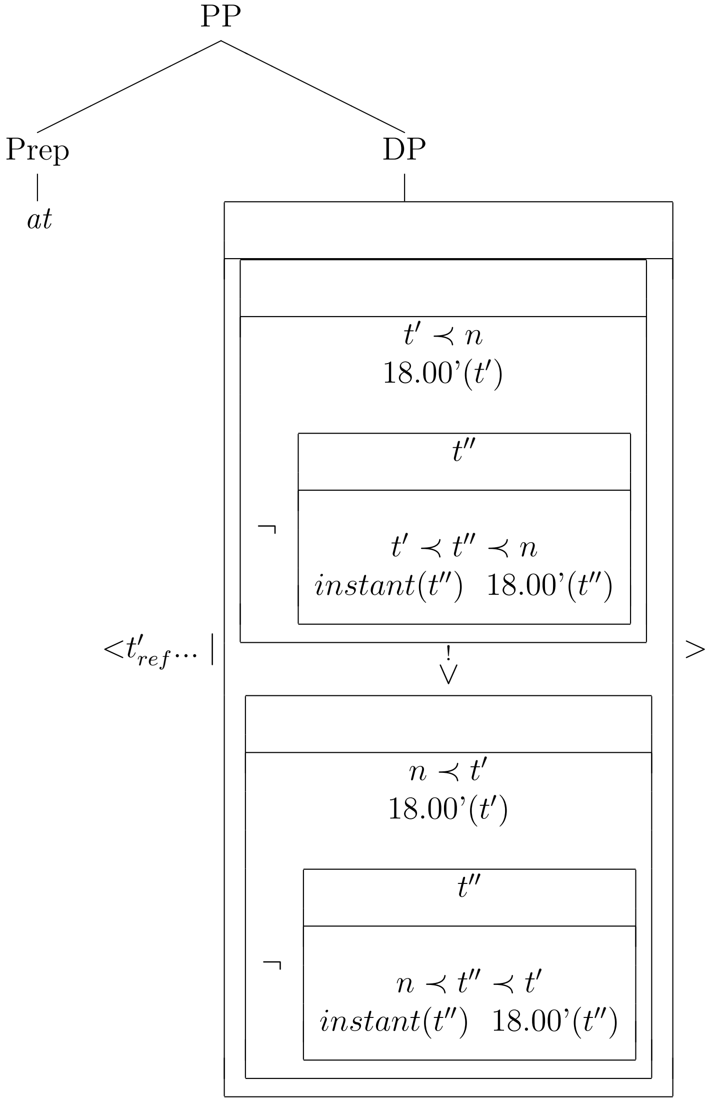

The transition from the DP representation to the PP representation involves the switch from the denotation of the DP to the location time that is made available by the PP. In the present case, where the preposition is *at*, this is a purely formal change, since the two times are the same. Nevertheless it is instructive to show the transition, since it instantiates a pattern of which we will see some further non-trivial instances in some of the sample constructions that follow. The result of the transition is shown in (3.265).

(3.265)

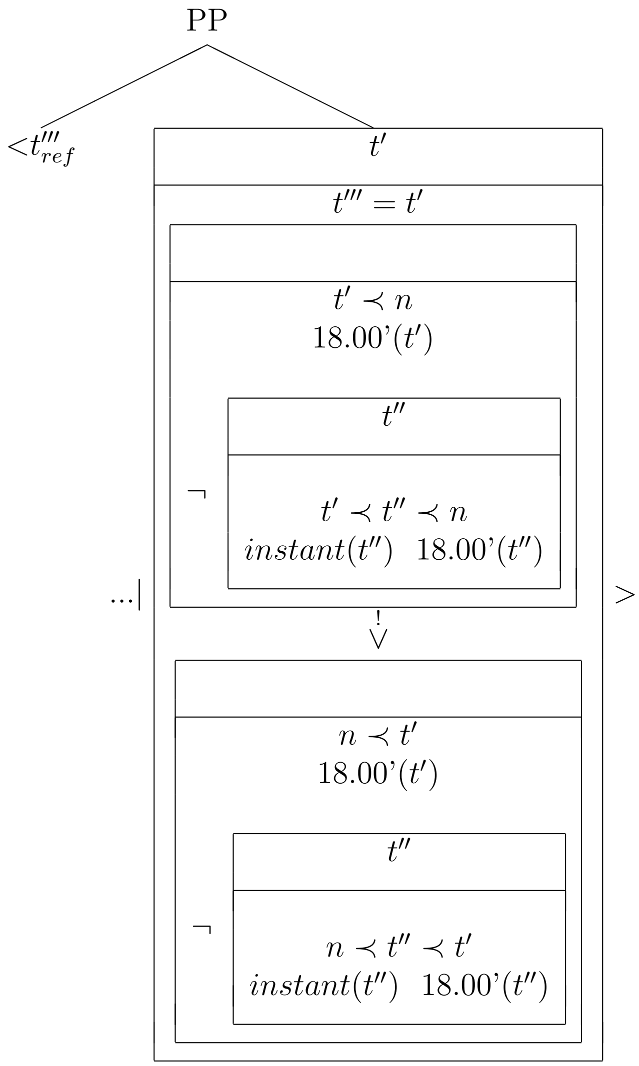

When the semantic representation in (3.265) is combined with the TP representation in (3.260.b), the referential argument t''' in (3.265) will serve as location time for the referential argument e of the TP representation, yielding the condition $e \subseteq t'''$. Since the TP contains the conditions $e \subseteq t$ and $t \prec n$, the second disjunct of the $\overset{!}{\vee}$ disjunction in (3.265) aborts because of inconsistency. So only the first disjunct, whose conditions are consistent with those of the TP representation, remains and we get as final representation for (3.260.a) the DRS in (3.266).

(3.266)

$$
\begin{array}{|c|}
\hline
t \;\; e \;\; x\, t' \;\; t''' \\
\hline
t \prec n \;\; e \subseteq t \;\; e \subseteq t''' \;\; t''' = t' \;\; t' \prec n \;\; \text{Frieda'}(x) \;\; \text{18.00'}(t') \\
\neg \begin{array}{|c|} \hline t'' \\ \hline t' \prec t'' \prec n \\ \mathit{instant}(t'') \;\; \text{18.00'}(t'') \\ \hline \end{array} \\
\text{“}\mathit{the\ shop}(z)\text{ ”} \\
e\text{: close'}(x, z) \\
\hline
\end{array}
$$

Next we look at an example with an absolute calendar term, viz.

(3.267) Mount Vesuvius erupted on the 24-th of August 79 A.D.

Before we turn to the semantics of this sentence first a remark about its syntax. One difference between this sentence and (3.260.a) is that this time the temporal adverbial does not occur in sentence-initial position. Is that a difference that should concern us? Answer: In general yes, in this particular instance no. As to the general answer: One important question about temporal adverbs is where they are syntactically adjoined. In (3.260.a) we have been assuming that its adverb is an adjunct to TP. For sentence-initial adverbs this assumption seems to be fairly uncontroversial. At the very least there is agreement that the attachment must be high up in the syntactic tree. (Exceptions, insofar as there are any at all, al seem to be connected with noticeable information structure effects, such as fronting of constituents that play a role of contrastive topic.) But for adverbs occurring in positions that are not sentence-initial the matter is less clear, and this is true in particular for adverbs that occur sentence-finally. Sentences with sentence-final temporal adverbs ar often ambiguous in ways that are best explained as ambiguities of adverb attachment. For instance, on its most plausible interpretation the adverbial *on a Saturday* in the sentence ‘He didn’t submit his abstract on a Saturday.’ is in the scope of negation and – at least given the assumptions about negation we made in Section 3.10.1 – this means that the adverb cannot be an adjunct to TP. In what follows we will ignore this complication, and assume that sentence-final as well as sentence-initial temporal adverbs are always TP adjuncts. On this assumption (3.267) gets the syntactic parse in (3.268).

(3.268)

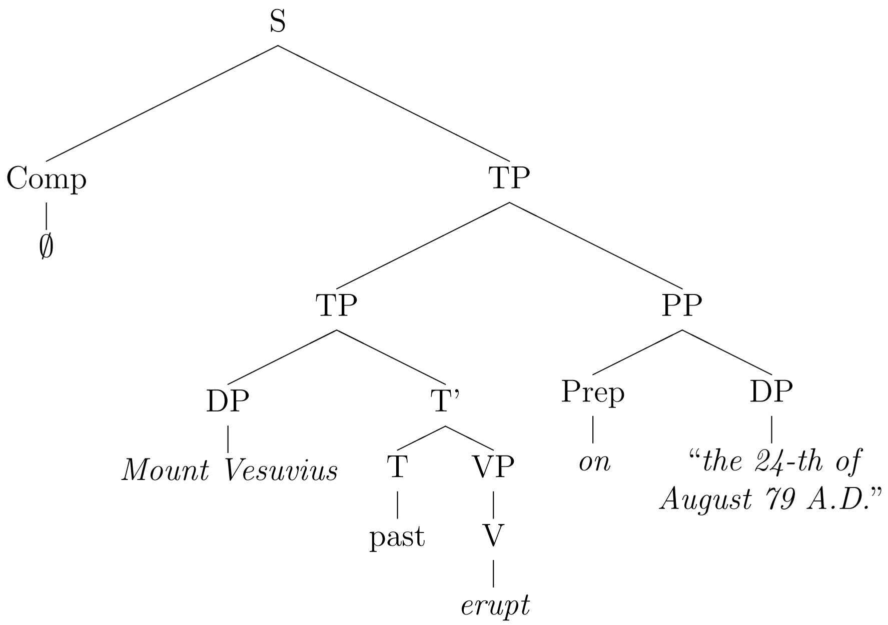

The difference between an absolute calendar term like *the 24-th of August 79 A.D* and an incomplete one like *18.00* is that the former denotes a particular time independently of context. It is able to do this because the NP constituent of the term expresses a predicate which, independently of context, has a unique satisfier: there is one and only one calendar day that fits the description '24-th of August 79 A.D.'. Behind this fact about unique satisfaction there is another aspect of terms such *24-th* and *August* to which we have not so far drawn explicit attention. This is that such terms are not only interpretable as predicates but as *relational* predicates, viz. as the predicate *24-th of* and the predicate *August of*, were the second (non-referential) argument place of the first predicate is to be filled by a calendar month and the non-referential argument place of the second predicate by a calendar year. (It is a curiosity of English syntax – for which I have no explanation; and I doubt that there can be an interesting explanation for it – that the *of* in *24-th of* is always overtly expressed, whereas the *of* in *August of* can be overtly expressed but need not be.) Because these predicates are relational they can be turned into complex non-relational predicates by filling their non-referential argument slots with suitable argument DPs. Thus the relational *August (of)* can be turned into the non-relational *August (of) 79 A.D.* through insertion of the DP *79 A.D.* and the relational *24-th of* can be turned into the non-relational *24-th of August 79 A.D* through insertion of the DP *August (of) 79 A.D.*. Since the DP *79 A.D.* denotes a particular calendar year (for reasons that we won't go into again), and the fact that the relational predicate *August* has the property that the relation it denotes is 'inversely functional' – that is, for each second argument (i.e. each calendar year) there is a unique first argument (the month of August of that year) – the non-relational predicate *August (of) 79 A.D.* has a unique satisfier, and that is then the referent of the DP *August (of) 79 A.D.*. Likewise, and for the same sort of reason, the non-relational predicate *24-th of August 79 A.D* has a unique satisfier, which therefore is the referent of the DP *the 24-th of August 79 A.D.*.

Below a few steps are shown of the representation construction for the temporal adverbial *on the 24-th of August 79 A.D.* (Once more, all reference to and representations of presuppositions have been suppressed in the DRS constructions that are presented in the current Section 3. All constructions are provisional in this respect.) We start with the LF shown in (3.269).

Note that the relation predicates *August* and *24th* are treated as combining into larger NPs with adjoined PPs. These adjunctions are to be thought of as triggering semantic operations that involve argument insertion – of the referential argument of the DP that is governed by the (overt or tacit) preposition into the non-referential argument slot of the predicate. Insofar as these combinations of relational nominal and adjoined PP are seen as syntactic realizations of argument insertion, they are not instances of adjunction in the semantic sense that we have so far been assuming sod far, that according to which the semantic representation of an adjunction involves unification of the representations of adjunct and adjunction site. To distinguish the present case of syntactic adjunction, whose semantic realization involves argument insertion rather than unification, for the kind considered hitherto, for which unification is the semantic realization, we will refer to the new kind as 'pseudo-adjunction'. But giving the creature a name isn't solving the real problem that this new form of adjunction presents: How do we know which syntactic adjunctions are pseudo-adjunctions and which 'real' adjunctions?

For our purposes in these Notes the following answer to this question will suffice. Pseudo-adjunctions are if (and only if) prepositions do not function as semantic constituents of their own – by expressing a relation that holds between the referential argument for the DP they govern and the referential argument of the adjunction site of their PP – but function rather as 1case markers' of sorts: expressions that indicate which argument position of the predicative expression to which the PP is attached is the slot for the referential argument of the DP. These case cannot be handled as cases of unification because the preposition does not introduce its own predicate with its own referential argument, and so there is no dref in such cases that could be unified with the referential argument of the adjunction site.

The LF for the temporal locating adverb *on the 24-th of August 79 A.D.* from which we will construct ist semantic representation is given in (3.269)).

(3.269)

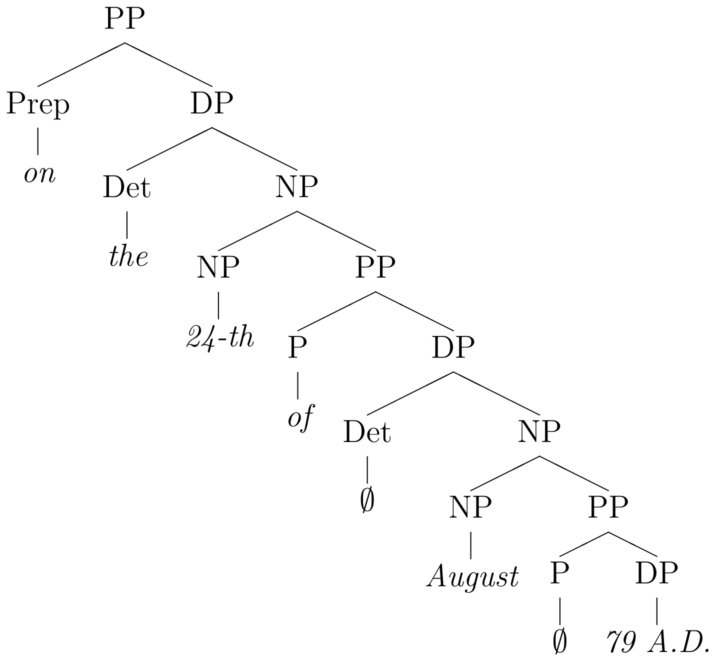

As representation for the 'proper name' *79 A.D.* we assume the structure in (3.270).

(3.270) $<t'_{ref} \mid \boxed{\text{79 A.D.'}(t')}>$

The next step is the one that constructs the representation of the PP directly containing the DP *79 A.D.* – thus, the lower of the two PPs in (3.269) – from the semantics in (3.270). This is a step in which nothing really happens: since the P node of this PP doesn't contribute a predicate of its own, the DP semantics is simply transferred to the PP node.

In the following step the semantics of the lower PP is combined with that of the NP *August* into the complex NP *August 79 A.D.*. The semantic representation of the NP *August* is obtained from that of the noun *August*, which we assume is interpreted as a relational predicate, denoted in our DRS formalism as 'August-of'. This NP representation is shown in (3.271).

(3.271)

$$<t^2_{ref} \mid \begin{array}{|c|}\hline t' \\ \hline \text{August-of'}(t^2, \underline{t'}) \\ \hline\end{array}>$$

Combining this representation of the NP *August* with that of its argument phrase *79 A.D.* leads to (3.273). Note that this step confronts us with a problem of which we have seen many instances before. Since the [NP PP]$_{PP}$ combination we are dealing with here is a case of pseudo-adjunction, and thus the corresponding step in the representation construction one of argument insertion, the familiar issue arises into which argument slot the referential argument of the DP should be inserted. Until now we relied on the assumption that the syntactic parser will make this information available. This is something we could do in cases of pseudo-adjunction too. However, for the cases of pseudo-adjunction we will consider argument insertion is unproblematic insofar as there will always be just one slot, viz. the non-referential argument position of the 2-place predicate introduced by the relational head noun of the NP, that is available for argument insertion. We can of course subsume the cases of argument insertion that are triggered by pseudo-adjunction under our general implementation of argument insertion by adding the relevant coindexations. That is, we can assume that the partially interpreted syntactic structure for has the form shown in (3.272).

(3.272)

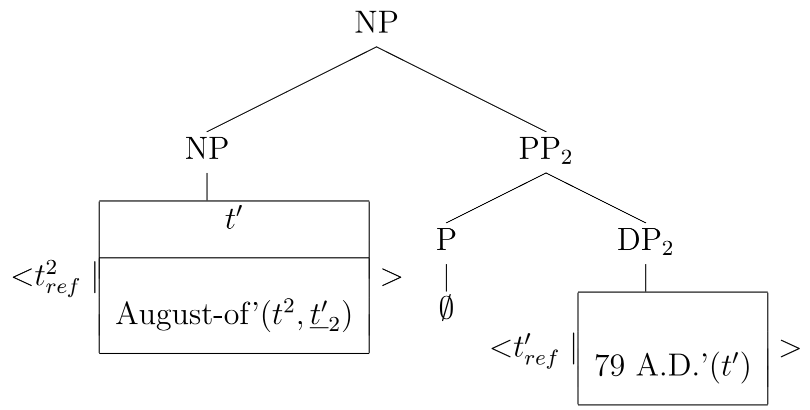

Merging the representations of the NP and the DP daughter of the upper NP node after inserting the referential argument of the 2-indexed PP into the 2-indexed argument slot of the NP representation leads to the representation in (3.273) for *August 79 A.D.*.

(3.273)

$$<t^2_{ref} \mid \begin{array}{|c|}\hline t' \\ \hline \text{August-of'}(t^2, t') \\ \text{79 A.D.'}(t') \\ \hline \end{array} >$$

The transition from the representation of *August 79 A.D.* as NP to its representation as DP doesn't show any overt differences. The only difference between DP representation and NP representation in a case like this is that the DP representation, qua representation of an expression of the category DP, can combine with other representations by way of inserting its referential argument into some argument slot. But the dref that is centrally involved in such an operation – the one that gets inserted into the coindexed argument slot of the predicate to which the DP is a syntactic argument – is the very same dref that plays the part of referential argument in the NP representation. In short, it is only the circumstance that the DP representation functions as the representation of a DP, and thus as involved in those compositional operations that are triggered by syntactic configurations that DPs (but not NPs) can occur in, that distinguishes such representations from the NP representations from which they are constructed, and from which they may be indistinguishable in their overt form. The upshot of this is that the representation of the DP *August 79 A.D.* is the very same as the NP representation in (3.273).

Once again the transition from this DP representation to the representation of the PP of which the DP is a daughter is trivial in that the DP representation is simply passed up from DP to PP. (The preposition *of* is semantically vacuous, just like the empty preposition of the PP *79 A.D.*.) Combining this PP representation with that of the NP *24-th* is another case of what we have just seen goes on when we looked at how the lower PP is combined with the NP *August*. The result is the one shown in (3.274).

$$(3.274) \quad <t^3_{ref} \mid \begin{array}{|c|} \hline t^2 \quad t' \\ \hline \\ \text{24-th'}(t^3, t^2) \\ \text{August-of'}(t'^2, t') \\ \text{79 A.D.'}(t') \\ \hline \end{array} >$$

Once again the representation for the DP *the 24-th of August 79 A.D.* is indistinguishable from (3.274) (just as the DP representation of *August 79 A.D.* is indistinguishable from its NP representation). After this the representation of the DP *August 79 A.D.* must be combined with the semantics of the preposition *on*. *on*, note well, should be treated as a sen,antically non-vacuous preposition, which contributes a temporal relation between the referential argument of the DP *August 79 A.D.* and the referential argument of the PP's adjunction site. In our example this will be the event e introduced by the verb, i.e. the famous eruption of Mt. Vesuvius. But when the representation for the PP is being constructed this information is not yet available, and in particular it is not yet known at that point whether the referential argument of the adjunction's site will turn out to be an event or a state. So we are running into the familiar situation that the representation of the PP will have to contain a $\overset{!}{\vee}$-disjunction with one disjunct for the case the eventuality will turn out to be an event and the other for the case where it turns out to be a state. We already dealt with this problem earlier in connection with the temporal preposition *at*. With *on* the problem appears to be a little different. When the referential argument is an event, then as before the relationship that *on* expresses between it and the referential argument t^4 of the Dp is that of the event being temporally included within this time. But when the referential argument is a state, the relation is somewhat different from the one we encountered in connection with *at*. I can be said to have had a tooth ache on Wednesday when it was only on Wednesday when I had my tooth ache, and not necessarily all of Wednesday. But it is also possible that my tooth ache already started on Tuesday, or that it continued on Thursday, or both.

The homogeneity principle for states, according to which for any state type S, any state s instantiating S and any time t temporally included in s there is a state s' that instantiates S and whose duration is t, enables us to capture what is common to these different possibilities in a single simple formula: there is a time t' included in t that is also included in the state represented by the referential argument. It should be clear enough how the lexical entry for *at* can be transformed into that for *on* in which this interaction of *on* with states is correctly captured. (In case you do not see this right away, write down the indicated entry for *on*, by way of an exercise.) The lexical entry for *on* is used in the representation of the (3.275), which ought to make the exercise even easier.

$$(3.275)\ < ev_{ref}, t^3 \mid \begin{array}{|c|}\hline t^2\ \ t' \\ \hline \begin{array}{|c|}\hline \\ \hline \text{Event}(ev) \\ ev \subseteq t^3 \\ \hline \end{array} \ \overset{!}{\vee}\ \begin{array}{|c|}\hline t^4 \\ \hline \text{State}(ev) \\ t^4 \subseteq t^3 \\ t^3 \subseteq ev \\ \hline \end{array} \\ \text{24-th-of'}(t^3, t^2) \\ \text{August-of'}(t^2, t') \\ \text{79 A.D.'}(t') \\ \hline \end{array} >$$

(3.276.a) shows the representation of the lower TP of (3.267). (There is nothing about the construction of this representation that calls for comments after all that we have seen.) (3.276.b) gives the final DRS for the sentence, after the lower TP representation has been combined with the representation in (3.275) for the temporal adjunct. This combination makes use of the fact that the referential argument of the lower TP node is an event, which makes it possible to eliminate the second disjunct of the $\overset{!}{\vee}$-disjunction in (3.275).

(3.276)

a.

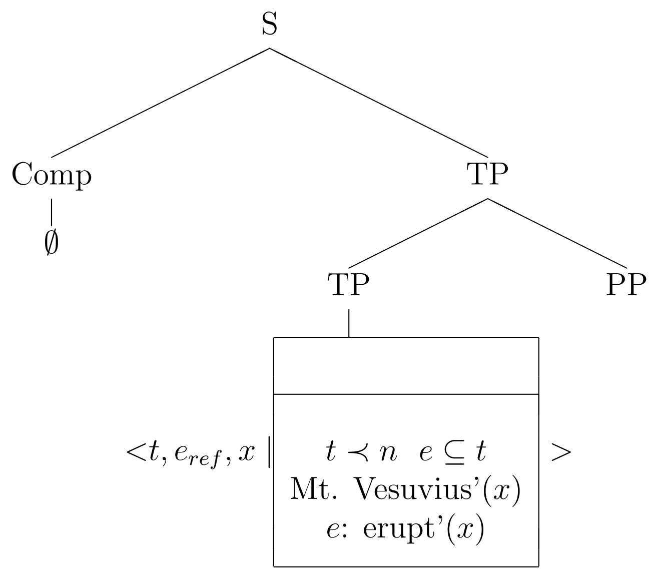

b.

$t \;\; e \;\; x \;\; t^3 \;\; t^2 \;\; t'$
24-th-of'(t^3, t^2) August-of'(t^2,t') 79 A.D.'(t') $t \prec n \quad e \subseteq t \quad e \subseteq t^3$ Mt. Vesuvius'(x) e: erupt'(x)

Next we turn to a sentence with a temporal indexical. The sentence is a minor variant of our last example.

(3.277) Mount Etna erupted last year.

The DRS construction for (3.277) is largely identical to the one we have just gone through. Some steps from this construction are shown in (3.278).

(3.278.a) gives the syntactic structure of the PP, (3.278.b) the representation of the upper NP node and (3.278.c) the representation of the PP.

(3.278)

a.

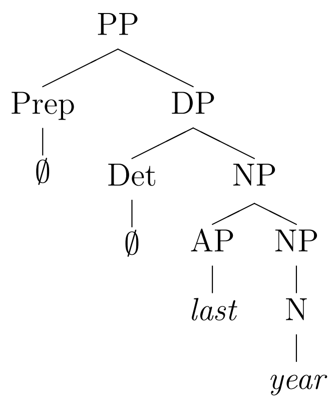

b.

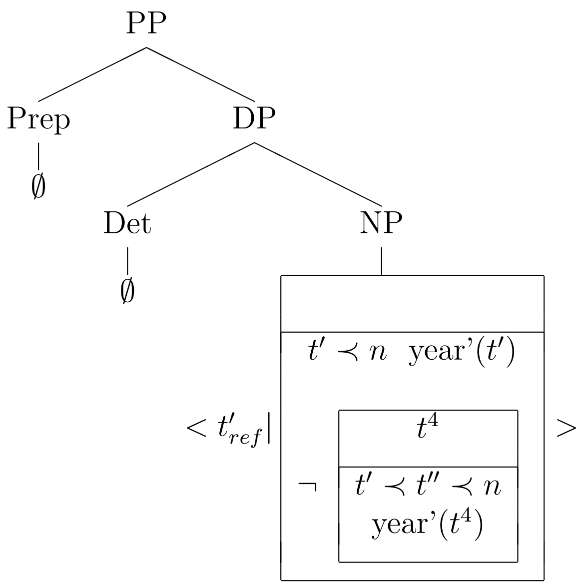

c.

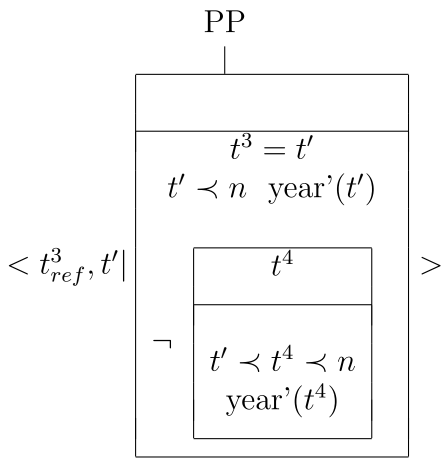

Comments:

1. The upper NP representation is obtained by combining the calendar predicate *year*(which at the same time is also the temporal granularity predicate in this case) with the indexical adjective *last*. When applied to the representation of an NP whose nominal head is a calendar related predicate such as *year*, *last* produces the effect described by its 'special purpose entry' given in (3.279), which reflects the indexicality of temporal adverbials beginning with *last*. (In a more comprehensive lexical entry for *last* the effect recorded in (3.279) would be just one item on a longer list, which would vary as a function of the input representation.)

(3.279) (‘special purpose’ lexical entry for *last* as indexical modifier of sortal calendar predicates)

last (adjective)

Sel. Restr: sortal calendar predicate

Sem.Repr:

$$<t \mid \begin{array}{|c|}\hline \\ \hline \mathrm{P}(t) \\ t \prec n \\ \hline\end{array} > \;\rightsquigarrow\; <t \mid \begin{array}{|c|}\hline \\ \hline \mathrm{P}(t) \\ \neg \begin{array}{|c|}\hline t' \\ \hline t \prec t' \prec n \\ P(t') \\ \hline\end{array} \\ \hline\end{array} >$$

The representations in (3.278.b) and (3.278.c) are the result of assuming that the TPpt of sentence (3.277) of which the PP is part coincides with the utterance time n. The lexical entry for *last* in (3.279) only covers this case. In a more comprehensive (and mor accurate) version of this entry the two occurrences of ‘n’ should be replaced by ‘TPpt’. For details see Section 4.

2. The predicate *year* shares with other calendar unit predicates the property that successive instances of it abut. This means that the relation between the years t' and t'' that in (3.278.b) and (3.278.c) is expressed by the condition that there is no year between them can expressed more simply as abutment: $t' \supset\subset t^2$. A Meaning Postulate for the calendar predicate *year* which states the equivalence of these two conditions will permit the simplification of (3.278.b) and (3.278.c) that can be obtained by switching to the abutment condition. We show the result only for the case of (3.278.c). It is displayed in (3.280).

(3.280)

$$\begin{array}{c} \text{PP} \\ | \\ < t^3_{ref}, t' | \begin{array}{|c|} \hline t^2 \\ \hline t^3 = t' \\ t' \prec n \quad \text{year'}(t') \\ \text{year'}(t^2) \quad n \subseteq t^2 \\ t' \supset\subset t^2 \\ \hline \end{array} > \end{array}$$

Combining the PP representation – either in the form (3.278.c) or in the form (3.280)

The remaining DRS constructions in this section are for sentences with *before*- and *after*-phrases and -clauses. The sentences are further variations of the Vesuvius/Etna outbreak sentences. Because the only ways in which the new variants differ from the earlier ones, we will be able to focus exclusively on their temporal adverbs. (3.281) lists the sentences we will be looking at.

(3.281)

a. Mount Vesuvius erupted after the 21-st of August 79 A.D.

b. Mount Vesuvius erupted three days after the 21-st of August 79 A.D.

c. Mount Vesuvius spewed small bits of rock (three days) before the big eruption.

d. Mount Vesuvius spewed small bits of rock (three days) before it erupted on the 24-th of August 79 A.D.

The LF for the PP of (3.281.a) is given (3.282).

(3.282)

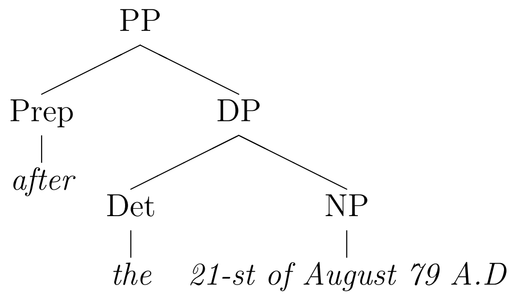

The internal structure of the NP *21st of August 79 A.D.* has been suppressed since it differs only trivially from the NP *24-th of August 79 A.D.*, with which we have been dealing in agonizing detail before, and the semantic representations differ only in their respective predicates *21st-of* and *24-th-of*. Also, as in the case of *24-th of August 79 A.D.* the representation of the NP *21st of August 79 A.D.* gets passed unchanged to the DP node.

It then gets combined with the preposition *after*. Like *at* and *on*, *after* is a genuine preposition, which expresses a relation between its referential argument and the argument that is contributed by its DP. In the case of *after* this relation is always temporal precedence, but a question may be raised about what the second argument can be. In Section 3.11.3 we hypothesized that the second arguments of the relations expressed by *after* and *before* must be either times or events, and we argued that when the complement DP appears to be describing a state, then interpretation will involve state-to-event coercion. For the case at hand, where the complement DP of *after* denotes a time, this issue doesn't arise; but it will later on, when we look at the semantic representation construction for (3.281.d).

These considerations lead to a semantic representation for the *after*-PP of (3.281.a) that has the form shown in (3.283).

$$(3.283)\quad < ev_{ref}, t^3 \mid \begin{array}{|c|} \hline t^2 \;\; t' \\ \hline t^3 \prec ev \\ \text{21st-of'}(t^3, t^2) \\ \text{August-of'}(t^2, t') \\ \text{79 A.D.'}(t') \\ \hline \end{array} >$$

The remainder of the DRS construction for (3.281.a) is as for the last two sentences and we omit it.

(3.281.b) differs from (3.281.a) only in that its temporal adverb involves a measure phrase modification of the preposition *after*. We assume the following syntactic structure for this PP.

(3.284)

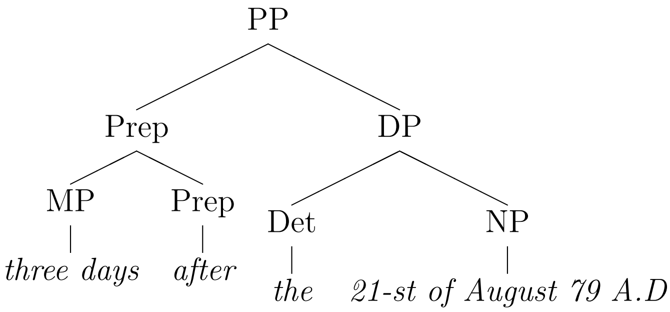

Recall that 'TMP' stands for 'Temporal Measure Phrase' (see Section 3.11.1). Dealing properly with the contribution that *three days* makes to the semantics of the PP in (3.284) would require a detailed formal development of the semantics of measure phrases in general. We won't engage in such a project here and instead adopt the following ad hoc substitute: we assume, consistently with what was said about the role of measure phrases earlier, that their semantics results from the combination of two parts, (i) a unit of measurement (uom) given by the nominal head (here: the temporal unit that has the length of a calendar day, a complete diurnal cycle of 24 hours) and (ii) a real number, denoted by the adjunct to the head (here: the number three). These two elements are combined into a predicate that is true of an entity of the relevant sort iff the quantity of that entity along the relevant dimension is correctly given by the specified number of units. In particular, in the case before us the relevant entities are temporal intervals and the predicate given by *three days* is true of an interval t iff t has the length of three calendar days. In our representations we represent such quantity-related predications in the form $\mu_U(x) = r$, where U is the given unit of measurement and x is the bearer of the predication. Thus, that t has a duration of three days is represented as '$\mu_{day}(t) = 3$'.

The next issue that has to be addressed in connection with the semantic representation of the PP in (3.284) is how the representation of the MP *three days* combines with the semantics of *after*. Earlier we observed that the combination of *three days* with *after* follows the general pattern of MP modifications of comparatives. Implicit in what we said at that point was that the comparative forms of adjectives are used to describe the distance, along the dimension indicated by the adjective, between two entities to which the adjective is being applied. A plain, unmodified comparative simply asserts that the first of these entities, the bearer of the complex predicate that is expressed by the AP whose adjective is in comparative form, exceeds the second entity (often mentioned explicitly in a *than*-phrase) by some unspecified positive amount. When the comparative form is modified by a MP, this means that the excess of the first over the second entity is of the magnitude denoted by the MP. To make sure that the combination of comparative and modifying MP makes the right contribution to the phrase of which they are part – in the present case: the PP in (3.284) – we have to make sure that the excess gets represented by a dref which can serve as argument for the predicate expressed by the comparative construction. (There may be a deeper explanation of how this excess is made available by comparative constructions, but if there is, I do not know what it is.)

Unfortunately the impression I created earlier that combinations of Temporal Measure Phrases (TMPs) and *earlier/later/before/after* are just a special case of combinations of MPs with comparatives is am oversimplification. Many TMPs have properties that can not be directly predicted from this assumption, and *three days* as it occurs in *three days after the 21st of August 79 A.D.* is a striking case of this. The problem has to do with the fact that *day* is ambiguous between a use as 'pure' temporal measure unit term and its run of the mill use in which it refers to days as opposed to nights.

Before I say more about this, let us first get another issue out of the way. In our discussion of *before* and *after* in section 3.11.3 we noted that *after*-phrases are sometimes subject to an ambiguity that has to do with whether *after* must be taken to express a precedence relation between the eventuality that the *after*-phrase is used to locate and the time or eventuality that the *after*-phrase makes available as second argument of this relation. I argued that this ambiguity arises only in comparatively special cases, viz. when the complement of *after* in the *after*-phrase describes what appears to be a state, where its occurrence as complement of *after* triggers state-to-event coercion, and where this coercion can be inchoative as well as maximizing. We found that the possibility of inceptive coercion is quite restricted (although we couldn't say exactly what the restrictions are), but in nay case this problem arises only when the complement of *after* has a default interpretation as state description. The case before us, in which the complement of *after* is the DP *the 21st of August 79 A.D.*, is not of this kind. In fact, we will not be looking at any examples in which the complement is state describing per default. So the possibility that the distance between the located eventuality and the time or event that the *after*-phrase makes available starts with the beginning of the latter need not be considered. From now in the Notes this possibility is set aside.

The problem, I just said, about the contribution that *three days* makes to the semantics of *three days after the 21st of August 79 A.D.* has to do with the different ways in which we use the noun *day*. In fact, there three possible meanings for *day* that can be made out: (i) *day* in the sense of a period lasting from dawn till night fall, (ii) *day* in the sense of calendar day (lasting from 0.00 a.m. till 12.00 p.m.) and (iii) *day* in the sense of a period lasting 24 hours, irrespective of where in the course of day or night its end points are situated. Our earlier discussion of the role of MPs in the semantics of comparatives modified by MPs suggests that it is the third meaning of *day* that should be relevant here. But more often than not that does' appear to be the relevant meaning, or at least not the only one. Examples where *daw* does seem to have its '24 hours' meaning are those in (3.285).

(3.285)

a. Fred arrived a day and a half after he called.

b. On Wednesday Fred called to say he would be coming. He arrived exactly one day after his call.

(3.285.a) is intuitively true only when the arrival time is roughly 36 hours after the time of calling. Here em day seems to play its part as temporal measuring unit term – 36 = 1.5 x 24. Likewise (3.285.b). This sentence seems to convey that the calle was roughly 24 hours after the call; it would, I think, be likely to be judged false when, say, Fred called Wednesday morning at 9.00 a. and then arrived Thursday at 10 p.m. But to get the 'period of 24 hours' meaning for *day* when it heads a TMP appears to be natural when it combines with a number expression that seems to denote a real or rational number, but not when it is understood as denoting a positive integer. In the two examples in (3.285.a), for example, the number of days referred to is one and a half. And in (3.285.b) the phrase *exactly one* contrasts one day with periods of slightly less and slightly more than a day. But when the number phrase modifying *day* is naturally understood as referring to one of the positive integers, as the *three* in the TMP *three days*, then the interpretation of the TMP is typically one based on the *counting* of days rather than the measuring of a temporal interval using *day* as unit of measurement. For instance, in the sample sentence that has led us into this discussion the temporal locating adverb *three days after the 21st of August 79 A.D.* locates the described event, the outbreak of Mt. Vesuvius, as somewhere within the third day after the 21st of August 79 A.D. – thus the 24th. This is compatible with the claim that is implied by the '24 hours' interpretation of *day* insofar as the temporal distance between the end of the 21st of August and the outbreak has to be of roughly the length of 3 x 24 hours. On the one hand a concession of the kind that is made by the 'roughly' of this statement is clearly necessary. The distance of the eruption of Mt. Vesuvius from the end of the 21st need not be exactly 72 hours in order that sentence (3.281) count as true; the distance can be anywhere between 48 and 72 hours. But on the other 'roughly' seems to concede too much, since a period of more than 72 hours (say one of 74 hours) does not seem to be compatible with what the sentence says.

All this goes to show that *three days after the 21st of August 79 A.D.* says just the same thing as *on the third day after the 21st of August 79 A.D.*. This is an indication that phrases of the form *three days after ...* – and likewise *three days before ...* – do not behave like TMP-modified comparatives in the sense of the discussion in Section 3.11.3. They illustrate the possibility of using what is from a more general perspective a grammatical construction designed for the expression of measured distances for the purpose of what is a related but nevertheless importantly different paradigm: that of counting. This phenomenon is not restricted to the term *day*; we also find it, almost as strongly, with *month* and with *year*, and to some extent with *week*. We do not find with the terms *hour*. *minute* or *second* – these are unambiguously units of temporal measurement.

For a systematic account of the semantics of temporal locating adverbs in which a comparative is modified by what looks like a TMP the phenomenon we identified in the last few paragraphs is an additional bother. Rather than treating all (real or apparent) TMPs as measuring the relevant distances of the located eventuality from the relevant 'origin of computation' we have to distinguish between true measurement interpretations and 'counting' interpretations. The latter, we have just seen, come into play when (a) the head of the TMP is a calendar related term like *day*, *month*, *year*, and perhaps *week* and (b) the indicated number is naturally interpreted as a positive integer. When not both of these conditions are satisfied, then, I propose, a measurement interpretation is appropriate. Let us make what is surely a simplifying assumption: that a counting interpretation is to be chosen when conditions (a) and (b) are both satisfied and a measuring interpretation otherwise.[69] Since for the case we are considering – the locating adverb *three days after*

[69]Such an interpretation still doesn't require that the relevant temporal distance be exactly the amount indicated by the number specification of the TMP. For instance, if Fred called at 9.00 a.m. on Wednesday and he arrived the next day at 8.00 p.m., or at 10.00 p.m., we might still consider (3.285.a) to be true. But this kind of 'imprecision' (as the official term has it) is something quite other than the difference between measuring and counting interpretations that has been the subject of this discussion.

the 21st of August 79 A.D. – both (a) and (b) are satisfied, this means that here it is the counting interpretation that is to be chosen.

The formal implications of this simplification are these: we first identify the origin of computation. In the case at hand this is the end of the period denoted by the 21st of August 79 A.D. Here we encounter another instance of the ambiguity associated wit the notion of a day: When does the 21st of August 79 A.D. end, at the onset of night or at 12.00 p.m.? Let us assume that it is the latter.[70] The next step in the interpretation of *three days after the 21st of August 79 A.D.* is the choice between a measuring and a counting interpretation. According to our simplifying assumption the choice has to be for the count in interpretation, since the conditions (a) and (b) of the last paragraph are fulfilled. Given this choice, the contribution made by th locating adverb *three days after the 21st of August 79 A.D.* is to the effect that the described eventuality (the eruption) is included within the n-th day after the origin of computation, where n is the number denoted by the number specification of the 'TMP'. In the case before us n = 3. So the eruption must have taken place within the third day from the end of the 21st of August.

The general principle here is as follows. Let α be a phrase in which *after* occurs with a complement and modified by an (apparent) TMP. Let t_0 be the origin of computation (determined by the complement DP or clause of *after*; in our example this is the DP *the 21st of August 79 A.D.*), let P be the lexical head of the 'TMP' (in our example this is the noun *day*) and let n be the positive integer denoted by the number specification of the 'TMP' (in our example this is 3). Then locating information contributed by α is that the described eventuality is included within the instance t of P that is separated from t_0 by a set T of instances of P that has cardinality n-1. (This specification is clear as it stands when n > 1. When n = 1, the set T will be empty; in this case the specification mans that t must abut t_0 on the right. For instance, if α were *one day after the 21st of August*, then t would have to be the day that abuts on the right the time 12.00 p.m. of the 21st of August.)[71]

[70]The assumption that the 21st of August 79 A.D. ends at the beginning or end of dusk cayuses additional hiccups in the interpretation of *three days after the 21st of August 79 A.D.*, which I will not elaborate on here, but which will be easier to see as soon as we have said more about what goes into the interpretation of this phrase.

[71]Someone might want to push the analysis even further by asking what it mean for a set to have cardinality k, where k is any natural number. For a proper answer to this question we would have to go back to a remark made earlier about the 'grammar of number': a subsystem of the grammar of English which deals with the syntactic forms of English

There is one last remark that we need to make about the interpretation of our sentence (3.281.b). It concerns the semantics of the verb *erupt*. The adverbial *three days after the 21st of August 79 A.D.* provides a temporal location for the eventuality that *erupt* introduces into the semantic representation of (3.281.b). But what is this eventuality? Surely it is an event. But which event? Is it the entire eruption, from the moment Mt Vesuvius started to spew actual lava (as well as vast quantities of ashes) until the point where this ended. Or is it the starting decent of this protracted process? Independently from the sentential contact in which the verb appears. This question is impossible to answer, apart from the acknowledgement that both options exist. But in the presence of a temporal locating adverb that specifies a certain time t within which the event described by the verb must be included some kind of charity principle seems to guide interpretation. When for instance t is of the order of magnitude of a day (which it is in our example) and the described event is of the kind all of which typically takes less than a day, then the natural interpretation will be that the entire event was included within t. (So, for instance, if volcanic eruptions typically lasted for a fraction of a day, then the natural interpretation of (3.281.b) would be that the entire eruption took place on the 24th of August 79 A.D.) But on the assumption that eruptions can and often do extend over several days, such an interpretation for (3.281.b) would be impossible or unlikely. In the light of this we can and prefer to interpreted the simple past tense *erupted* in (3.281.b) as referring to the onset of the eruption, rather than to all of it. But note well that the exact wording is important here. The sentence 'The last big eruption of Mt. Vesuvius took place on the 24th of August 79 A.D.' cannot be interpreted in this way. This sentence can only mean that all of the eruption occurred on the 24th, not just the beginning of it. The deverbal noun *eruption* can only denote all of the eruption, not just its onset.

---

number terms. On the basis of a syntactic definition for a chosen set of number terms we could define a denotation semantics of these terms. But since we have decided not to formulate a syntax for a system of number terms we are lacking the foundation for such a definition of the semantics of number terms also.

The problem of number terms semantics also arises, by the way, in connection with ordinal number expressions – *first*. *second*, *third*, *fourth* and so on – as distinct from the cardinal expressions that have been in the focus of our attention. Here too what is strictly speaking needed is (i) a syntactic definition of a comprehensive set of ordinal number expressions and (ii) a definition of the semantics of such expressions that builds on this syntactic definition. One of the ultimate benefits of such syntactic and semantic characterizations of cardinal number expressions on the one hand and ordinal expressions on the other would have to be an account of why *three days after* and *on the third day after* make the same semantic contribution.

In our formal representation of (3.281.b) we will ignore this complication and represent the event located by *three days after the 21st of August 79 A.D.* simply as the event introduced by the occurrence of *erupt*.

At long last we are ready to present the semantic representations of the adverbial *three days after the 21st of August 79 A.D.* of (3.281.b), as a temporal location description of its referential argument *ev* ((3.286.a) below) , and of sentence (3.281.b) as a whole in (3.286.b). After all that has been said these representations should be very nearly able to speak for themselves. We forebear a stepwise construction of the semantic representation of *three days after the 21st of August 79 A.D.*, but leave this as an exercise for the reader.

(3.286)

a.

$$< ev_{ref} | \begin{array}{|c|} \hline t^4 \ \ T \ \ t_0 \ \ t^3 \ \ t^2 \ \ t' \\ \hline \begin{array}{|c|} \hline \\ \hline \text{Event}(ev) \\ ev \subseteq t^4 \\ \hline \end{array} \ \overset{!}{\vee} \ \begin{array}{|c|} \hline t^6 \\ \hline \text{State}(ev) \\ t^6 \subseteq ev \\ t^6 \subseteq t^4 \\ \hline \end{array} \\ T = \{t^5\text{: day'}(t^5) \ \& \ t_0 \prec t^5 \prec t^4\} \quad |\, T \,| = 3 \text{ - } 1 \\ \text{end-of}(t_0, t^3) \quad \text{12.00 p.m.}(t_0) \\ \text{21st-of'}(t^3, t^2) \\ \text{August-of'}(t^2, t') \\ \text{79 A.D.'}(t') \\ \hline \end{array} >$$

b.

$$\begin{array}{|c|} \hline t \ \ e \ \ x \ \ t^4 \ \ T \ \ t_0 \ \ t^3 \ \ t^2 \ \ t' \\ \hline T = \{t^5\text{: day'}(t^5) \ \& \ t_0 \prec t^5 \prec t^4\} \quad |\, T \,| = 3 \text{ - } 1 \\ \text{end-of}(t_0, t^3) \quad \text{12.00 p.m.}(t_0) \\ \text{21st-of'}(t^3, t^2) \\ \text{August-of'}(t^2, t') \\ \text{79 A.D.'}(t') \\ \\ t \prec n \quad e \subseteq t \quad e \subseteq t^4 \\ \text{Mt. Vesuvius'}(x) \\ e\text{: erupt'}(x) \\ \hline \end{array}$$

All of the notation used in (3.286) has been used before, though some of it not since the end of Section 2. This is in particular of the set notation '$X = \{x: \Phi(x)\}$' which expresses that X is the set consisting of all x that satisfy the condition Φ and the notion '$|X|$' for the cardinality (= number of elements) of the set X.

For the last two sentences in (3.281) we only consider the versions without the measure phrase *three days*. The representation constructions for the versions with the measure phrase are left to the reader.

The main difference between the PP of (3.281.a) and that of (3.281.c) is that the DP of the former denotes a time whereas that of the latter denotes an event. However, this difference doesn't amount to very much – we already noted this – so long as the eventuality of (3.281.c) can be understood as making its semantic contribution by way of its duration; moreover, given the convention we have been using all along, according to which temporal relations involving eventualities can be presented as if they were relations between times, with the eventualities tacitly contributing their durations, we can write the contribution expressed by *before the big eruption* as the conjunction of the condition '$t^3 \prec e$ ', where e represents the big eruption and t^3 is the location time, and a $\overset{!}{\vee}$ condition for the relation between t^3 and the described eventuality ev. For the PP of (3.281.c) (without *three days*) the upshot of this is a semantic representation of which a simplified representation is given in (3.287) – simplified because we are lacking the proper means of handling definite descriptions like *the eruption*. (In a way that is of course also true for the definite descriptions *the 24th of August 79 A.D.* and *the 21st of August 79 A.D.*, where we have been gating the unique reference for granted (as guaranteed by the properties of our calendar). The temporal relation contributed by *before* is of course the converse of that contributed by *after*.)

Once more I present just the semantic representation of the temporal adverb *before the big eruption*, in (3.287.a), and of the entire sentence, in (3.287.b).

(3.287)

a.

$$
\begin{array}{|c|}
\hline
t^3 \quad e \\
\hline
\\
\begin{array}{|c|}
\hline
\\
\hline
\text{Event}(ev) \\
ev \subseteq t^3 \\
\hline
\end{array}
\;\overset{!}{\vee}\;
\begin{array}{|c|}
\hline
t^4 \\
\hline
\text{State}(ev) \\
t^4 \subseteq ev \\
t^4 \subseteq t^3 \\
\hline
\end{array}
\\
\\
t^3 \prec e \\
\text{“the big eruption”}(e) \\
\hline
\end{array}
$$

b.

$$
\begin{array}{|c|}
\hline
t \quad e' \quad x \quad t^3 \quad e \\
\hline
\\
t^3 \prec e \quad e' \subseteq t^3 \\
\text{“the big eruption”}(e) \\
\text{Mt. Vesuvius'}(x) \\
e' : \text{“spew small bits of rock”}(\text{x}) \\
\hline
\end{array}
$$

In the last example from the four listed in (3.281) the temporal adjunct is a subordinate clause. As noted earlier, in our set-up there is not much difference between the contributions that are made by temporal subordinate clauses and by temporal PPs with eventuality denoting DPs: both types of adverbial make an eventuality dref available for unification with the referential argument of the adjunction site and in both cases this referential argument represents the described event.

In order for this parallelism to work out the way it should, it must be assumed that temporal subordinate clauses are adjuncts to TP, just as we have been assuming for all temporal adverbs that have been considered so far. Thus the LF for (3.281.d) is assumed to be as in (3.289). But before we give this structure we first have to settle on a proposal for the representation of temporal subordinate clauses. The one we will adopt is shown in (3.288) for the subordinate clause of (3.281.d). The category label ‘TSC’ (for ‘ Temporal Subordinate Clause’) denotes a sub-category of the category S. This subcategory is suited for adjunction to the TP of the clause within which the TSC is embedded. The filler of the Comp position – here the conjunction *before* – determines the semantic relation between the TSC and its host clause.

(3.288)

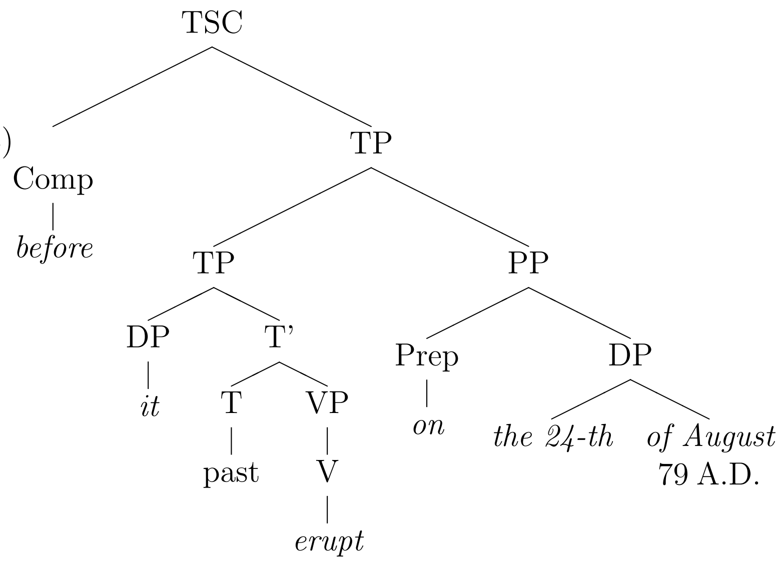

(3.289)

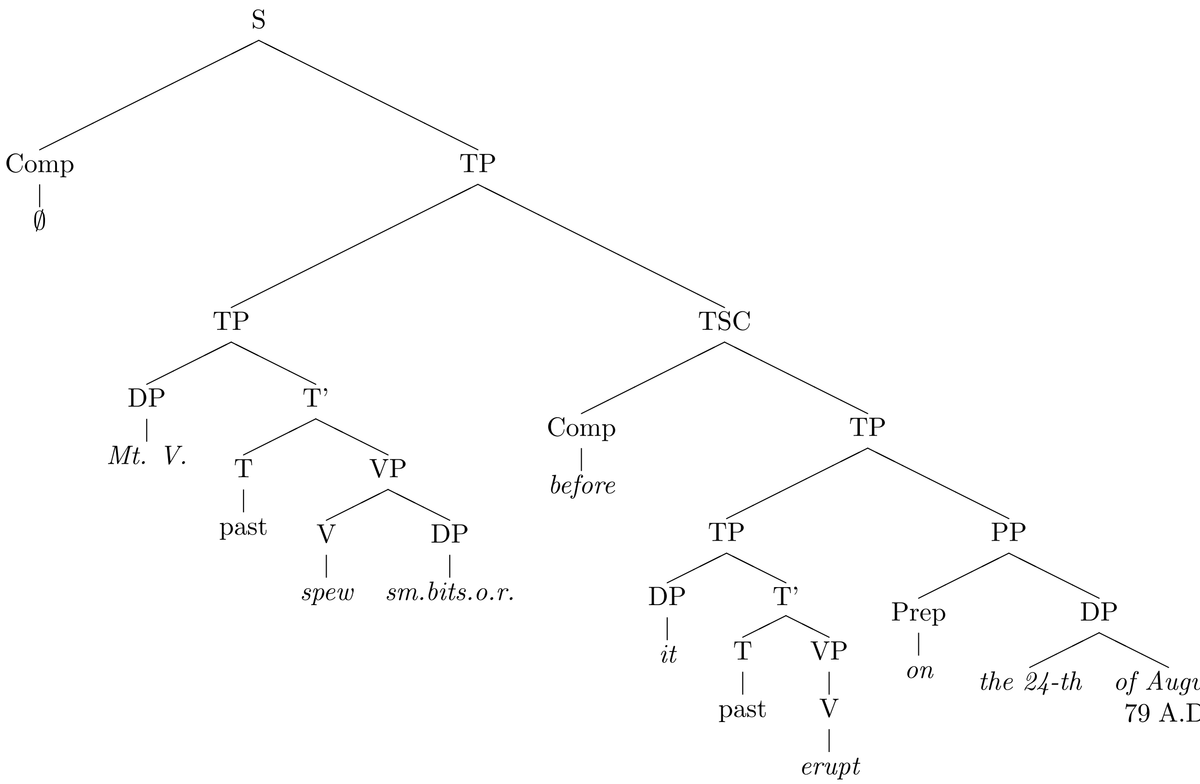

To give an idea of how the representation construction works in this case we first show, in (3.292.a), the semantic representation of the upper TP of the subordinate clause and then, in (3.292.b), the representation of the entire TSC. The construction for the upper TP node of the LF for the TSC in (3.292.a) is identical with that of sentence (3.267) and is presented here without further comment.

The transition from the representation in (3.292.a) to the semantic representation of the TSC confronts with a new problem. More accurately, it focus us upon a problem that isn't exactly new, but that we could more easily ignore before than we can at this point. It is a problem that arises generally for modified and unmodified expressions headed by *before* and *after* (and also for the corresponding expressions that are headed by *earlier* and *later*). So far the expressions of these types that we considered were phrases in which the complement of *before* or *after* was a DP. We have been treating all these phrases as adverbs – as temporal locating adverbs, to be more precise – which appear in the sentences of which they are part as TP adjuncts. As such they had to be treated semantically as predicates of eventualities, as representations whose referential arguments are eventuality drefs, which can be unified with the referential arguments of the adjunction site. But these are not the only syntactic configurations in which phrases of this form can occur. We also find them as constituents of expressions like those in (3.290).

(3.290)

a. some time after the 21st of August

b. a summer many years before the war

c. the last war before World War II

d. my hangover after last night's party

In all these examples the *before* or *after*-phrase is not an adjunct to a TP but to a nominal constituent (an NP, presumably). We ignored such cases, and arguably on fairly goods grounds, because we have been concerned throughout Section 3.11, with temporal *adverbs*. That the phrases in question also play the grammatical role of adjectives was a matter beyond out horizon.

In principle we could persist in this attitude also vis-—'a-vis temporal subordinate clauses like that of (3.281.d). They too can occur as adjuncts to nominal as well as to verbal projections. For instance, the TSC of (3.281.d) can also be incorporated into sentences like those in (3.291).

(3.291)a. some time after John had gone to Cambodia

b. a summer many years before the war

c. the last war before World War II

d. my hangover after last night's party

Confronted with the possibilities attested in (3.291) we could of course respond in the same way as we tacitly did in relation to *before*- and *after*-phrases: *before*- and *after*-clauses can have either the status of verbal or of nominal adjuncts; here we are only interested in their role as verbal adjuncts; so all we are concerned with is the semantics they must have when used in this capacity. If that is the line we take, then the semantic representations of *before*- and *after*-clauses should have the form of eventuality predicates, with an eventuality dref as referential argument. And by the same token, if that is the line we take, then there must be something in the syntactic structure of an adverbial TSC that makes it an adverbial TSC as opposed to an adjectival one: there are two syntactic TSC types – we might label them '<TSC,Adv> and <TSC,Adj> – that we can form by combining a TP with *before* or *after* as complementizers. In line with this proposal the label of the top node in (3.288) should be '<TSC,Adv>' and not just 'TSC'. And the construction of the semantic representation of the <TSC,Adv>-node of (3.288) should turn the semantic representation of the TP node in (3.292.a) into the representation in (3.292.b). This representation can then be unified with the lower TP node of the main clause in (3.289), as shown in (3.292.c).

An alternative way of analyzing the semantic contribution of the *before*-clause in (3.281.d) would be to assume that there is only one syntactic category TSC and that expressions of this category can be adjoined both the verb al and to nominal projections. In this case one would have to assume that the referential argument of the semantic representation of a TSC can be unified both with the possible referential arguments of the semantic representations of verbal and of nominal adjunction sites. Since the referential argument of the semantic representation of nominal adjunction site can be a dref that represents a time, as in (3.291.a,b). A similar line of attack would of course also be possible for *before*- and *after*-phrases: here too we could adopt a single category of adjectival or adverbial adjunct, the semantic representation of which would be of sufficiently general type to allow for unification with the referential arguments of all possible adjunction sites. This alternative approach would have the merit of capturing the generalization that at lest for the expressions in question what can be an adjectival adjunct can also be an adverbial adjunct and conversely. The only constraints on the referential arguments from the various adjunction possible sites is that they represent entities that can stand in temporal relations.

Although it would not be too hard to adopt this slightly more general approach, we stick to the less general path along which we have been proceeding, in which both *before*- and *after*-phrases and *before*- and *after*-clauses are treated as adverbials only. Hence the representations in (3.292.a,b.c).

$$(3.292)\text{a.}\quad < t,\ e_{ref},\ t^3,\ t^2,\ t',\ y| \begin{array}{|c|}\hline \\ \hline \text{24-th-of'}(t^3, t^2) \\ \text{August-of'}(t^2, t') \\ \text{79 A.D.'}(t') \\ t \prec n \quad e \subseteq t \quad e \subseteq t^3 \quad \text{“it”}(y) \\ e\text{: erupt'}(y) \\ \hline \end{array} >$$

$$\text{b.}\quad < ev_{ref}| \begin{array}{|c|}\hline t\ \ e\ \ t^3,\ t^2,\ \ t'\ \ y \\ \hline \text{24-th-of”}(t^3, t^2) \\ \text{August-of'}(t^2, t') \\ \text{79 A.D.'}(t') \\ t \prec n \quad e \subseteq t \quad e \subseteq t^3 \quad \text{“}it''(y) \\ e\text{: erupt'}(y) \\ \\ ev \prec e \\ \hline \end{array} >$$

$$\text{c.}\quad \begin{array}{|c|}\hline t^4\ \ e'\ \ x\ \ t\ \ e\ \ t^3\ \ t^2\ \ t'\ \ y \\ \hline \text{Mt.Vesuvius'}(x) \\ \text{24-th-of'}(t^3,\ t^2) \\ \text{August-of'}(t^2, t') \\ \text{79 A.D.'}(t') \\ \\ t \prec n \quad e \subseteq t \quad e \subseteq t^3 \quad y = x \\ e\text{: erupt'}(y) \\ t^4 \prec n \quad e' \subseteq t^4 \quad e' \prec e \\ \\ e'\text{: ‘spew-small-bits-of-rock'}(x) \\ \hline \end{array}$$

N.B. In (3.292.c) the pronoun *it* from the TSC has been resolved to its antecedent *Mount Vesuvius* in the main clause by identifying the dref y introduced by the pronoun with the referential argument x of the DP *Mount Vesuvius*. This of course is an operation for which we have no formal justification at this point.

3.10.6 Quantifying Temporal Adverbs and other Devices of Temporal Quantification

All the temporal adverbials that have been mentioned up to this point can be described as 'locating adverbs' in the following sense: each instance of them provides a time t or an event e that helps locate the eventuality ev of the clause of which that instance is a constituent, in the sense that ev must coincide with t/e or precede it or follow it or precede or follow it by some specified amount of time. Not all temporal adverbs are quite like this, however. There are also temporal adverbs that quantify over times, and which, you might say, stand to the adverbs we have looked at so far in the way that quantifying DPs like *every volcano* or *most abstracts* stand to referring DPs like *Mount Vesuvius* or *the worst abstract about temporal adverbs*.

For a start, consider the following set of sentence variants.

(3.293) Louise answers her letters always/usually/often/regularly/sometimes/ occasionally/rarely/never by e-mail.

These are well-formed and meaningful sentences no matter which of the adverbs *always,..., never* we consider; but each sentence has a different meaning. The differences have to do with how often and when Louise must answer her letters by e-mail in order for the sentence to be true. What some of the adverbs in (3.293) have to say about this seems quite clear and crisp, and can be expressed in our representation formalism straightforwardly. These are the adverbs *always*, *sometimes*, *never*, the adverbial counterparts of the nominal quantifiers *every*, *some* and *no*, which can be expressed straightforwardly within the first order predicate calculus (a version of which is included within our DRT-based representation formalism). The other adverbs from the enumeration in (3.293) have no simple truth-conditional characterizations. How often should Louise answer her letters by e-mail in order that the *often*-version of (3.293) can be counted as true? Or how many such e-mail replies by Louise must there be to make the *regularly*-version true, or the *occasionally*-version, or the *rarely*-version? These questions are reminiscent of those we asked when we discussed habitual and dispositional sentences in Section 3.7. And they are just as hard to answer as the questions that baffled us there. In Section 3.7 we adopted a portmanteau solution for habituals and dispositionals, by adopting the operators HAB and DISP but then saying next to nothing about the truth conditions of semantic representations containing these operators. Here, we adopt the same evasive policy, by simply setting these difficult quantifying adverbs aside and focusing on the three – *always*, *sometimes* and *never* – whose truth-conditional contributions that can be expressed with the formal tools that our semantic representation formalism (the DRSs language we are using) provides us. We will briefly return to the recalcitrant quantifying adverbs in the final subsection of Section 3.11, in the context of a more general discussion of the vagueness of temporal expressions. (But please, do not expect solutions to the problems that make me set aside those adverbs right now!)

The sentences in (3.293) are like the habitual and dispositional sentences of Section 3.7 in that their tense is the simple present. There is an important point to this choice in that, as we will see below, not all temporal adverbs that would intuitively seem to qualify as quantificational adverbs go with the present tense, or go well with it. That the ones occurring in (3.293) are perfectly natural in present tense sentences singles them out as quantificational adverbs of a certain kind. (More on that later.) But there is also a drawback to using present tense sentences as illustrations of the use of quantificational adverbs, which we can see more clearly when we consider their past tense counterparts. (3.294) gives the simple past tense counterparts of the sentences in (3.293) for the adverbs to which we have decided to restrict attention in this exploration.

(3.294) Louise answered her letters always/sometimes/never by e-mail.

When we ask ourselves under what conditions the sentences in (3.294) are true, one of the things that strike us is that we cannot answer the question until we have been told more about the part of the past that the sentences are supposed to be talking about. In this respect the sentences in (3.294) differ from those in (3.295), which provide an answer to this question.

(3.295) In the nineties/When she was an assistant professor, Louise answered her letters always/sometimes/never by e-mail.

For these sentences it is clear which part of the past they are targeting – for the first it is the last decade of the 20-th Century and for the second it is the period when Louise was an assistant professor. Because of this additional information the sentences in (3.295) with *never* and *sometimes* have clear and easily articulable truth conditions. For instance, the *never* version with the adverbial *in the nineties* is true if and only if there was no time during the nineties when Louise answered a letter by e-mail (and likewise for the sentences with *sometimes*). For the *always*-sentences in (3.295) there is a further complication in that here we need a non-vacuous restrictor of the quantification that *always* expresses. The sentence itself doesn't provide such a restriction, so context is needed to supply one. We will address this problem below.

In the sentences in (3.294) no definite period of time is specified within which the quantification is to apply. In this regard these sentences are, as they stand, underspecified. They will only take on definite truth conditions when used in contexts from which such a 'temporal quantification frame' can be inferred; all that can be inferred from the sentence itself is that this quantification frame must be located in the past of the utterance time n. The sentences in (3.293) pose a similar problem. They too fail to mention a temporal quantification frame, and only tell us that this frame must include n. And with these present tense sentences there is the additional problem that in English and many – perhaps all – other languages it is difficult to refer to such n-crossing intervals concisely. (We can say things like 'Nowadays, that is during this second decade of the 21-st Century, Louise answers her letters by e-mail.', but this is round-about and sounds a little awkward.) In fact, such present tense sentences are often used with the intention to *not* make the frame time fully explicit, perhaps because we have a sense that we do not have enough of a grip on the future to justify a commitment as regards for how long the generalization we are asserting is going to be valid. In this regard the sentences in (3.293) are much like the habitual and dispositional sentences of Section 3.7.

Even though the sentences in (3.295) are explicit about their quantification frame, there is another problem about their truth conditions. We will refer to this problem as the 'distribution problem (of adverbial quantification)'. It is a problem that generally arises for adverbial quantification, but not for the nominal quantification that we have so far been looking at exclusively in these Notes, in which the quantifying constituent is a DP. Quantification by means of DPs, we have seen, is unambiguous in the following respect: the way in which DP quantification is expressed in natural language syntax makes it clear which part of the information provided by the quantifying sentence goes into the restrictor of the quantifier – in DRT: the representing duplex condition – and which part goes into the nuclear scope. (This is arguably not completely true, since DP quantification too has its share of scope ambiguities, but once these ambiguities have been resolved in the way we have been assuming they can and should be – viz. in the LF that serves as input to construction of the semantic representation –, then the question which part of the sentence material goes into the restrictor and which part into the nuclear scope is fixed.) With adverbial quantification things are less straightforward. Some indication of the complexities of the distribution question is a famous pair of examples due to Rooth (1985), reproduced here as (3.296.a,b).

(3.296)(Italics indicate focal stress.)

a. In St. Petersburg officers always escorted *ballerinas*.

b. In St. Petersburg *officers* always escorted ballerinas.

The observation here is that (3.296.a), in which the focal stress is on *ballerinas*, has for its natural interpretation that during the relevant period any occasion when an officer escorted someone, that someone was a ballerina. In contrast, the natural interpretation of (3.296.b) with its stress on *officers*, is that whenever someone escorted a ballerina, that someone was an officer. At the present time the theory of Information Structure (of which the semantics of focus is one of the chapters) has developed to a point where it allows for a reasonably good understanding of how focus affects the interpretation of adverbial quantification (see in particular (Beaver & Clark 2008)). A rule of thumb, which explains the difference in truth conditions between (3.296.a) and (3.296.b), is that a focus constituent of a clause with a quantifying adverb goes into the nuclear scope of the quantifier, while the background goes into the quantifier's restrictor. But the rule doesn't always make the right predictions, and moreover, sensitivity to focus-background contrast is only one of a range of different mechanisms that are relevant to the distribution problem.

Another point that is relevant in this connection has to do with the difference between speech and writing. When the sentences in (3.296) are spoken, the prosody of the pronunciation, which will place focal stress on the italicized part, will make clear which part of the sentence material is meant to go into the restrictor and which into the nuclear scope. A reader therefore has to rely on other clues to determine which parts of the clause go into the restrictor and which into the nuclear scope. Often it is contextual information that provides the clues needed to make this choice. Another possibility is that readers often construct *underspecified* representations for such sentences –

underspecified in that they leave open whether certain constituents go into the restrictor or the nuclear scope of a quantifier they represent. Such representational underspecificity can and often will be resolved at some later stage, when the information needed for its resolution has become available, or when the need for it arises (for instance because the interpreter realizes that the differences in truth conditions that come with different possible resolutions are important for certain inferences he wants to draw). Underspecification and the reduction or eliminations of underspecified representations is a subject beyond that scope of these Notes. (For accounts of underspecification within the setting of DRT see for instance (Reyle 1993), (Reyle et al. 2007).) The problems connected with the representation of adverbial quantification are not restricted to questions of how to distribute overt sentence material over restrictor and nuclear scope. Consider the sentences in (3.297). (3.297.a) is one of the sentences from (3.295) and (3.297.b) is an elliptic version of that sentence.

(3.297)a. In the nineties Louise always answered her letters by e-mail.

b. In the nineties Louise always answered by e-mail.

Intuitively sentence (3.297.a) seems to mean that whenever during the nineties Louise received a letter, then she answered it by e-mail: the restrictor speaks of events of Louise receiving letters and the nuclear scope of her answering those letters by e-mail. But what are the principles that guide this identification of restrictor and nuclear scope? A related, and even more urgent question arises in connection with (3.297.b). Here it is no longer clear without help from the context what information should go into the restrictor: the information that Louise had received a letter? the information that she had been contacted in some other way? or the information that she had been contacted in some way that required a response? In the absence of a context that allows us to differentiate between these and other possibilities there just is no way of telling what restrictor may have been intended.

In fairness it should be conceded that most nominal quantifications also have to rely on context for their intended interpretations. It has been argued that every occurrence of a quantifying DP comes with an invitation to domain restriction: an invitation to amplify the quantifier restriction that the DP overtly supplies with additional material that is culled from the context in which the DP is used. There is a substantial literature on this topic. For a widely known approach see (**?**)StaSza:oqdr). But this kind of problem is clearly different from the problems posed by (3.297.a) and (3.297.b).

Although adverbial quantifiers often lead to quandaries over the distribution of sentence material, English and other languages have certain grammatical forms that come with strong presumptions as to where their material is meant to go. Probably the most familiar example of this are *when*-clauses. Consider (3.298.a). Here the *when*-clause material should go into the restrictor and the main clause material into the nuclear scope (while *in the nineties* gives the quantification frame). A variant of (3.298.a) is shown in (3.298.b). It shows how when the quantification is universal, the functions of quantifying adverb and *when*-clause can be combined into the single word *whenever*; as nearly as I can tell, (3.298.a) and (3.298.b) are alternative ways of expressing the same content.

(3.298)a. In the nineties, when Louise got a letter, she always answered by e-mail.

b. In the nineties, whenever Louise got a letter, she answered by e-mail.

However, *when*-clause material doesn't always go into the restrictor. As shown in (3.300) below, the contribution made by a *when*-clause need not always be the restrictor but can also be the quantification frame. But what *when*-clauses cannot do is contribute their material to the nuclear scope of the duplex condition that is introduced by a quantifying adverb. In this regard temporal subordinate clauses with *before* and *after* are different. They have a tendency to be interpreted as nuclear scope material. (This, by the way, is equally true of PPs in which *before* or *after* is the prepositional head; this is another bit of evidence that the constructions in which *after/before* is complemented by an ordinary DP and those where the complement is a clause aren' really different.) Examples are the sentences in (3.299). Both the sentence (3.299.a) and the sentence (3.299.b) are naturally understood as saying that whenever in the nineties John went to the airport to catch a flight, he did so always less than two hours before the departure of that flight.

(3.299)a. In the eighties John always went to the airport less than two hours before his flight departed.

b. In the eighties John always went to the airport less than two hours before (the departure of) his flight.

But *before*- and *after*-clauses and -phrases are no less immune to the effects of information structure than other constituents. For instance, a natural interpretation of (3.300.a) (if perhaps not the only one) is one in which the adverb – *before noon*, *after breakfast*, *in the morning* – is the material that makes up the nuclear scope, while the material between it and the coma which separates main clause from *when*-clause goes into the restrictor. The sentence suggests that Louise has been running for a long period of time, which includes the time when she was a student and that when she was a student the running she did then was done before noon/after breakfast/in the morning. (The *when*-clause in (3.300.a) serves to determine the quantification frame.) Moreover, the adverbs *before noon*, *after breakfast* and *in the morning* also imply that the quantification is over the days that make up the period contributes by the *wham*-clause; this is because the nouns]em noon, em breakfast, *morning* are all naturally interpreted as relational nouns; trey raise the question ‘ noon/breakfast/morning of which day?’. (Quantification by temporal adverbs generally requires a temporal interval as quantification frame as well as a ‘granularity’ in the form of a partition of they interval into suitably sized convex portions; the quantification is then over this partition (i.e. over the set of all these portions of the interval). For details see ((Reyle et al. 2007)).)

(3.300)a. When she was a student, Louise always ran before noon/after breakfast/in the morning.

b. When she was a student, Louise always had a run before noon/after breakfast/in the morning.

c. When she was a student, Louise always ran around the block before noon/after breakfast/in the morning.

d. In her old age Louise always worked before lunch and had a nap after lunch.

Judgments about these sentences may vary somewhat. Here are, for what they are worth, my own intuitions. In (3.300.a) the adverb *before noon/after breakfast/in the morning* is naturally understood as the material that makes up the nuclear scope and the remainder of what follows the comma as material for the restrictor. (3.300.b) and (3.300.c) can also be interpreted this way, but here there is (for me) a strongly competing interpretation, according to which Louise had the habit of running around the block before noon etc. when she was a student, without any implications about her running during other periods of her life. According to this interpretation not only the adverb goes not the nuclear scope but also the contribution made the verb phrase, *have a run* in (3.300.b) and *run around the block* in (3.300.c). More precisely, this means that in these interpretations the nuclear scope gets the eventuality drefs introduced by the verbs of (3.300.b) and (3.300.c) together with their characterizations, as a 'having-a-run' event or as a 'running-around-the-block' event. The restrictor in these interpretation is made solely of conditions that restrict the quantified dref t to taking values that are calendar days within the quantification frame that is supplied by the *when*-clause.

(3.300.d), finally, is easily read in a way that puts the adverbs *before lunch* and *after lunch* into the restrictors of the two quantifications expressed by its two conjuncts. (This is thus the reading according to which the sentence says what Louise did on each of the days belonging to its quantification domain before lunch was work and that what she did after lunch on each of this days was have a nap.)

The difference between (3.300.a) on the one hand and (3.300.b,c) on the other is of interest insofar as it shows how the distribution problem can be affected by subtle differences in aspect – *run* is non-telic (an 'activity verb' in Vendler's terminology) – whereas *have a run* and *run around the block* are telic ('accomplishment verbs', in Vendler's terminology). On the other hand the difference between (3.300.a,b,c) in all of which the locating adverb is naturally interpreted as part of the nuclear scope, and the mentioned reading of (3.300.d), which puts the locating adverb into the restrictor, shows that such adverbial constructions are subject to the effects of Information Structure no less than other constituents. Finally, the sentences in (3.300) give a flavor of the complexities when quantifying temporal adverbs and temporal looting adverbs cohabit in the same clause.

Next we turn to an issue that is somewhat tangential to the general purpose of Section 3.11 – that of giving an aperçu of the different kinds of temporal adverbs. Not all quantifying adverbs are temporal adverbs. As Lewis observes in a seminal paper on natural language quantification ((Lewis 1975)), the equivalence between, for instance, *every* and *always* as general devices of universal quantification (rather than quantification in some restricted domain like that of temporal quantification) was taken for granted by some of the leading logicians of the first half of the 20-Century. (Lewis draws particular attention to Russell.) Among Lewis' own examples of non-temporal uses of *always* are mathematical sentences like the following.

(3.301) A quadratic equation always has at least one root.

The quantifications in such sentences have nothing to do with time (certainly not in a literal sense). Lewis concludes from examples like this one that *always* and other quantificational adverbs can bind variables that range over any kinds of entities; times are among those kinds, but do not enjoy a preferential status. Lewis also observes that as they are used in natural language sentences, adverbial quantifiers like *always* often bind several variables at once – so-called 'non-selective binding' in his terminology – as we can see for instance in the donkey sentence variant 'If a farmer owns a donkey he always beats it.' On Lewis' analysis of this sentence the *always* of this sentence binds both the variable that represents the farmer and the variable that represents the donkey. (A variant of the 'non-selective binding mechanism' can be found in DRT, where, to mention just one example, the antecedent DRS of the condition representing a conditional may contain several drefs, all of which get 'universally' bound as part of the verifiability conditions for conditional DRS conditions.) However, in the treatment of quantifying adverbs like *always* that we adopt here we restrict attention to their use as temporal quantifiers, which bind one dref at the time and where the bound drefs are invariably drefs representing times.

Not only can adverbs like *always* be used to express quantification over things other than times; for their part, nominal quantifiers, of which we have only considered non-temporal examples up to now, can be used to quantify over times just as they can be used to quantify over other kinds of things. Two examples are given in (3.302).

(3.302)a. When she was a student, Louise played volleyball every week.

b. When she was a student, Louise played volleyball every Wednesday.

c. When she was a student, Louise always played volleyball on Wednesday.

(3.302.a) asserts that every calendar week within the period when Louise was a student contained an event of her playing volleyball. This, in other words, is a quantification over successive, week-size portions of that period. Much the same is true of (3.302.b). Here the restrictor of the quantifier *every* isn't the predicate 'week', but instead the predicate 'Wednesday', so the domain of quantification is the set of all Wednesdays within the given period. Compare these two sentences with (3.302.c). Intuitively (3.302.c) seems to be saying much the same thing as (3.302.b). But it says it in a different way. As in some of our earlier examples in this section the PP *on Wednesday* is naturally understood as part of the nuclear scope of the quantificational structure imposed by *always*. But if that is where this PP goes in the semantic representation of (3.302.c), what is there to fill the restrictor?

The answer that comes to mind (and the only one that comes to mind in the absence of overriding context information) is that the times that are being quantified over are the calendar weeks of the period when Louise was a student – just, in other words, as in (3.302.a). But while in (3.302.a) this aspect of interpretation is made explicit by the form of the nominal quantifier, in (3.302.c) it is implicit and must be somehow inferred from the information that is overtly present in the sentence. At an informal level it is clear what is involved in this inference: if the nuclear scope speaks of something happening 'on Wednesday', then the times represented by the dref that is bound by the temporal quantifier must be portions that each contain a unique Wednesday, which the 'Wednesday'-condition in the nuclear scope can then select as the unique Wednesday within any one of these portions. Thus the inference takes the form of a kind of 'abduction': by assuming that the restrictor of the quantification contains the condition that the times quantified over are calendar weeks, the occurrence of *on Wednesday* in the nuclear scope can be justified; and this would seem the simplest and most natural way of providing a justification for the occurrence of *on Wednesday* there. In fact, the predicate *Wednesday* is similar in tho regard to the predicates *noon*, *breakfast* and *morning* in (3.300): just as these nouns point towards a quantification granularity of days, *Wednesday* points to a granularity of weeks.

By way of a summary of the informal discussion in this section of quantifying temporal adverbials we note three features that distinguish such quantificational devices from the nominal quantifiers that have been considered in various earlier parts of these Notes. Each temporal adverbial quantification involves:

(i) a *temporal quantification frame*: a period of time over which the quantification extends in the sense that the values of the quantificationally bound time dref must all be included within the frame. Often this period is implicit, and this is so in particular when the sentence containing the quantifying adverbial is in the present tense. But a temporal quantification frame is always part of the interpretation of such adverbials. When no frame is mentioned explicitly or reconstructible from the context, the sentence comes across as vague or as semantically underspecified.

(ii) the *distribution* of explicit and implicit content material over the restrictor and the nuclear scope of the quantification. The principles according to which restrictor and nuclear scope are filled appear to be diverse and are still poorly understood.

(iii) Often temporal quantification involves a choice of *granularity*: the quantification is over the parts of a certain partition of the temporal quantification frame, where these parts are all of a certain (typically calendar-related) size – the 'grain size' of the given granularity. Currently our understanding of the principles that govern the determination of granularity is also no more than partial.

3.10.7 Representation Constructions for some Adverbial Temporal Quantifications

One of the problems with which adverbial quantification confronts us, we noted, is how the material from a clause containing an adverbial quantifier is to be distributed over its restrictor and its nuclear scope. When it comes to the computation of the semantic representations of such clauses within the framework we have been developing and using, the first question we need to answer is whether the information about what goes into the nuclear scope and what into the restrictor is explicit in the LF for the clause, or has to be inferred, by whatever means, when the semantic representation is derived from that LF. What is more, answers to this question need not be as straightforwardly black or white as the question suggests. It could also be that the LFs for such clauses provide some clues that help divide the clause material into restrictor material and nuclear scow material, while still leaving some decisions to the construction algorithm.

We saw that among an important factor in the distribution over restrictor and nuclear scope is Information Structure, and more particularly that the distribution often follows focus-background division. (Recall the discussion around the examples of Rooth in (3.296.a,b).) Since this division is usually marked by prosody in spoken English, it would not be an unreasonable assumption that the LFs for such sentences make this information explicit. In fact, assumptions along these lines are well-known from the literature, using some forms of Focus marking of syntactic constituents that are prosodically marked as in focus (among them Rooth's own proposal in ((Rooth 1992))). But we also noted that in written language the focus-background division is just as important for restrictor-nuclear scope distribution. But here the division is usually not prosodically marked. And usually this information can somehow be recovered, as people usually manage to do when they read out a written text and for the most part get the prosodic focus markings right as they go along. As fas I know there is no real understanding within the linguistics community of speakers are able to do this. But there can be little doubt that they succeed as well as they do because they rely on semantic information that in the kind of processing architecture we are assuming will for the most part be available *after* the conversion from LF into semantic representation has taken place, or at least a good part of it.

In the light of these observations it is clear that what can be offered here in the way of DRS construction for sentences with adverbial quantification will have to be quite limited. To get a better idea of what we may be able to do let us turn to a concrete example. We start with sentence (3.298.a), which is repeated below.

(3.298.a) In the nineties, when Louise got a letter she always answered by e-mail.

We assume for this sentence the LF shown in (3.303). (The display is split in two, since otherwise the structure wouldn't fit onto the page.) This LF assumes that the quantifying adverb em always is a TP adjunct, and thus high up in the tree, and that the adverbial that specifies its quantification frame is adjoined above it. The material that has to be distributed over the restrictor and the nuclear scope of the quantification expressed by *always* is to be found in the adjunction site of *always*.

(3.303)a.

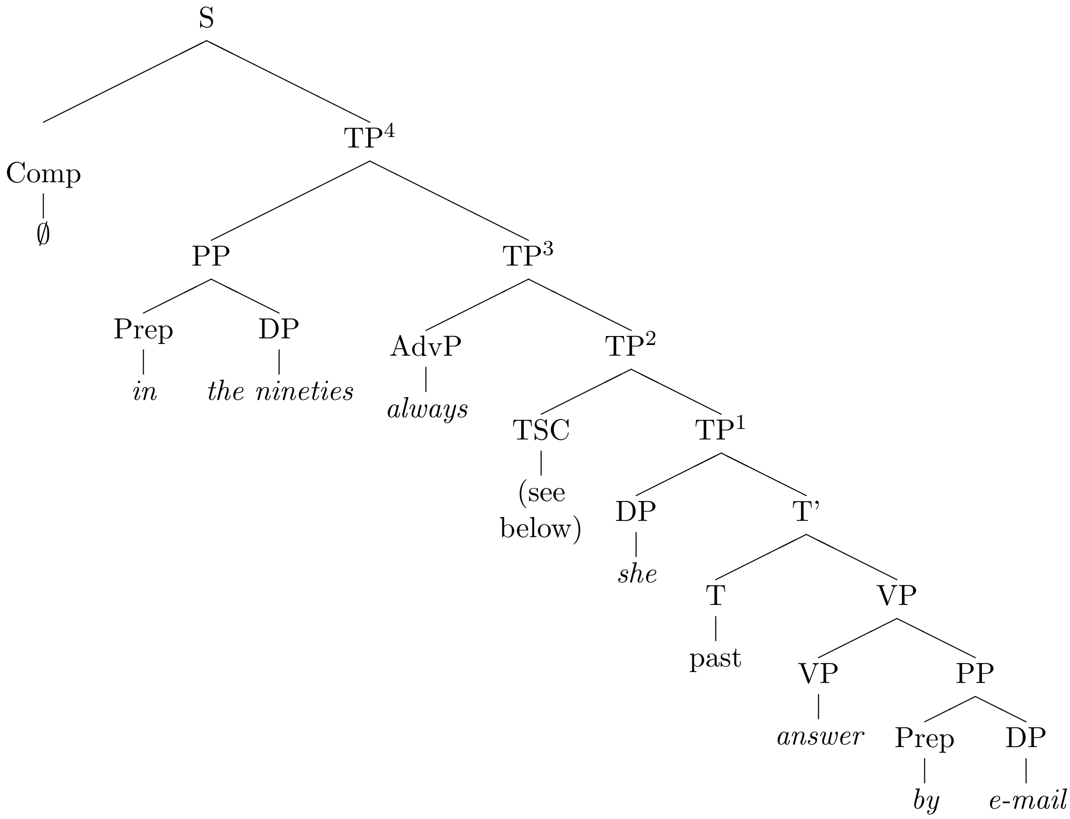

b.

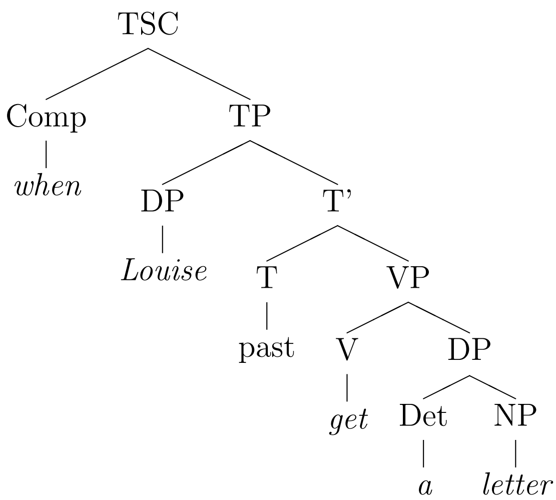

Intuitively two distributions seem possible for this LF. The first is that which puts the *when*-clause (whose highest node is the TCS node) into the restrictor and the remainder of the adjunction site of *always* into the nuclear scope. A second possibility, intuitively less prominent perhaps, is that according to which only the PP *by avail* goes into the nuclear scope and all the other parts of *always*'s adjunction site into the restrictor. This reading does not require that each time during the nineties when Louise got a letter she replied to it. Without explicit contextual information that licenses such a reading it is not easy to get it, and perhaps even impossible. But (3.304) gives a context, provided by its first sentence, in which this reading does seem possible.

(3.304) When people contacted Louise she usually replied. When she received a letter she answered by email.

How can these two readings be formally obtained from the LF in (3.303)? The proposal that follows is clearly incomplete, for the reasons mentioned in the opening paragraphs of this section. But let us forge ahead and discuss the limitation after we are done.

We assume that LFs for sentences involving adverbial quantification are first transformed into structures in which the distribution over restrictor and nuclear scope is made explicit. From such structures the semantic representation can then be constructed in much the way that we have seen throughout Section 3. I will refer to the results of these transformations as LF's. But this should not be construed as an expression of the view that these are 'syntactic' structures. They are simply the first step on the way from syntactic to semantic representations.

An LF' is obtained from an LF by moving the constituent or constituents that will end up in the nuclear scope from their position in the 'scope' of the adverbial quantifier (i.e. the adjunction site of the adverb) into a position to the right. That is, the adjunction site is split into two parts, a restrictor part and a nuclear scope part; the constituent or constituents in question are moved into the nuclear scope part while what is not moved from the adjunction site of the quantifying adverb forms the restrictor. The LF' codes the source positions of the constituents that are moved in the transition from LF to LF' through co-indexation between the moved constituents with copies of their top nodes that are left in situ.

Here we will only consider cases where just one constituent is moved into the nuclear scope part. The LF' for the first of the mentioned two readings for (3.298.a) results when the TP^1 node is moved into the nuclear scope part, so that the TSC is all that remains of the adjunction site. The result is shown in (3.305). Again the tree is broken up, for technical reasons, and this time into three parts.

(3.305)a.

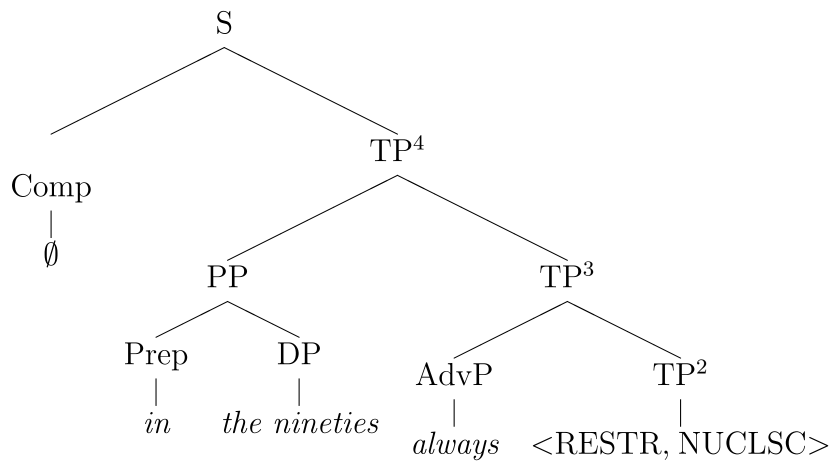
S
Comp
∅
TP^4
PP
Prep
in
DP
the nineties
TP^3
AdvP
always
TP^2
<RESTR, NUCLSC>

b.

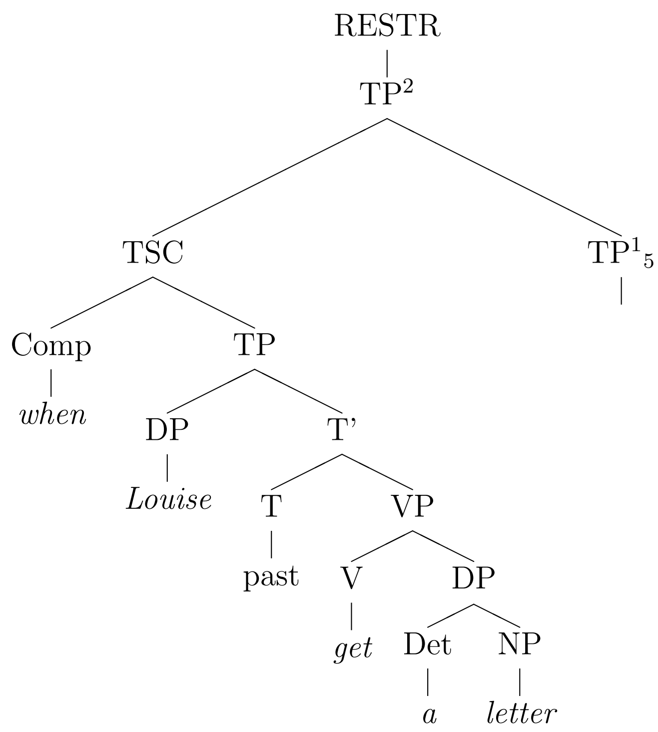
RESTR
TP^2
TSC
Comp
when
TP
DP
Louise
T'
T
past
VP
V
get
DP
Det
a
NP
letter
TP$^1{}_5$

c.

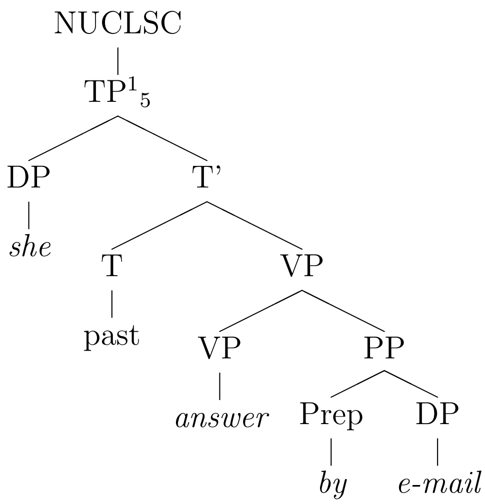

The semantic representation for the NUCLSC part of (3.305.a), shown in (3.305.b), can be constructed using principles that have been discussed. The result is shown in (3.306.a). The semantic representation for the RESTR part must capture the semantic contribution of *when*. We have seen that this is problematic in that the semantics of *when* is not purely temporal but involves a causal momentum as well. In the representation of the RESTR part in (3.306.b) I have adopted the simplification that when a *when*-clause that functions as temporal locating adverb itself describes an event (as is the case for the *when*-clause) of (3.298.a), then (a) when the eventuality that is temporally located by the *when*-clause is an event, then this eventuality follows the *when*-clause event and (b) when the located eventuality is state, then this state temporally includes the *when*-clause event.

There is also another question about the representation of the *when*-clause in (3.306.b). This question is not about the specific meaning of *when* (as distinct from other temporal conjunctions) but with the role of TSCs as restrictors to adverbial quantifiers. In our informal discussions in the last section we were confronted more than once with the question what kind of 'variable' (or dref, in the formal framework in which we are working) is bound by the quantifier. According to what our discussion has been implying, there are cases in which this variable isn't even explicitly represented in the semantic representations from which the duplex condition is to be constructed that represents the adverbial quantification. Examples are those where quantification over the calendar days within the temporal quantification frame is inferred from the occurrence of relational nouns like *morning* or *breakfast*. These examples indicate that a fairly complex algorithm (or set of algorithms) can be involved in determining the dref that is to figure as 'bound variable' (i.e. as the dref appearing in the central diamond of the duplex condition; for more on this see the paragraph immediately below the diagrams in (3.306)). I have no good idea what this general algorithm or algorithm set is like. So on this point the construction resented here is necessarily going to be a kind of stop-gap.

Among the cases that such an algorithm (or one of the algorithms in the set) would have to be able to deal with are those of which our present sentence (3.298.a) is an instance – those in which the bound variable is provided by a TSC and where, moreover, this dref is the referential argument of the highest TP node of that TSC. In the example before us this is the dref e that represents the event of Louise answering (the letter she had received). This presents us with a formal representation problem (as well as the more fundamental problem how this dref is selected as bound variable): We now have a competition between two drefs as 'referential argument' of the TSC representation that occupies the RESTR part of (3.305.a), (a) the referential argument ev that is waiting for unification with the referential argument of its adjunction site (the TP^1 node of (3.303.a)) and (b) the dref e that will be needed as 'bound variable', in the same way that the referential argument of a quantifying DP is used as bound variable in the duplex condition to which such a DP gives rise. As things stand I have no good solution to this representational problem either. I will assume, as a more or less ad hoc solution, that the store of the TSC representation in (3.306.a) contains both the drefs ev and e and that e is marked with the subscript $_{qbv}$ (for 'quantificationally bound variable') to make sure that this is the dref that will go into the central diamond, while the $_{ref}$-indexed dref ev will be the role of the referential arguments of temporal locating adverbs, viz. that of unifying with the eventuality for which the adverb provides a temporal location.

One final remark on the unification of the referential argument ev of the TSC and the referential argument e' of the TP^1 node. So far such unifications were past of an operation in which they were followed by merge to the two representations to which the unified drefs belonged. But in the construction that we are considering here there should of course be no merge. In this regard the present construction, in which the two representations to which the unified drefs belong end up in the two DRSs that make up the duplex condition, is to the cases of unification considered in Section 3.11 hitherto like the cases of quantified donkey sentences (like 'Every farmer who owns a donkey beats it.') stand to donkey discourses (like 'John owns a donkey, He beats it.'); in the latter cases there is merge, in the former there is not.

This difference between the present case and the earlier cases in which the referential argument of a locating adverb unified with the referential argument of its adduction site brings out with perhaps greater clarity that the unification operation and the merge operation, which until now may have seemed formal operations that were part of a single interpretation step, are importantly distinct. Unification is still as before: the two drefs are unified, with the effect that the two representations from which they come now de facto share the same dref. But there is now also a slight complication about the formal implementation of the unification. On the one hand the referential argument e' of the representation of TP^1 node has to unify with the referential argument ev of the TSC and thereby partake in the condition that relates ev to the referential argument e of the representation for the TP of the TSC. (The relevant condition in the present case is '$ev \prec e'$'.) But on the other hand e' has to be 'declared' in the nuclear scope DRS of the duplex condition – it should end up in the Universe of this DRS; for if not, then, as isn't hard to see, our representation wouldn't capture the right truth conditions. As rings stand these are incompatible requirements. They can be made coherent only by transferring the condition or conditions that contain ev before unification – here the one condition '$ev \prec e'$'. – from the TSC representation to the TP^1 representation. This specification of what has to be done with the TSC representation and the TP^1 representation when they serve as joint input to the adverbial quantifier *always* has an ad hoc flavor. It too should be reconsidered in the context of a more comprehensive and systematic account of adverbial quantification.

(3.306)a. $$< ev_{ref}, e_{qbv}, y, x \mid \begin{array}{|c|}\hline t' \\ \hline \begin{array}{|c|}\hline \\ \hline \text{Event}(ev) \\ e \prec ev \\ \hline \end{array} \overset{!}{\vee} \begin{array}{|c|}\hline \\ \hline \text{State}(ev) \\ e \subseteq ev \\ \hline \end{array} \\ t' \prec n \quad e \subseteq t' \\ \text{Louise'}(x) \quad \text{letter'}(y) \\ e\text{: get'}(x, y) \\ \hline \end{array} >$$

b. $$< e'_{ref},\ z \mid \begin{array}{|c|}\hline \\ \hline t \prec n \quad e' \subseteq t \quad \text{“she”}(z) \\ e'\text{: answer'}(z) \quad \text{by-e-mail'}(e') \\ \hline \end{array} >$$

The next step is a new one. Some aspects of it have already been discussed and what will be said in this paragraph will in part be a repetition of that; but there is no real harm in this, I believe. In this step, the quantifying adverb *always* is combined with its sister, the TP^2 node of (3.305). The semantic representation of this sister node is the pair of semantic representations of the structures labeled ‘Restrictor’ and ‘Nuclear scope’. It is part of the lexical semantics of *always* (and other temporal quantifying adverbs) that this is the right kind of input for it. What a quantifying adverb does with such an input is by and large familiar at this point from our earlier discussions of nominal quantification in Section 9.3.1: the quantifier introduces a dref a new ‘quantificational’ state which is characterized by a duplex condition. This dref becomes the referential argument of the output representation. The duplex condition has, as always, a semantic quantifier in its central diamond position, which captures the specific content of the quantifying adverb (as distinct from other such adverbs). The restrictor and nuclear scope DRSs of the duplex conditions are determined by the two parts of the input on which the quantifier operates. In the case of quantifying DPs the first part is given by the NP of the DP and the second part by the DP’s sister. In the case of an adverbial quantifier the parts are the representations of the RESTR and the NUCLSC constituents of the adjunction site of the adverb in the LF’. In both cases one constraint on the possible values of the dref bound by the central quantifier is that they must be temporally included in the duration of the quantificational state (see section 3.9.1 for discussion).

In the light of these extensive (and not optimally satisfying) comments on the operations that are involved in going from the representations in (3.306) to the representation of the TP^3 node of (3.305.a).

(3.307) $$<s_{ref} \mid \begin{array}{|l|}\hline s: \begin{array}{|c|}\hline t' \;\; e \;\; y \;\; x \\ \hline t' \prec n \quad e \subseteq t' \\ \text{Louise}(x) \\ \text{letter'}(y) \\ e\text{: get'}(x,y) \\ e \subseteq s \\ \hline \end{array} \;\; \underset{e}{\langle \forall \rangle} \;\; \begin{array}{|c|}\hline t \;\; e' \;\; z \\ \hline t \prec n \quad e' \subseteq t \quad e \prec e' \\ \text{“she”}(z)\text{”} \\ e'\text{: answer'}(z) \\ \text{by-email'}(e') \\ \hline \end{array} \\ \hline \end{array} >$$

Intuitively this last representation is unsatisfactory insofar as it fails to establish enough of a connection between the nuclear scope part pf the dullex condition and its restrictor part. Intuitively it is clear that the pronoun *she* should be understood as anaphoric to *Louise*. We can correct this by adding the condition '$z = x$' to the nuclear scope DRS (much as we have been doing on a couple of previous occasions, although the proper setting for this si waiting for Section 4). But there is also another obvious connection between restrictor and nuclear scope. The 'answering' events of which the nuclear scope speaks are clearly events of answering the letters spoken of in the restrictor. This relationship may be somewhat concealed at the level of surface grammar but it must be plain to anyone reading sentence (3.298.a). What really ought to have been done to capture this connection between restrictor and nuclear scope is to treat the occurrence of *answer* in sentence (3.298.a) as an instance of a transitive verb – perhaps all occurrences of *answer* should be analyzed this way – with an unfilled direct argument slot, the filler of which has to be recovered from context. When the interpreter is offered consists of sentence (3.298.a) on its own one, then the restricting *when*-clause is all that he gets by away of context; in such a situation everything points to the received letter as antecedent and nothing points away from it. So the letter represented as y is the obvious filler.

The instruments needed to properly model this second instance of context-based resolution are even more obviously missing from the tool box with which we are making do right now. But we can add conditions that fill the gap almost as easily as we can do that to capture the anaphoric relation between *she* and *Louise*. (3.308) shows the improved version of (3.307).

(3.308)

$$<s_{ref} \mid \boxed{ s: \begin{array}{|c|}\hline t'\ \ e\ \ y\ \ x \\ \hline t' \prec n \quad e \subseteq t' \\ Louise(x) \quad letter'(y) \\ e: get'(x,y) \\ e \subseteq s \\ \hline \end{array} \ \left\langle \begin{array}{c} \forall \\ e \end{array} \right\rangle \ \begin{array}{|c|}\hline t\ \ e'\ \ z \\ \hline t \prec n \quad e' \subseteq t \quad e \prec e' \\ z = x \\ e': \text{answer'}(z,y) \\ \text{by-email'}(e') \\ \hline \end{array} } >$$

The remainder of the construction consists in locating the state s of (3.308) with the help of the PP *in the nineties*, whose LF is shown in (3.309) and for which a simplified semantic representation is shown in (3.310). The contribution that the PP *in the nineties* makes to the semantics of sentence (3.298.a) is correctly captured by the principle that locating a state takes the form of a condition to the effect that the time specified by the adverb is included within the state. Applying this principle to obtain the representation of the TP^4 node of (3.305) and then transferring, as last step of the construction, the remaining drefs in the store to the Universe of the maine DRS, we obtain the structure in (3.312).[72]

[72] Once again we are skirting a problem that has been loitering on the outskirts ever since we introduced the habitually predicate HAB in Section 3.7.2. We skirted the problem again in Section 3.9.1 where we looked at the interaction between nominally expressed universal quantification and temporal reference. In the cases considered in those sections and also in the one that is before us now we have been dealing with states that are, in some quite broad sense, quantificational. When such states get located by a temporal adverb, like *in the nineties* in our present example, then the effect of this is typically that the adverb gives the exact duration of the state and not just some lower bound to that duration, in the way it does according to (3.312). To do justice to this intuition the locating condition for the state s in (3.312) should not have been '$t^2 \subseteq s$' but '$dur(s) = t^2$'.
Here are two reasons why location conditions of this latter kind are the right ones for quantificational states. First, in general the weaker condition according to which the time of adverb is included in the quantificational state will in general not give the right truth conditions. Consider for instance the sentence 'Last month Louise went for a run on most days'. Suppose that in fact Louise went for a run on slightly less than half of the days of last month, but that she went for a run on every day preceding last month and on every day following last month. Then there would be a state whose duration would span the period starting one week before last month and ending one week after it which satisfies the condition that within it on most days Louise went for a run. So a representation in which the quantification state is only restricted by the condition that its duration must include last month would be verified by this scenario. But intuitively it is ascension in which the sentences is false.
The second reason is of a different kind. It is the intuition that temporal quantification always involves a temporal frame. More on this below. Universal quantification is special

(3.309)

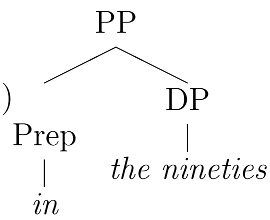

(3.310) $< t^2 \mid \boxed{\text{“the nineties”}(t^2)} >$

(3.311)

$$\begin{array}{|c|}\hline t^2 \quad s \\ \hline \\ \text{“the nineties”}(t^2) \qquad t^2 \subseteq s \\ s\text{:} \begin{array}{|c|}\hline t' \;\; e \;\; y \;\; x \\ \hline t' \prec n \quad e \subseteq t' \\ Louise(x) \quad letter'(y) \\ e\text{: } get'(x, y) \\ e \subseteq s \\ \hline \end{array} \; \left\langle \begin{array}{c} \forall \\ e \end{array} \right\rangle \; \begin{array}{|c|}\hline t \;\; e' \;\; z \\ \hline t \prec n \quad e' \subseteq t \quad e \prec e' \\ z = x \\ e'\text{: answer'}(z, y) \\ \text{by-email'}(e') \\ \hline \end{array} \\ \hline \end{array}$$

Let us now have a brief look at the other interpretation we identified as possible for sentence (3.298.a) (if dispreferred except in special contexts). this was the interpretation in which it is only the manner of replying – by

in that here the problem just described does not arise: whether the universal quantification is said to hold throughout a state that temporally includes a given time t or through out a state whose duration is exactly t comes to the same thing. This is why we can get by with the representation in (3.312) as it is. But this is, you might say, a case of lucking out.
Implementing the stricter location conditions for quantificational states for which this footnote makes a case would not be difficult. One way to do this would be to change the construction principles that introduce quantificational states in such a way that the state drefs they introduce are marked as representing quantificational states (e.g. by a subscript ‘$_{quant}$’ or something to that effect). The rule that governs the combination of temporal locating adverbials with their adjunction sites can then be restated in a way that is sensitive to the presence of this mark.

email – that makes up the nuclear scope of the quantification; all other material in the adjunction site of *always* becomes part of the restrictor. The LF' structures are given in (3.312).

(3.312)a.

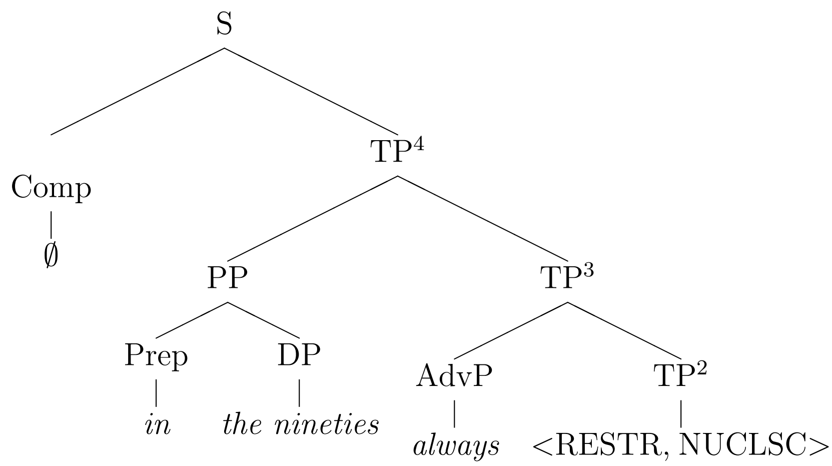

b.

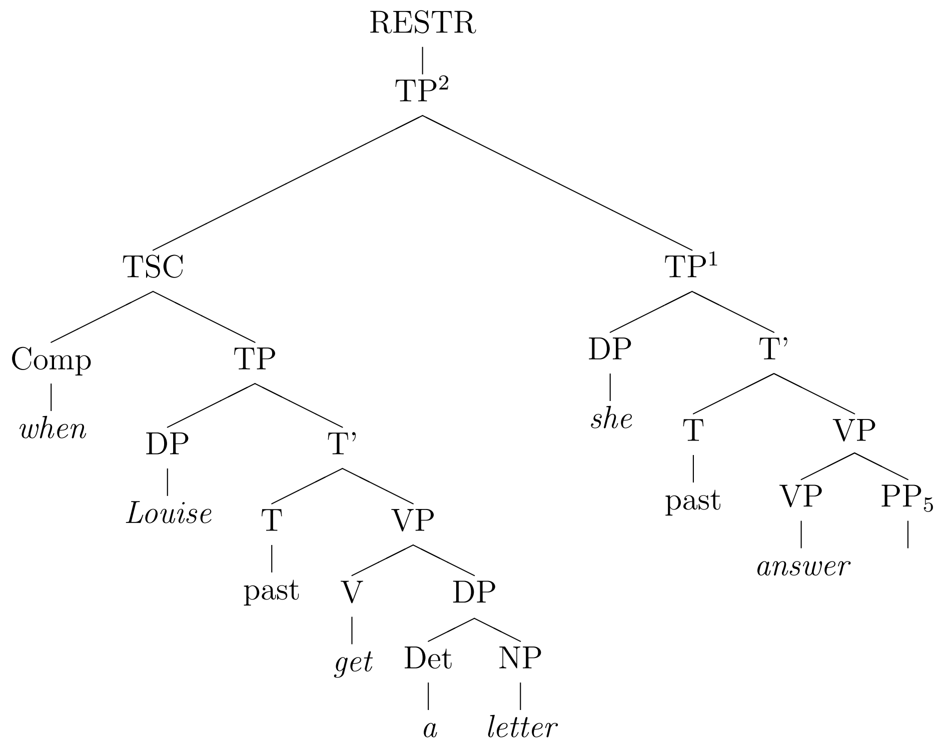

c. [NUCLSC [PP_5 [Prep *by*] [DP *e-mail*]]]

This time the representation of the RESTR node can proceed as in LFs where there is no separation of the kind we find in LF' structures. In particular, the unification of the TSC representation and the representation of the TP^1 node now proceeds as in the constructions we considered before we started our discussion of adverbial quantification. This time the construction step that is effected by the division into the RESTR and the NUCLSC part is the one in which the representation of the PP *by email* is unified with its adjunction site in the LF (3.303) (the VP of the main clause). Note that this unification occurs at the same point in the construction at which it would have occurred if no splitting into RESRT and NUCLSC had taken place in the transition from LF to LF'. To give an ida of the details, (3.313.a) gives the semantic representation of the NUCLSC part before unification of PP representation with the representation of its adjunction site, while (3.313.b) shows the relevant part of the representation of the RESTR part at that same stage. (The missing part of this diagram, the semantic representation of the TSC, is the same as in the previous construction. It is given once more explicitly in (3.313.c).)

(3.313)a. [NUCLSC [PP_5]]

$$< ev'_{ref} \mid \begin{array}{|c|}\hline \\ \hline \text{by-e-mail'}(ev') \\ \hline \end{array} >$$

b.

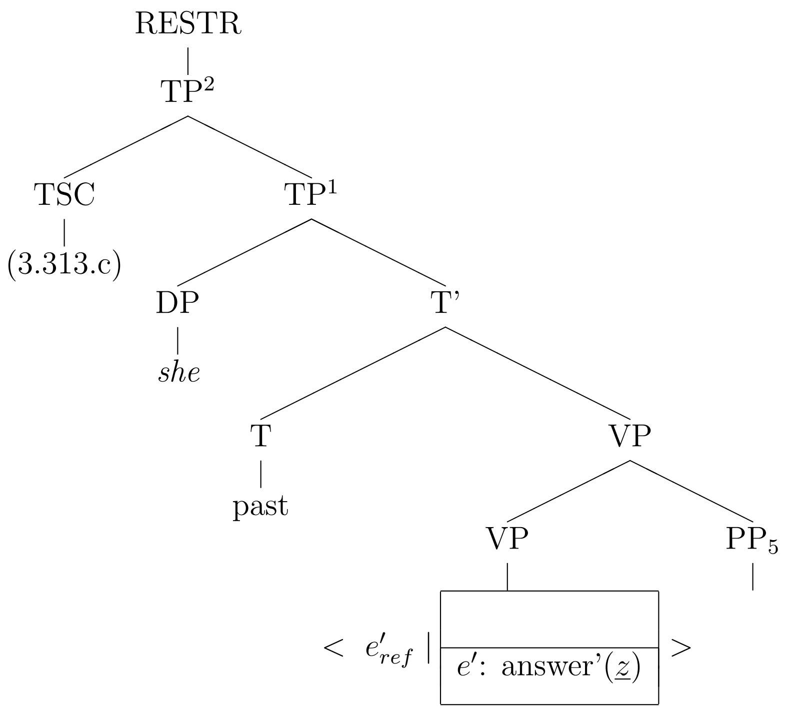

c.

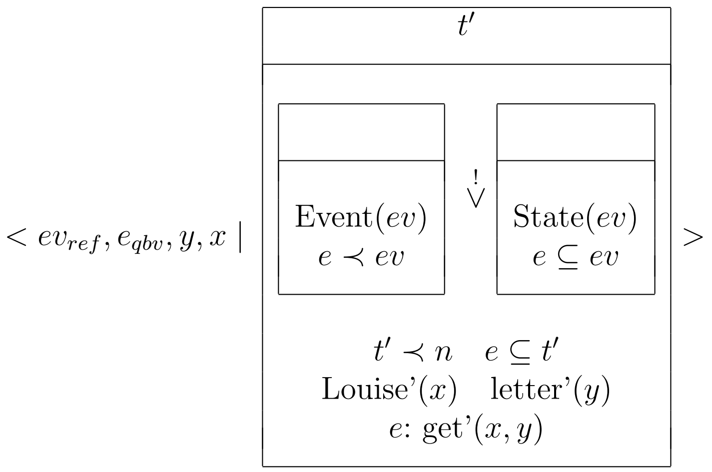

At the point where the representation constructions of RESTR part and NUCLSC part have reached the stages shown in (3.313) the representation of the PP under NUCLSC is combined with that of the lower VP under RESTR. The result of this is simply the unification of the two referential arguments, with the effect that ev' in the PP representation gets replaced by e' and the store of the PP representation, now empty, disappears. This operation only produces an actual change in the PP representation, the result of which is shown in (3.314).

(3.314)

$$\begin{array}{c} \text{NUCLSC} \\ | \\ \text{PP}_5 \\ | \\ \begin{array}{|c|} \hline \\ \hline \text{by-e-mail'}(e') \\ \hline \end{array} \end{array}$$

The remaining steps hold no further surprises after all we have said about the representation construction for the other reading of (3.298.a). The final result is (3.315).

(3.315)

$$\begin{array}{|l|} \hline \qquad\qquad t^2 \quad s \\ \hline \qquad \text{“the nineties”}(t^2) \qquad t^2 \subseteq s \\ s: \begin{array}{|c|} \hline t' \;\; e \;\; y \;\; x \;\; t \;\; e' \;\; z \\ \hline t' \prec n \quad e \subseteq t' \\ \text{Louise'}(x) \quad \text{letter'}(y) \\ e\text{: } get'(x,y) \\ e \subseteq s \\ \\ t \prec n \quad e' \subseteq t \quad e \prec e' \\ z = x \\ e'\text{: answer'}(z,y) \\ \hline \end{array} \left\langle \begin{array}{c} \forall \\ e \end{array} \right\rangle \begin{array}{|c|} \hline \\ \hline \text{by-email'}(e') \\ \hline \end{array} \\ \hline \end{array}$$

For our DRS constructions we consider the third version of sentence (3.300.b). The sentence is repeated here and given its own label.

(3.300.b.3) When she was a student, Louise always had a run in the morning.

We saw that the sentences in (3.300.b) and (3.300.c) are also ambiguous with regard to what goes into the restrictor and what into the nuclear scope. This is true in particular for (3.300.b.3). One of its readings is that in Louise's student days the following always happened: she had a run in the morning. In this reading all the material of *Louise had a run in the morning* is in the nuclear scope and it may not be immediately clear what goes into the restrictor.According to the second reading it is only the qualification *in* *the morning* that goes into the nuclear scope, while the otter material from *Louise had a run in the morning* goes into the restrictor.

In both these readings the *when*-clause plays the part of quantification frame and not of restrictor and provider of the dref that is bound by the quantifier. That is a difference with (3.298.a), on both of the readings for it that we have considered. But otherwise there is a fairly close similarity between the second of the two readings for (3.298.a) and the second of the two just mentioned readings for (3.300.b.3). In both cases the nuclear scope contains just the information contributed by a PP modifier of the main clause VP. Our formal discussion of this second racing for (3.300.b.3) can rely for much that needs to be said about what the treatment above of the second reading of (3.298.a) has already made explicit, which means that we can be fairly brief. We start with this reading.

As we assumed for (3.298.a), the two readings for (3.300.b.3) have taken to have a single LF, which can then be transformed into two different LF' representations for the two readings. (3.316) gives the shared LF.

(3.316)

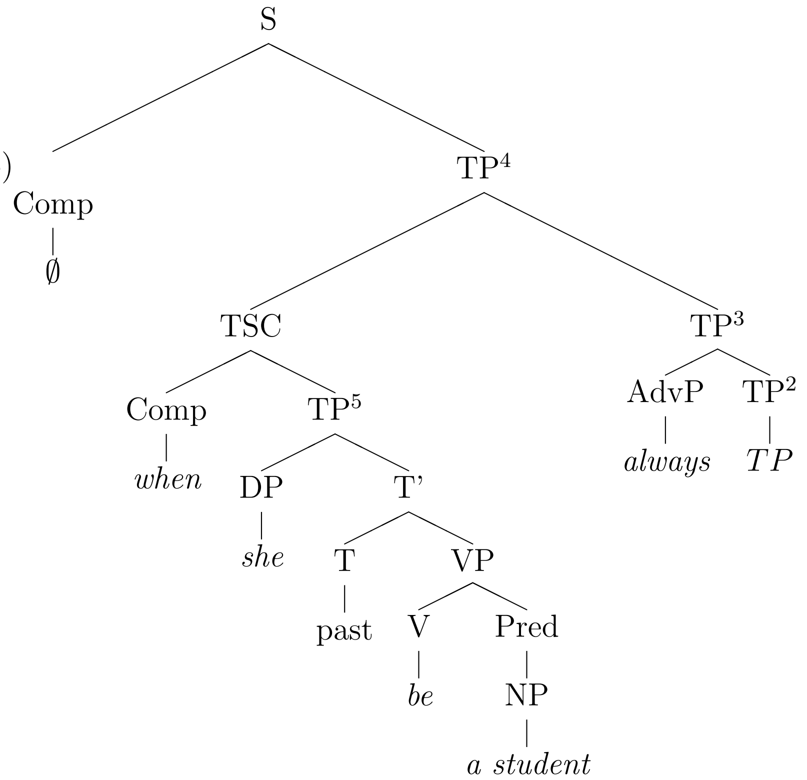

(3.317)TP =

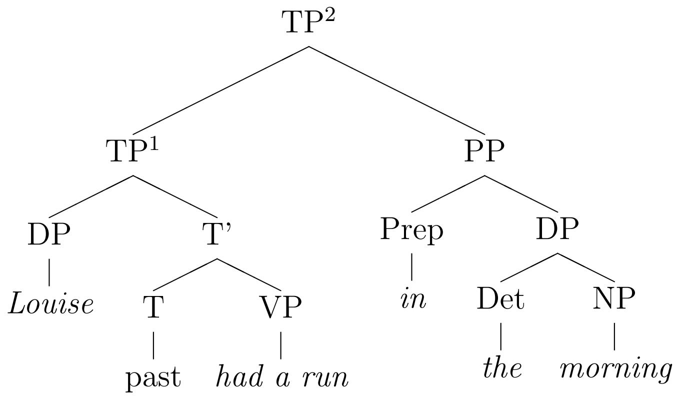

Note that in this LF the TSC is attached above *always*. This reflects its playing the part of quantification frame in (3.300.b.3). I will return to this difference between (3.316) and (3.303) later on. (3.318) shows the LF' for the second reading of (3.300.b.3), in the same format in which we presented the LF's for(3.298.a) .

(3.318)a.

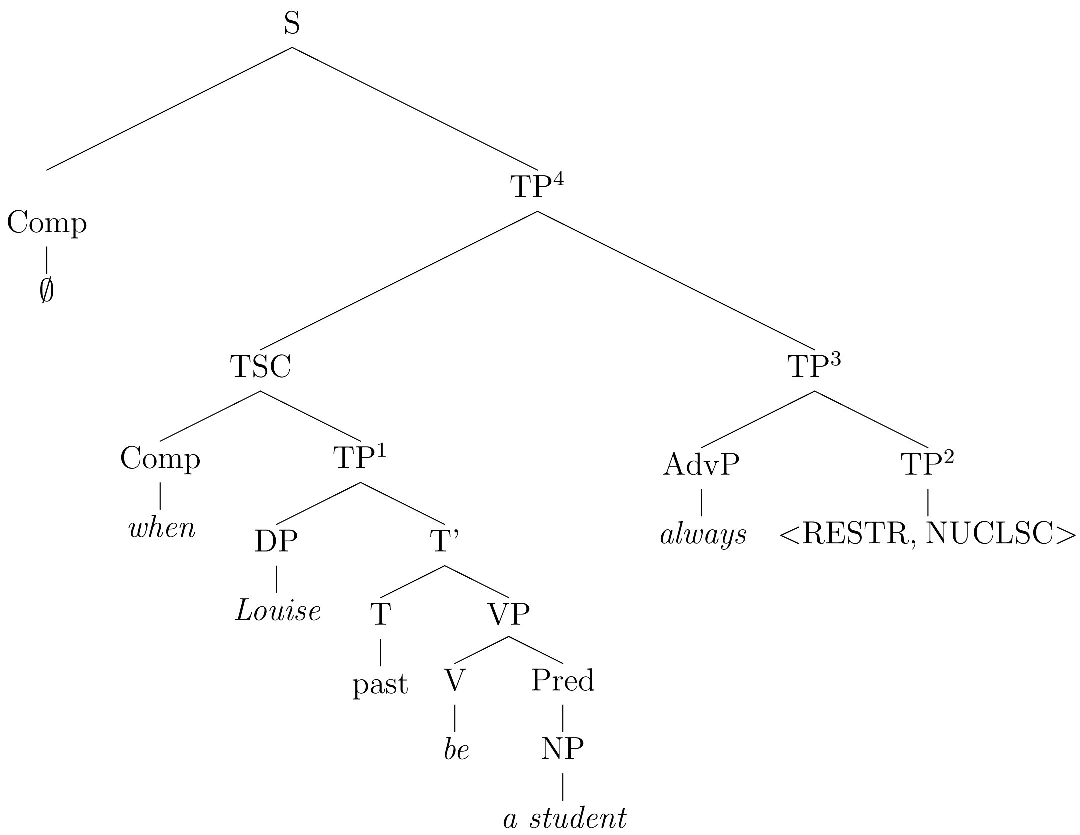

b.

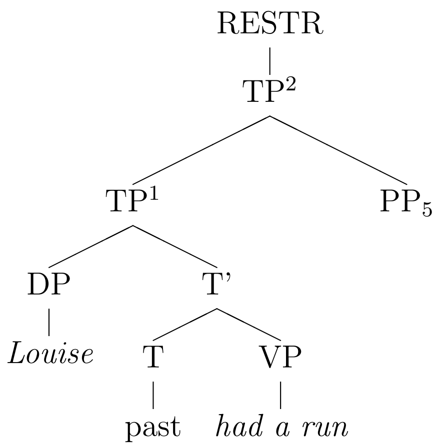

c.

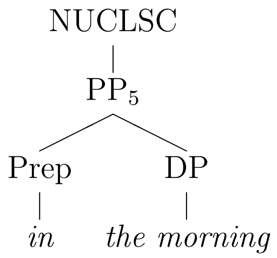

The construction of the semantic representation from the LF' in (3.318) closely follows that of the second reading for (3.298.a). We make two simplifications, treating the VP *have a run* as an unanalyzed verbal predicate and the PP *in the morning* also as an unanalyzed eventuality predicate. We will come back to the analysis of *in the morning* below, but will not look any further into the structure of *have a run*.

The derivation of the semantic representation from (3.318) is like that for the second reading of (3.298.a) in that unification between the PP in the NUCLSC part and the VP in the RESTR part occur before the RSTR and the NUCLSC part play their roles as input to the representation of the quantification. The representation of the quantification once again raises the question what should serve as dref bound by the quantifier *always*. We already noted that this is a difficult issue, for which the treatments in this section do not offer any real solution. A general principle is that this dref should be available in the restrictor part. In the cases of representation construction for sentences with adverbial quantifiers we have so far looked at the quantificationally bound dref was introduced by some part of the sentence material in the RESTR part, and that is also the case for the present construction (although this isn't always so, as we will see presently). In the case before us the dref that should be bound by the quantifier is the event dref e' introduced by the ‘verb’ *have a run*. (But exactly how this dref is selected is, I repeat a story that these NOtes fail to tell.)

Our next diagrams show the representation constructions for the node TP^3 and for the TP^1 node of the TSC. Note that this time the TSC is a stative clause, so its referential argument is a state dref.

(3.319)a. $$< s_{ref}, z \mid \begin{array}{|l|}\hline \\ s: \begin{array}{|c|}\hline t' \ \ e' \\ \hline t' \prec n \quad e' \subseteq t' \\ \text{“she”}(z) \\ e': \text{have-a-run}(z) \\ e' \subseteq s \\ \hline \end{array} \left\langle \begin{array}{c} \forall \\ e' \end{array} \right\rangle \begin{array}{|c|}\hline \\ \hline \text{in-the-morning'}(e') \\ \hline \end{array} \\ \\ \hline \end{array} >$$

b. $$< s'_{ref}, x \mid \begin{array}{|c|}\hline t \\ \hline t \prec n \ \ t \subseteq s' \\ \text{Louise}(x) \\ s': \text{student'}(x) \\ \hline \end{array} >$$

The next step is to combine the representation in (3.319.a) with the representation of the *when*-clause, which is providing the quantification frame for the quantifier *always*. We already touched on the question how quantification frame phrases or clauses combine with the quantifications for which they are the frames when looking at the representation constructions for (3.298.a), in particular in footnote 72. In that discussion the question was raised what the temporal relation should be between the time that is made available by the quantification frame phrase and the quantification state that is introduced by the quantifier. In particular, I there raised the question whether the quantification state s should be represented as including the time t^2 denoted by the PP *in the nineties* or whether the relationship should be tighter, viz. that t^2 is the duration of s. In the discussion of this point in footnote 72 I touched on the principle that temporal quantifications, adverbial temporal quantifications among them, always need a temporal quantification frame, as an essential aspect of their semantics.[73])

[73]The need for a temporal quantification frame arises most clearly when tense places

Let us assume that the interpretation of temporal quantifications requires the identification of a temporal quantification frame (with quantifications temporally surrounding the TPpt a likely exception). And that when a part of a sentence containing such a quantification is identified as the specifier of its temporal quantification frame, then its contribution takes the form of a temporal identity condition (such as 'dur(s) = t^2'). Then how are we to formally state the principle that this is so and that consequently such an identity condition must be added when the representation of the temporal quantification frame phrase or clause is combined with the representation of the quantification? A minimal requirement should be that the syntactic structure from which the representation is constructed – the LF or LF' – must make it possible to recognize that the phrase or clause plays this particular role. The LFs for quantifying sentences on which we have been relying so far there is no unambiguous indicators to this effect. (We have done no more than prepare the ground for the intended constructions by adjoining the temporal quantification frame phrase or clause in a position that facilitates the representation construction, for instance immediately above the quantifying adverb in our LFs for the present sentence and the previous one. But according to our general DRS construction conventions there is nothing in the adjunction position as such that unequivocally identifies the expression as a temporal quantification frame clause or phrase.)

We could consider at this point adopting a new syntactic construction to deal with temporal quantifications and their temporal quantification frames. For instance, one might contemplate syntactic structures for quantificational adverbs and the expressions contributing their quantification frames along the lines of (3.320).

---

the quantification state in the past pr the future of n. For present tense sentences with adverbial or other kinds of quantifiers the condition that the quantification state includes n is arguably all that is needed. (This is what we have been assuming in earlier representation constructions for sentences involving quantifications, for instance in Section 3.9.1.) Perhaps a case could be made for the position that even in present tense quantifying sentences there must be some implicit quantification frame – that the speaker must have some interval surrounding n in mind when she uses such a sentence; at least she should have some conception of the order of magnitude of this interval. Personally I am not persuaded that this is right, but I leave it as an open question. If the present tense cases are really different from the non-present ones in that they commit to no more than that the quantification state includes n, then these cases have to be handled separately. (In fact, the right generalization here would be that no more than a temporal inclusion specification is needed whenever the quantification state is understood as including the TPpt. For details see Section 4.

(3.320)

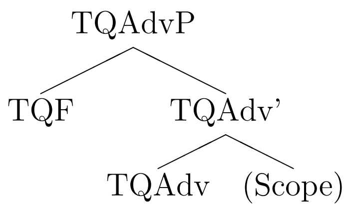

A syntactic configuration like this one makes explicit (a) that the TQAdv constituent (the temporal quantificational adverb) combines with its input in the right way, which lead to a representation with a referential argument for a correctly characterized quantification state; and (b) that representation of the structure labeled TQF (for 'Temporal Quantification Frame') specifies the duration of that state as the temporal interval that it itself provides.

I am inclined to think that some such ploy will have to be adopted eventually. But I shy away from such a non-trivial revision of our syntax-semantics interface at this point. So we will stick with the syntactic structures we have been using throughout Section 3, and in particular with the LF and LF' for the sentence with which we are dealing.

De facto this means that we have to assume, when combining the TCS with the quantification representation in (3.319.a), some special construction principles. What we are really facing at this point is a construction step in which the *when*-clause is identified as temporal quantification frame constituent, whose function it is to locate the quantification state represented by its sister node. We can think of this total operation as involving two separate operations. The first turns the semantic representation of the TP node that is sister to the Comp node occupied by*when* into a representation whose referential argument is a time dref t^2 that stands for the duration of the referential argument of the TP representation. The next operation then identifies t^2 with the duration of the quantification state from the representation of the sister node to the TSC. (3.321.a) gives the representation of the *when*-clause constructed according to this recipe and (3.321.b) the result of combining this representation with the representation in (3.319.a). We have also transferred the drefs that still occupied store positions at this point to the right DRS Universe, so that (3.321.b) is the final representation.

$$(3.321)\text{a.} \qquad < t^2_{ref}, x \mid \begin{array}{|c|}\hline t \quad s' \\ \hline t \prec n \quad t \subseteq s' \\ \text{Louise}(x) \\ s'\text{: student'}(x) \\ t^2 = \text{dur}(s') \\ \hline \end{array} >$$

$$\begin{array}{|c|}\hline t \quad s' \quad t^2 \quad s \quad x \quad z \\ \hline \\ t \prec n \quad t \subseteq s' \\ \text{Louise}(x) \\ s'\text{: student'}(x) \\ t^2 = \text{dur}(s') \\ \\ s\text{:} \begin{array}{|c|}\hline t' \quad e' \\ \hline t' \prec n \quad e' \subseteq t' \\ z = x \\ e'\text{: have-a-run}(z) \\ e' \subseteq s \\ \hline \end{array} \; \left\langle \begin{array}{c} \forall \\ e' \end{array} \right\rangle \; \begin{array}{|c|}\hline \\ \hline \text{in-the-morning'}(e') \\ \hline \end{array} \\ \\ \text{dur}(s) = t^2 \\ \hline \end{array}$$

So far we ignored whatever internal structure there is to the VP *have a run* and the PP *in the morning*. As I said earlier, we'll leave things with the VP as is, since delving into its internal structure won't serve any real purpose here. But it is going to be useful to say a little more about the PP.

The first point to be made is that *morning* is most naturally treated as a relational noun. When someone speaks of 'the morning', the natural question that provokes is: 'Morning of which day?'. Indeed, this is how in the majority of cases the reference of an occurrence of *the morning* or *this morning* or *that morning* is resolved: one tries to identify some particular day such that the referent is that day's morning. The story about the DP *the morning* of the PP *in the morning* in (3.300.b.3) is much like this. But there is a further twist to it: the day in question is the day at which the relevant value for the event dref e' occurred. We cannot give all the formal details that are involved in this representation. They involve a formal procedure which identifies with the dref e' a dref d for the unique day on which e' occurred and then resolving the referent m of *the morning* as the unique entity that stands in the 'morning' relation to d. The methods developed in Section 4 will enable us to make this precise. Here we only show the result, which replaces the DRS in (3.321.b).

(3.322)

$$
\begin{array}{|c|}
\hline
t\ \ s'\ \ t^2\ \ s\ \ x\ \ z \\
\hline
t \prec n\ \ t \subseteq s' \\
\text{Louise}(x) \\
s'\text{: student'}(x) \\
t^2 = \text{dur}(s') \\
s\text{:}\ \begin{array}{|c|}
\hline
t'\ \ e' \\
\hline
t' \prec n \quad e' \subseteq t' \\
z = x \\
e'\text{: have-a-run}(z) \\
e' \subseteq s \\
\hline
\end{array}
\ \left\langle \begin{array}{c} \forall \\ e' \end{array} \right\rangle\
\begin{array}{|c|}
\hline
d\ \ m \\
\hline
\text{day'}(d) \quad e' \subseteq d \\
\text{morning'}(m, d) \\
\hline
\end{array} \\
\text{dur}(s) = t^2 \\
\hline
\end{array}
$$

What remains is the other reading of (3.300.b.3), in which all the material in the sister TP TP^2 to the *when*-clause belongs to the nuclear scope of the quantifier and there is thus no overt material to go into the restrictor. That is the LF' structure for this reading of (3.300.b.3) is as in (3.323).

(3.323)a.

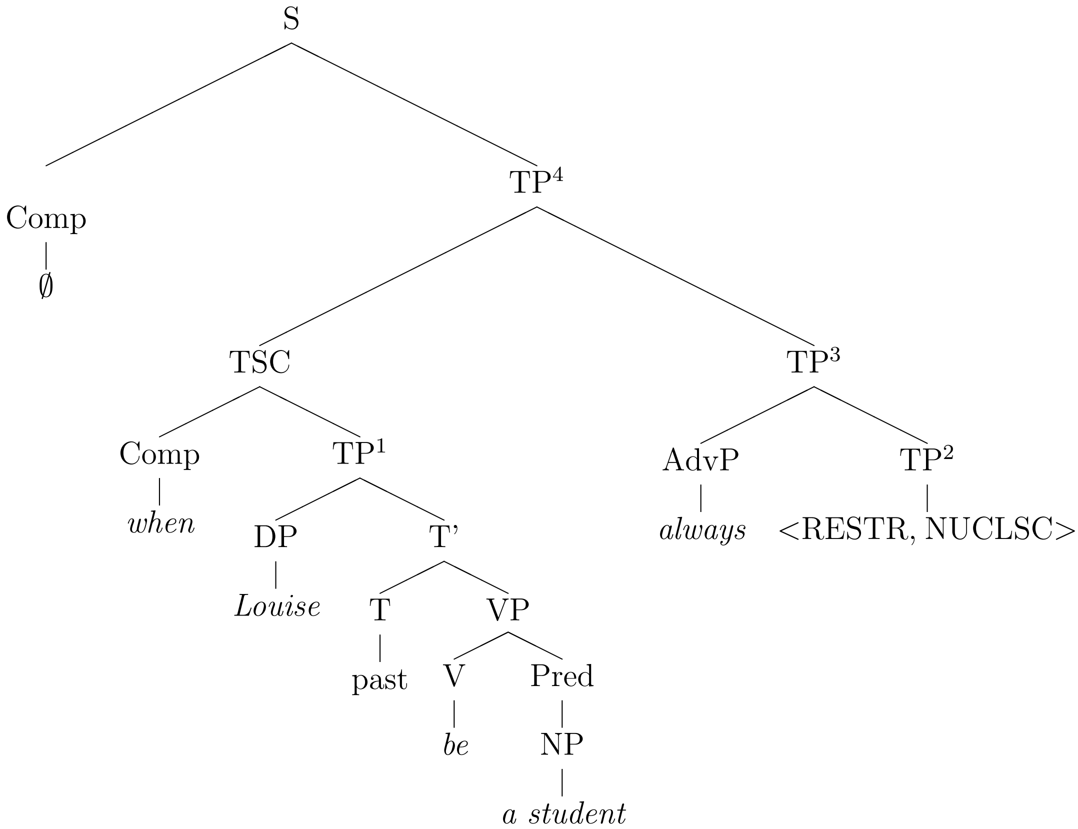

b.

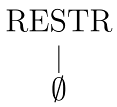

c.

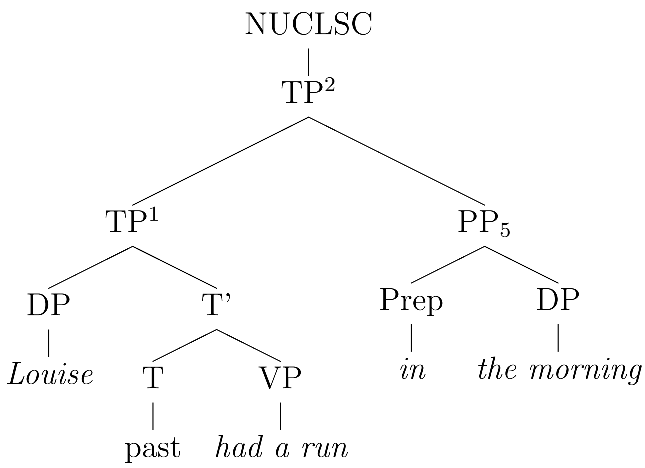

The problem with this structure is a general prohibition against empty restrictors, for one thing because it is the restrictor of an adverbial quantification that is responsible for supplying the quantificationally bound dref of the duplex condition and for another because the restrictor has to provide some kinds of constraints on the quantificationally bound dref.

In the present case the need to accommodate material for the restrictor that fills these gaps goes hand in hand with the need to resolve the definite description *the morning* in the nuclear scope part. As we have just seen, this definite description is naturally resolved in a 'dependent' way, viz. by interpreting the morning as dependent on the dref that is quantificationally bound. And the third factor that conspires with these two is that, as we observed earlier, temporal adverbial quantifications often involve a granularity determining partition of the quantification frame. Since the different mornings that are the possible denotations of *the morning* will have to be mornings of different days, the same kind of dependent resolution of *the morning* that we argued for in relation to the other reading of (3.300.b.3) should be what the present interpretation should be aiming for too.

We can solve solve all three problems – the need for a quantificationally bound dref, with suitable constraints on its possible values, determining the granularity of the quantification and a suitable resolution of *the morning* – by accommodating as quantificationally bound dref a dref d that ranges over the calendar days within the quantification frame and interpreting *the morning* as the morning of d. This accommodation leads to the representation in (3.324) for the TP^3 node of (3.323.a).

(3.324) $$\begin{array}{|c|}\hline t\ \ s'\ \ t^2\ \ s\ \ x\ \ z \\ \hline \\ t \prec n\ \ t \subseteq s' \\ \text{Louise}(x) \\ s'\text{: student'}(x) \\ t^2 = \text{dur}(s') \\ \\ \hline \end{array}$$

$$< s_{ref} \mid \begin{array}{|l|}\hline \\ \hline s : \begin{array}{|c|}\hline d \\ \hline \text{day'}(d)\ d \subseteq s \\ \hline \end{array} \ \genfrac{}{}{0pt}{}{\forall}{d}\ \begin{array}{|c|}\hline t'\ \ e'\ \ m \\ \hline t' \prec n \quad e' \subseteq t' \quad e' \subseteq d \\ z = x \\ e'\text{: have-a-run}(z) \\ \text{morning'}(m, d) \\ \hline \end{array} \\ \hline \end{array} >$$

The remainder of the construction is identical to that for the other reading of (3.300.b.3) and left as a (boring) exercise.

Of all the DRS constructions presented in Section 3 those from this last subsection (Sn 3.11.7) have been the arguably the most complex. They have also been the least satisfactory in that, to a hitherto unseen degree, man of the construction principles we have appealed to have been ad hoc: no clear formally precise generalizations can be inferred from the applications in the constructions shown. There were several reasons for this. One is the difficulty to determine the function of certain sentence constituents. An example of this are the different roles that are played by the *when*-clauses of (3.298.a) and (3.300.b.3) – quantification restrictor in the one case and temporal quantification frame in the other. We touched on the question how LF notation could be adapted so as to make this distinction explicit. But the real problem is how and when this information can be supposed to become available: intuitively the choice between restrictor and quantification frame often has to be made on the basis of information about the meaning of a given sentence, and in the architecture we are working with here that information will be available only after much of the interpretation process, which converts LF into semantic representation, has already taken place, and not before it starts. As we noted before, the best that an architecture like ours can do to cope with problems of this sort is to let the parser make guesses, let the construction algorithm do its work for each of those guesses and then check the results for coherence and plausibility. But such a model of the interpretation process is not only suspect because of its computational inefficiency. At the present time it also seems beyond direct implementation, because we do not seem to have operational definitions of coherence and plausibility.

The most important obstacle to a general set of construction set of principles that covers at least the examples looked at in the president subsection is the need to appeal, in our last DRS construction, for the second reading of (3.300.b.3). Not only did this accommodation involve the reference resolution of the definite description *the morning* – this is a matter that we will be able to deal with to a large extent in Section 4. But what seems a much less tractable problem is process of accommodating just the right dref and DRS conditions to comply with the ('non-emptiness') requirements for the restrictors of duplex conditions. Such accommodations often seem to rest on certain forms of abductive inference. But I don't thin anyone at the present time has a clear conception of what these abductive process are in general.

Because too many of the construction principles that we have relied upon in this subsection have not been articulated with sufficient precision and sufficient generality, the rules involved in the interpretation of adverbial quantifiers has not been included in the Summary of the rules for constructing semantic representation that is given in Section 3.11.9, the final subsection of Section 3.

3.10.8 Temporal Adverbs and Vagueness

In our informal discussion of quantifying temporal adverbs, we set aside the majority – among them *often*, *rarely* and *regularly* – because they present problems that we had no means to deal with. These problems come on top of those that we then did proceed to address, problems that have to do with quantification frame, granularity and the filling of restrictor and nuclear scope of the duplex condition induced by the quantifying adverb. The additional problems presented by *often*, *rarely*, *regularly* and their kin have to do with the *vagueness* of these quantifiers. Recall the package of sentences presented in (3.293). Past tense versions of those sentences with an explicit quantification frame, are given in (3.325).

(3.325)a. In the nineties Louise answered her letters often by e-mail.

b. In the nineties Louise answered her letters rarely by e-mail.

c. In the nineties Louise answered her letters regularly by e-mail.

The question we posed earlier in relation to such sentences was: How many instances of the event type described – that of Louise answering letters by email – must there be within the quantification frame in order for each of these sentences to be true? The answer to such questions has to be in the spirit of something like: 'That depends.' But what do such answers depend on?

The problems raised by these questions are instances of a much more general problem about meaning in natural language, that of *vagueness*. Vagueness makes itself felt in meaning and truth conditions in relation to almost everything that plays a part in the determination of the truth conditions of sentences and the concepts they and their constituents express. For instance, vagueness is manifest in the extensions of nearly all predicate words of natural languages – that is: of nearly all their nouns, verbs, adjectives and adverbs. It is particularly prominent with the last two types of words. This is clearly reflected by the literature on vagueness, in which there has traditionally been a heavy focus on adjectives, as the paradigm examples of vague predicates. (On the whole the vagueness literature hasn't had all that much to say about adverbs; but the reason for this is, I presume, that the bulk of adverbs function as predicates of events, and events are problematic in their own way because of the vagueness of their identity conditions. When we talk about predicates of people, for instance, such as *tall*, *clever* or *bored*, then this second source of vagueness, concerning the sameness or distinctness of the entities to which a vague predicate can be applied, can be safely set aside and we can focus more easily on the first.

Common to all who have had things to say about vagueness the vagueness of a predicate P manifests itself as indeterminacy of P's extension: for some entities the question whether P is true of them cannot be answered. A second aspect of the vagueness of natural language predicates is that the question what can be ascertained to belong to the extension of such a predicate depends on the context in which it is used. It continues to be a point of debate whether the context dependence of vague predicates is an integral part of their vagueness or an additional and separable factor. I myself am among those who see the context dependence of vague predicates and their vagueness as inseparable. It is in this spirit that this mini-section on vagueness should beread.

In the domain of temporal adverbials we find quite a few items that are vague in the sense of adjectival vagueness referred to in the preceding two paragraphs. Those that fit what we have just been saying about vague predicates most closely are non-quantifying adverbs like *soon* or *recently*, or PPs like *a long time ago*. How soon is *soon*? How recent must an event be to qualify as one that can be described with the help of the adverb *recently*? How long ago must an event or state of affairs be to be describable as 'a long time ago'? Once again, uniform answers are not be had. To the extent that answers can be given, they will depend heavily on what kinds of eventualities are being located as near or far in time; and even when we know what eventualities are being talked about, or in what terms they are being described in a sentence containing a vague temporal adverbial, the answer may still depend on yet other contextual factors. And even when all contextual information is in, the extensions of such adverbials will still not be fully determined.

The claim that no context will resolve all questions about the extension of a vague predicate P at once must be distinguished from the one whether for some particular candidate for membership in the extension P there will be contexts in which it is settled whether that candidate does or does not belong to the extension. (For instance, there may be no context which settles for every person whether that person belongs to the positive extension of *bald*, but it could nevertheless be the case that for every person d there are some contexts in which it is determined whether or not d is bald.) In fact, predications involving vague predicates have a stipulative dimension to them: By asserting '$P(d)$' a speaker can *create* a context in which the predication is true (because by asserting it she has brought it about that (in the context thus created) the claim is true). In this way it is possible to create contexts in which '$P(d)$' has a definite truth value for any d to which P can be applied. On the other hand speakers can finesse the problem of indeterminacy by simply not applying vague predicates to entities in contexts in which the truth value of $P(d)$ is not determined.

These two aspects of vagueness – our tendency to avoid real or potential instances of indetermination, as if we were tiptoeing around a patch of quicksand, and the possibility of stipulating predicate satisfaction (or non-satisfaction) at least for the purpose and duration of a given conversation – are among the reasons why vagueness is much less of an obstacle than one might come to think on the strength of awareness of how ubiquitous it is. It is also for this reason that semanticists have been on the whole on fairly safe ground in ignoring vagueness when exploring other issues in natural language meaning. This is what we have been doing also in these notes, and it is the policy to which we will stick also in what it still to come.

We started this interlude to vagueness with a reference to vague quantificational adverbs – those temporal adverbials that provoked a comment on their vagueness in Section 3.11.4. At first blush the vagueness displayed by such adverbs may not seem to fit the notion of predicational vagueness that we have been talking about in the present section all that well. But a closer look suggests that they too can be seen as involving vagueness of predication. The entities about which such adverbs can be construed as making predicational claims are their quantification frames. What for instance is it for a claim to the effect that within the duration of some quantification frame t it was often the case that P to be true? How many instances of P within t must there have been?

The vagueness of quantificational claims made by sentences involving adverbial quantifiers like *often* or *rarely* is also found with habitual sentences, discussed in Section 3.7.2: How often must something be done over a give period of grime t in order to count as having been done habitually during T? However, the claims made by habitual sentences are arguably not just distributional. Typically they aren't understood as just claims about the *actual* distribution of instances of the relevant eventuality description over the quantification frame, but as also carrying implications about what the distribution over the frame might or would have been if the world had been different from what it actually is, or was. This counterfactual dimension of the claims that habitual sentences make about their quantification frames is even more obvious for sentences whose representations involve the disposition operator DISP, but it also undeniably present in sentences containing the adverb *habitually*. Exactly which quantificational adverbs have such an intensional dimension to their semantics is not all that easy to settle – compare for instance *regularly*, *usually*, *habitually*, *often* and *rarely* from this point of view. But it seems plausible that *habitually* isn't the only intensional one on this short, very incomplete list. (At least *regularly* and *usually* would seem to be strong candidates as well.)

In addition, the duration of the temporal quantificational frame will often be vague as well. We noted that this is always so with present tense sentences involving temporal quantifying adverbs. But past and future tense sentences will often have vague or underspecified frames too, especially when they lack frame adverbials that specify their frames, so that the frame must be inferred from the context. Sentences with adverbial temporal quantifiers can thus be vague along several dimensions at once. Thismakes such sentences an interesting topic, although, as a matter of actual fact, I do not know that they have been studied from this angle, as part of the theory of vagueness.//

But to repeat, vagueness is not among the topics of these Notes, and we will not return to it again.

3.11 Summary of the Representations and Rules of Section 3

In the course of Section 3 our representation structures have become gradually more complex and so have the rules used to construct them. Representations and construction rules will become yet more complex when we incorporate the representation and construction of presuppositions into them in Section 4. But before these additional complications are added, it will be good to summarize the current status quo. We start with the simpler part, the general form of the representations. But we include among the notational devices that can enter into the representations of the formalism we define in this subsection not only the representational constructs that have been considered so far in Section 3 but also allow for the representation of pluralities, which played some part in Section 2, but not in Section 3.

3.11.1 Representations

1. The representations that are constructed by the construction algorithms of Section 3 consist of a *DRS*, preceded by a *store*. The store is a finite, possibly empty list of discourse referents, with or without annotations, and the DRS is a pair consisting of (i) a set of discourse referents (the *universe* of the DRS) and (ii) a set of *DRS conditions* (*its condition set*). Representations with empty stores are identified with the DRSs that are their second components. The annotations of discourse referents occurring in stores are given in the form of subscripts. They may be (i) canonical natural number terms ('0', '1', etc.) or (ii) the label 'ref'. Numerical subscripts are also referred to as *indices* (used for coindexation with other elements of the structures in which they occur: other drefs, slot markers (see below) and noun phrases).

2. Discourse referents come in different *sorts*: (a) *entities* in general, (b) *times*, (c) *eventualities* (subdivided into two sub-sorts, (d) *events* and (e) *states*).

The sorts of times and eventualities are mutually disjoint and are sub-sorts of the general sort of entities. Likewise the sub-sorts of events and states are mutually exclusive. (The inclusion and exclusion relations between the sorts are reflected by the structure of the Universes D_e – the universes of the type of individuals – of the models for our DRS language. These Universes are subdivided into sorts, those of times, eventualities, events and states among them.)

Special symbols are used for discourse referents that represent entities of the special sorts ‘time’, ‘eventuality’, ‘event’ and ‘state’. These symbols all involve a letter that ‘wears the represented sort on its sleeve’ – ‘t’ for times, ‘e’ for events, ‘s’ for states and ‘ev’ for eventualities. The letters that occur as parts of the symbols used as discourse referents for entities in general – i.e. discourse referents that are not restricted to the representation of entities of a special sort are chosen from the end part of the alphabet: ‘x’, ‘y’, ‘z’, ‘u’, ‘v’, ‘w’... . In general, a discourse referent symbol will consist of just a letter or of a letter followed by a numerical subscript, or followed by one of more primes or by a numerical superscript (to be thought of as a compact way of representing larger sequences of primes). Note well: the subscripts mentioned in this last sentence, which are integral parts of discourse referent symbols, must be sharply distinguished from the indices mentioned earlier. These indices are also graphically presented as subscripts. But unlike the subscripts just mentioned, indices are independent elements in their own right, which have their own roles to play in the course of representation construction: Some of the construction rules make explicit use of them. (For details see the section on rules below.)

There are also discourse referents for *pluralities* (= mereological sums) of entities of the various sorts for which there are individual discourse referents. The letters that are part of the symbols used for these plural discourse referents are upper case letters, but otherwise they are subject to the same restrictions as the symbols for discourse referents for individual entities of the different sorts. The sortal implications of the letters that are part of discourse referents for pluralities are the same as in the case of discourse referents for single entities. Thus ‘E’,‘E'’, ‘E^2, ... all represent pluralities of events, and so on.

Lower case discourse referents are also referred to as *singular* or as *individual* discourse referents, upper case discourse referents as *plural* discourse referents.

3. DRS conditions come in a number of different forms:

(a) atomic conditions, which consist of an n-place *predicate constant* combined with n *terms*(each of which must conform to the sortal restrictions imposed by the predicate on the argument position filled by term). With a few exceptions the argument terms of atomic DRS conditions are of one of two kinds: either (i) discourse referents or (ii) *slot markers*. Slot markers serve as indicators that the argument slot they occupy is still waiting for a 'real' argument. The symbols used for slot markers are underlined and consist of a letter with or without a numerical subscript.

The forms of atomic DRS conditions follow the somewhat idiosyncratic conventions that have become widely used within DRT over the years: for predicates corresponding to most nouns and adjectives the notation is one of the standard notation used in Predicate Logic, that of an n-place predicate followed by a sequence of n argument terms that is enclosed within parentheses and in which the terms are separated by commas. (Examples are 'philosopher'(x)', 'mother'(x, y)', 'happy'(x)', unknown-to'(x, y)','above(x, y), expressing that x is a philosopher, x is mother of y, x is happy, x is unknown to y, x is above y.) But for predicates whose referential argument is an eventuality – those that are used to represent verbs and most prepositions – the referential argument term is placed in front of the predicate separated from it by a colon, while the non-referential argument terms are placed behind the predicates enclosed in parentheses ('e: sleep'(x)', 's: hate'(x, y)', 's: PROG($^\wedge e.$ $\begin{array}{|c|}\hline e \\ \hline e\text{: sleep'}(x) \\ \hline\end{array}$)', etc.). Finally, there are some special 2-place predicates which combine with their arguments in agreement with the 'infill' principle, according to which the predicate is placed between its two argument terms, as in '$x = y$', where '=' occurs between1x' and 'y'. Other predicates for which the infill notation is adopted are '$\prec$', '$\succ$', '$\subseteq$', '$\supset\subset$'. (This ends the most fussy and tedious part of this summary.)

Among the predicate constants of our language there are those that select for times and/or eventualities. This is true for the 2-place predicates '$\prec$'', '$\subseteq$', '$\supset\subset$ each of whose arguments can be either a term of sort 'time' or a term of sort 'eventuality', and also for 2-place predicate 'res', whose first argument must be of the sort 'state' and the second argument of the sort 'event'. Furthermore, the predicates representing event verbs will require of their first arguments that they be of the sort 'event' and the predicates representing state verbs will require of their first arguments that they be of the sort 'state'. (But note well, this is not an exhaustive account of selection restrictions. Predicate words of natural languages typically impose various restrictions on the sorts of their arguments. When formal counterparts of these words are introduced into our representation language, so as to make it possible to represent the semantic contributions that are made by these words, then it will always have to be part of these introductions what the selection restrictions on the arguments of these predicates are.)

Besides discourse referents and slot markers, atomic predications may sometimes also contain *complex terms* as argument terms:

(i) When α and β are singular or plural discourse referents, then $\alpha \oplus \beta$ is a *term* (denoting the mereological sum of the denotations of α and β).

(ii) When A is a plural discourse referent, then $|A|$ is a term (denoting the cardinality of the plurality denoted by A).

(iii) Constant terms for the natural numbers. The use of such terms has come up more than once, but we haven't been very explicit about what they like. At this point we should make a decision. Let us adopt the usual decimal notations for this purpose. That is, we use '0' as the constant term that denotes the number zero, '1' as the term denoting the number one, '2' as the term denoting the number two, .., '15' as the term denoting the number 15 and so on. (As regards the model theory for our formalism this entails that every model M contains a copy of the natural numbers and that the number terms '0' etc denote in M the corresponding elements from M's copy of the natural number set.)

Terms of types (ii) and (iii) only occur in identity conditions, of the form '$|A| = \mathbf{n}$', where A is a plural discourse referent and $\mathbf{n}$ a number term.

(iv) If α is a discourse referent of the sort 'eventuality' (or of one of its subsorts), then 'dur(α)' is a term of sort 'time'; 'dur(α)' denotes the period of time that is occupied by the denotation of α.

3.b Besides atomic DRS conditions the Condition Sets of DRSs may also contain complex conditions. Complex conditions are defined recursively, with the set of atomic DRS conditions as basis for the recursion and recursive clauses that take the form of applications of certain syntactic operations. Complex conditions have always been part of DRT, from its first formulation in (Kamp 1981*b*). But in earlier DRS formalisms the constituents of complex conditions were DRSs. In the representation formalism we have been developing here the constituents are in general representations in which DRSs are preceded by a store.[74]

[74]Our representations may have several stores, not only a *main store* – one that precedes all of the remainder of the representation – but also stores in subordinate positions. This

Apart from having to provide for stores, however, the formation rules are the same as in the earlier formulations of DRT, at least for the operators that are considered in these earlier formulations: $\neg$, $\Rightarrow$, $\vee$, $\boxed{}\langle\forall\rangle\boxed{}$. (That arbitrary DRSs can be the inputs to the operations that create complex DRS conditions, whereas these conditions can in their turn occur as members of DRSs, means that the definition of complex conditions and DRSs must be cast in the form of a *definition by simultaneous recursion*; for details see (Kamp & Reyle 1993).) The formally simplest of these complex DRS conditions are those formed with the help of the *Negation operator* $\neg$, which combines with a single input representation. The other operators are binary: they take two representations as inputs. We have already seen a number of examples of how complex DRS conditions with two input representations are formed. A special case is constituted by the duplex condition, which are used in DRT to represent quantification and in which the quantification operator binds a discourse referent that may have occurrences in both of the input representations. (This is the norm for duplex conditions that arise in the construction of NL sentences with quantifiers). In the early applications of DRT the most prominent duplex conditions are those representing universal quantification. But duplex conditions can in principle be formed for any binary quantifier Q (with the requirement of course that explicit verification conditions are stated for these duplex conditions that correctly capture the semantics of Q). In the Notes we have so far encountered one type of non-universal duplex condition, that in which Q is the existential quantifier. But other quantifiers are expressible in natural languages as well, for instance those for which English has the words *many* and *most.* Duplex conditions that capture the semantic of such quantifiers could in principle be added to our representation formalism as well. And they sure will have to be eventually.

There is no need to state the formation rules for all the operators listed in this paragraph. By way of example we just state the formation rule for negation conditions and that for duplex conditions whose operator is a binary quantifier Q.

---

possibility arises because the governing operators of complex DRS conditions can take scope over the stores of their operanda. In Sections 3.9 and 3.10 we noted some problems with DRSs that contain stores in such subordinate positions. The assurance was given that those difficulties would be resolved in Section 4. At this point the only thing to do is to define the formation rules for complex DRS conditions so that they cover the examples of representations that have been given in Sections 3.9 and 3.10.

3.b.1 (i) Suppose that $<ST \mid \; K>$ is a representation with store ST and DRS K.

$$\text{Then } <ST', \mid \begin{array}{|c|}\hline \\ \hline \\ \neg <ST,K> \\ \hline\end{array} > \text{ is a } \textit{complex DRS condition.}$$

(ii) Suppose that $<ST_1 \mid \; K_1>$ and $<ST_2 \mid \; K_2>$ are representations with stores ST_i and DRSs K_i and that Q is a generalized quantifier and that ev is an eventuality discourse referent, that ST'' is some store and that $ST''' = ST \setminus \{x\}$. Then

$$<ev_{ref}, ST'' \mid ev: \begin{array}{|c|}\hline \\ \hline \\ < ST''' \mid \begin{array}{|c|}\hline \\ K \cup \begin{array}{|c|}\hline x \\ \hline \\ \\ \hline\end{array} \\ \\ \hline\end{array} > \\ \\ \hline\end{array} \left\langle\begin{array}{c} Q \\ x \end{array}\right\rangle \begin{array}{|c|}\hline \\ \hline < ST' \mid K'> \\ \hline\end{array} >$$

is a *complex DRS condition.*

3.b.2 Next the complex conditions that take eventuality descriptions as inputs. We have encountered four operators that operate on such inputs, $PROG$, HAB, $DISP$ and *Res*. Strictly speaking, these 'operators' are higher order predicates (syntactically as well as semantically). So we could have treated the conditions to which they give rise as predications with higher order argument terms, and thus as atomic conditions. That would have required a different definition of terms, however – one according to which argument terms can also be formed through application of the intensional abstraction operator ^ to an eventuality description of the form $< ev_{ref}, .. \mid K >$.

The alternative way of dealing with such DRS conditions is to treat $PROG$ and consorts as operators that require eventuality descriptions as inputs, and that yield as outputs conditions which take the form of relating their referential arguments – we have seen that for all of $PROG$, HAB, $DISP$ and *Res* the referential argument is a state – to intensional abstractions obtained via application of ^. This is the perspective that informed the syntax and semantics we have been assuming for these operators. There is no reason why we should not hold on to it here.

If we take this line, then conditions formed with the help of $PROG$, HAB and $DISP$ can be defined as follows:

Suppose that $< ev_{ref}, ST \mid \ K>$ is a representation with store ev_{ref}, ST and DRS K, and that s is a state dref. Then

$< s_{ref}, ST \mid s : OP(^\wedge ev.K)$ is a *complex DRS condition*, where OP is one of $PROG$, HAB and $DISP$.

Complex conditions formed with the help of *Res* can be defined in much the same way:

Suppose that $< ev_{ref}, ST \mid \ K>$ is a representation with store ev_{ref}, ST and DRS K, that e is an event dref and that s is a state dref. Then

$< s_{ref}, ST \mid s : Res(e,^\wedge ev.K)$ is a *complex DRS condition.*

3.b.3 The last type of complex condition that is part of our representation formalism is that formed with the help of the disjunctive underspecification operator $\overset{!}{\vee}$. We encountered only instances in which $\overset{!}{\vee}$ forms disjunctions between DRSs. But there is in principle no reason why underspecification disjunctions could not also arise whose disjuncts are representations with non-empty stores. It is this more general characterization we adopt here:

Suppose that $<ST_1 \mid \ K_1>$ and $<ST_2 \mid \ K_2>$ are representations with stores ST_i and DRSs K_i. Then $<ST_1 \mid \ K_1> \overset{!}{\vee} <ST_2 \mid \ K_2>$ is a *complex DRS condition.*

The above covers the full range of possible representations of our formalism. As noted, some of the complexity that such representations can have serves to support the proper construction of representations from linguistic input and should no longer be present when a representation construction has been completed. More precisely, *completed representations* are distinguished by the following restrictions:

(i) they contain no non-empty stores: (ii) they contain no slot markers; (iii) they contain no elements of underspecification (i.e. no occurrences of $\overset{!}{\vee}$).

4. While representation construction is in progress, representations of the forms defined in 1.-3. above occur as decorations of nodes of syntactic trees. So what would be strictly speaking needed at this point is a formal definition of decorated syntactic trees, as functions from tree-structured node sets to the decorations of the nodes from those sets. Such a definition would have to combine (a) the formal definition sketched out above of the semantic representation formalism we have been using with (b) some formal definition of the repertoire of possible syntactic trees. In other words, this would require a formal specification of the generative grammar we are taking for granted in these Notes – of that grammar which is implemented by the parser that we assume delivers the LFs from which our semantic representations are computed. But spelling out such a grammar (and a parser that implements it) is precisely what we are trying to avoid in these Notes. So a formal definition of the decorated trees that can arise in the course of semantic representation construction is out of the question here.[75]

3.11.2 Rules

The construction rules we have been using in our sample DRS constructions can be divided into two groups. The first group consists of rules that are triggered by general syntactic configurations (and not by feature values associated with functional heads). It consists of three rules:

1. Lexical Insertion

2. Argument Insertion

3. Adjunction

Argument Insertion comes into two forms:

a. Argument Insertion involving non-quantificational argument phrases

[75]The best one could hope to do without getting embroiled in a formal definition of a generative syntax for English (or for a natural fragment of English that covers all examples we have been looking at in these Notes) would be to adopt is simple definition of a very liberal notion of 'syntactic tree', which will cover all that a a proper grammar for such a fragment could be expected to generate, but at the price of wildly overgenerating, and to then define a repertoire of 'possible' trees partially decorated with semantic representations of the kind we have defined. Given our definition of semantic representations such a project would no doubt be feasible. But it would be hard to see much interest to it.

b. Argument Insertion involving quantificational argument phrases

The second group is made up of rules that are triggered by feature values. It is at this point an open question how large this group will have to be eventually. We will briefly address the question whether the rules in this group can be reduced to a small number of rule types at the end of this section. The feature driven rules that we have encountered so far are:

4. Tense Feature rules

5. Aspect Feature rules

6. Frequency Adverb rules

7. Store Resolution

8. $\overset{!}{\vee}$ Resolution

In addition to these there are rules that have to do with the distinction between singular and plural. Strictly speaking these do not belong to the present list, since plural noun phrases have not been included in the bottom-up construction method that we have been pursuing in Section 3; it only figured in Section 2, where DRSs were constructed top-down. But such rules will of course be needed when the results from Section 2 about the semantics of plural and singular are incorporated into the bottom-up construction method, which they will have to eventually.

We now discuss the different rules in more detail. I start with the rules in Group I.

CR1. Lexical Insertion.

Applications of the rule of Lexical Insertion are replacements of lexical items from the given fragment by their semantic representations. So nearly all that can be said about the semantic import of its applications is contained in the lexical entries by which the lexical items are replaced. In the fragment we have so far considered, all lexical items function like predicates in the sense that their lexical entries are representations involving a store containing a dref marked as referential argument (with the subscript $_{ref}$) followed by DRS in which this dref figures as an argument in some DRS condition.

Lexical Insertion literally consists in replacing an occurrence of the given lexical item that is attached to a given syntactic node with the semantic representation from its lexical entry (as new attachment to that node), with one caveat: the drefs occurring in the semantic representation of the entry may have to be replaced by others of the same sort to avoid 'bound variable clashes' in cases where those drefs already have occurrences elsewhere in the representation that is being constructed.

CR2. Argument Insertion. Applications of this rule are triggered by syntactic configurations consisting of a mother node and two daughter nodes, in which one of the daughters is an argument phrase that is linked to an argument position somewhere in the representation associated with the other daughter. (In the fragment we have looked at, the argument phrase is always a DP.) Given the way we have implemented our syntax and semantics, such configurations can always be recognized because the category label of the one daughter is coindexed with an argument slot in the sister representation. One of the operations involved is always the insertion of the referential argument of the representation associated with the first daughter into the coindexed slot in the representation associated with the second daughter.

But this is not the only operation that is involved in applications of Argument Insertion. What else may be involved will depend on whether the DP is quantificational or non-quantificational. When the DP is non-quantificational, then the remaining operations take the form of merging the two representations, which consists in merging their stores and merging their DRSs. Furthermore, the referential argument of the DP representation is stripped of its referential argument status.

When the DP is quantificational, then a new DRS is introduced with an empty Universe and a Condition Set with just one condition. This condition has the form 's : DC, where DC is a duplex condition in which: (i) the central quantifier is the one denoted by the Determiner of the DP; (ii) the dref bound by the central quantifier is the referential argument of the DP representation, which is removed from its store and placed below the central quantifier in the duplex; (iii) the remainder of the store of the DP representation and the DRS following it are placed in the restrictor slot of the duplex condition and the store and DRS of the complement representation into the nuclear scope slot.

CR3. The Adjunction rule is triggered by all and only syntactic adjunction configurations (those in which an adjunct is syntactically combined with its adjunction site). The underlying concept of this rule is Unification: the referential argument of the adjunct is identified ('unified') with the referential argument of its adjunction site, whereupon the two representations are merged. In the implementation we have chosen the unification of the two drefs leads to replacement of the referential argument of the adjunct by that of its adjunction site: the referential argument of the adjunct is removed from its store and all its other occurrences are replaced by the referential argument of the adjunction site.

It is important to be aware of the conceptual difference between the rule of Adjunction and the rule of Argument Insertion. Argument Insertion has close affinities with functional application – the complement representation is 'applied' to the DP representation; part of this conception is that the two representations are of different logical types (that of predicate or functor and that of argument). Adjunction on the other hand treats the two representations that serve as inputs to its applications as of the same type – both are treated as 1-place predicates.

[N.B. make sure that Adjunction can be applied in the general form in which it is stated here also for temporal adjuncts: the referential argument of the adjunction site should be the location time, not the eventuality.]

CR4. So much for the rules of the first group. The second group is more open-ended and what will be said about it here is certainly not going to be the last word. Furthermore, specifications of the rules in this group vary with the feature values that trigger them.

CR4.1 Given the way in which construction rules triggered by feature values have been specified in our implementation, they are much like the rule of Lexical Insertion in that more or less all of the information needed to apply the rules is encoded in the 'lexical entries' for the feature values that trigger them. In fact, neither in the case of Lexical Insertion nor in that of the rules triggered by such feature values is there a need to say more than what we are told by the entries themselves. The 'lexical entries' for the triggering feature values in question are stated in the form of representation transformations with schematic references to the input representations to which the rules they trigger are to be applied. The only matter that needs to be spelled out is what it means to instantiate the schematic references in such an entry by actual semantic representations. But this is straightforward and more words would only tend to obscure that matter.

CR4.2 What has been said about feature value triggered rules in the last paragraph straightforwardly applies not only to the feature values determined by the T- and Asp-nodes, but also to some cases in which Comp acts on its complement representation, viz. those in which Comp determines an overt complementizer, such as *if*. (In this case the application of the rule produces a conditional DRS condition with the representation of the *if*-clause (the representation of the complement to the Comp node that dominates *if* as complementizer) as antecedent and the main clause representation as consequent. But note that in what has been presented here Comp nodes are treated quite differently depending of the type of clauses they head. Relative clauses have been treated as adjuncts as have temporal subordinate clauses headed by conjunctions like *after* or *before*. But there are many more conjunctions that form subordinate clauses besides the few we have discussed, and it isn't clear from the little that has been said about the ones for which an explicit treatment has been proposed, what treatment or treatments should be adopted for the others. All in all the semantics of adjunction is still a kind of building site of the theory we are in the prices of developing.

CR4.3 Another rule triggered by Comp is the rule of 'of Store Resolution' that is triggered by the empty Comp nodes of main clauses. As we have seen from the applications of the rule, this rule can be seen as a principle in its own right, which is unlike any of the other construction rules we have encountered: It is a kind of closure operation which transforms all drefs in stores into existentially bound variables, at levels determined by the positions of the stores that contain them. This transformation is effected by transferring the drefs in question from their stores to DRS Universes at the right levels (and then eliminating the now empty stores), thereby transforming the representation to which the rule applies into a DRS. In principle this rule could be stated in the form of a schematic transformation, as the 'lexical entry' for the covert 'main clause' feature value that the Comp feature assigns to the Comp nodes of main clauses. But such a formal assimilation of the Store Resolution rule to the rules triggered by other Comp feature values would detract from the special character of Store Resolution.

CR4.4 The rule of $\overset{!}{\vee}$ Resolution is very different from all the others. Its applications involve consistency checks on DRSs that result from merging $\overset{!}{\vee}$-disjuncts with other representations with which representations with $\overset{!}{\vee}$-disjunctions are combined in the course of representation constructions. The instances of consistency checking that played a role in those sample constructions in which $\overset{!}{\vee}$-disjunctions were involved were very simple. But that doesn't alter the fact that applications of $\overset{!}{\vee}$ Resolution involve a kind of procedure that one might feel belongs to a different level of semantic processing from that at which the other construction rules are applied that have just been listed in this summary. Rather, consistency checking would seem to belong to the same processing level that is also responsible for the resolution of presuppositions.

That should have been a reason to postpone $\overset{!}{\vee}$ Resolution to this later processing level and to postpone its discussion to Section 4. But that, though plausible in principle, would have been awkward: it would have meant dragging $\overset{!}{\vee}$-disjunctions along through further construction operations, subjecting each $\overset{!}{\vee}$-disjunct separately to each of the remaining construction rules, and end up with a final DRS for the sentence as a whole with what in some cases would have been multiple $\overset{!}{\vee}$-disjunctions, and some of them quite large, all waiting to be reduced to single disjuncts at the next processing level. Building such semantic representations would have added a lot of extra notational clutter, which would have made the representations harder to survey and the relevant points about their construction harder to make clearly and concisely.

3.11.3 Methodological Afterthoughts

What we have at this point is a set of rules of quite moderate size. But it is, we saw, an open-ended set and right now it is not easy to estimate how many further rules may need to be added as the fragment of English we want to cover will be extended. For all we know this will eventually lead to much larger rule sets.

Would that be a bad thing? How important is it for the set of rules that a theory of linguistic structure – the rules that determine the production and interpretation of language – to be small in size? There are two ways in whig this question may be understood and different sections of the linguistic community have understood it differently.

it is one of the recurrent themes in the philosophy of science that theories that manage to make do with fewer rules are to be preferred over theories that need more rules to make the same predications. But it isn't clear to what extent such general Occamist preferences are relevant to our case. A linguist may want the theory of a language she is developing (or of a language fragment or of several languages or language fragments) to be a true account of the linguistic competence of of its speakers, to the point that the rules her theory specifies should be the rules that speakers actually make use of when producing or interpretation utterances. On this view the question "Is fewer rules better?" can be answered only on the basis of what the set of rules is that speakers acquire when they learn their language and put to work when they use it. The real question is: "What are the rules that define linguistic competence?" If that isa small set, then a theory that specifies that set of rules, and thus a small set, will be better than one that specifies more, but only because the second theory has failed to come up with the right set. If the true set of rules is a large set, then a theory that specifies that set will be better than one which specifies a smaller set, even if this second theory makes the same predictions about production and interpretation as the first.

From such a psychologically realistic attitude towards the significance of the rules put forward by linguistic theories the central question is – to repeat – what the rules that characterize linguistic competence really are. But unfortunately that is a question to which it is very hard to gain proper empirical access. We are fairly good, and getting better, at determining what the possible meanings are that speakers can get for particular sentences and bits of discourse. (Introspection goes what is really a surprisingly long way, even if it certainly doesn't go far enough. But over the past decades various branches of linguistics – psycholinguistics, corpus linguistics, field linguistics – have made significant advances over what can be settled by simple introspection.) So if two theories make different predications about, say, what possible readings speakers associate with a given sentence or discourse bit, then there is good chance that we can settle which theory is wrong and which is right on this point. (Of course they could both be wrong but they couldn't both be right.) But when the theories make the very same predictions about possible readings for all sentences and discourse bits to which they both apply, but do so by postulating different rules, which imply different ways of arriving at these predictions, we are facing a quite different situation. Over the past decades psycholinguists have developed a battery of techniques for finding out more about the processing aspects of how speakers handle language, but we are still very far away from a situation in which we can, when confronted with the question which rules a speaker makes use of arriving at the different readings for a given sentence or discourse bit, we simple can take the experimental method off the shelf that we can expect to give as an answer. (Quite possibly this is a situation we will never reach.)

In the light of these considerations it isn't obvious that the strong concern there appears to be in certain quarters of the linguistic community to keep the set of semantic composition rules as small as linguistic theory can get away with hasn't got much to speak for itself. Such a concern can be detected among some of those semanticists whose work follows the central tenets of Montague Grammar, where the principle of *functional application* has been a central focus; often the methodological commitment seems to be that all that is compositional to linguistic meaning involves functional application in an essential way and that it is one of the semanticist's central tasks to explain how functional app;location is involved in cases where that is not immediately obvious. More explicitly, whenever syntax presents us with a binary mother-daughters configuration – a syntactic mother node with two daughter nodes, then the compositional semantic of this configuration is that according to which the semantics of the mother node is obtained by the function that is the semantics of one of the daughters to the argument provided by the other daughter. (This commitments is often reinforced by the view that syntactic mother-daughters configurations never involve more than two daughters.)

Those working in a Montagovian paradigm are likely to protest. Referring to their advocacy of fictional application as the only orb primary principle of semantic composition in the context of the present discussion is misleading, they might say, because psychological relevance is not what their theories are after and that is so in particular for the rules and principles they propose. For them the question whether fewer rules are better can only be a question of Occamist parsimony. And of course, on that score fewer rules em is better.

DRT is different. Here the processes involved in human interpretation of linguistic signals has been a motivating concern from the start (although it is certainly possible to adapt DRT as a method for describing the form-meaning relation without any commitments to psychological relevance). DRT does therefore expose itself to questions about the cognitive adequacy of the rules it proposes. This renders it vulnerable to the problem that it must acknowledge questions as empirically significant for which no empirical verification methods exist now and for which no such methods may ever be found. Related to this problem is another, which we have also touched upon here and there. A host of psychological results, some of it going back to the sixties and seventies, suggests something that should have been fairly plausible also in the absence of any explicit experimental evidence – different aspects of linguistic processing – phonological parsing, syntactic parsing and the construction of semantic representations – all happen simultaneously and on line. So the kind of pipeline model that is an essential feature of the architecture we have adopted – first a complete syntactic parse is established and this structure, the complete LF, then serves as point of departure for the construction of the semantic representation – is unrealistic from a psycholinguistic point of view in any case. For this reason question about psychological adequacy of DRT must be handled with the greatest circumspection, especially when it is not inter sentential connections that are at stake, but the processing of the internal structure of individual sentences.

3.11.4 Model-theoretic Semantics of the Representations

From the beginning DRT has taken the line that a model-theoretic semantics must be given for the representation it postulates. But this commitment has for the most part been restricted to completed DRSs, those that the theory assigns to fully interpreted sentences and bits of discourse. With the adoption of the Bottom Up construction of semantic representations, however, the commitment has changed somewhat. The Bottom Up construction method is meant to model aspects of the compositional nature of linguistic meaning more faithfully than this was possible in the previously used Top Down approach. The representations of sentence constituents that are now built from the representations of their constituents now are to thought of, as they typically are in Formal Semantics, as representations that determine their own denotations; so it is reasonable to demand that the model theory provides those representations with semantic values in models as well (and not just the DRSs that are the final outcomes of the composition process).

In relation to the DRS constructions described in Section 3 the demand for model-theoretic values for the intermediate representations (and not only for the final products) comes to this. What are the semantic values determined by representations of the form $< ST \mid K >$, in which a DRS K, which may contain argument slot symbols as well as drefs, is preceded by a store ST? At first sight this problem might seem much harder than it is for completed DRSs. But on closer inspection the problem turns out not to be much more difficult. The main point is that the semantic values that structures of the form $< ST \mid K >$ determine in different models should be defined as functions whose values are the values of DRSs, such as K. In fact, these functions are best described as 'two-level' functions – functions that when applied to a proper input (a tuple consisting of arguments for the argument slots of the function) yield another function, which yields when it is applied to a tuple of arguments appropriate for it, as its value the semantic value of a DRS. This 'two level' conception of the functions in question is motivated by the observation that the possible semantic values of the DRS K of a semantic representation $< ST \mid K >$, that is, the different truth values that K can determine in different models, depend on to kinds of 'unknowns': (i) the values represented by the drefs in the store ST and (ii) the values that could be represented by drefs that could be inserted into the argument slots in K.

This description of what the semantic value of K in $< ST \mid K >$ in a model M depends on not only motivates the assumption that the semantic value of $< ST \mid K >$ in M is a function from certain arguments to truth values of K in M, but also to see this function as 'two-leveled' because the arguments on which the values of K in M depends can be divided into two classes: (i) values of drefs in the store and (ii) values for argument slots. We will assume that the two level function that is the values of $< ST \mid K >$ in M takes for the argument slots as its 'outer' arguments, and values for the drefs in ST as 'inner arguments'. That is, the value of $< ST \mid K >$ in M is a function from tuples consisting of possible value for the argument slots in K to functions from tuples of values for the drefs in ST to truth values in M for K.[76]

Apart from the complication hinted at in the last footnote the stipulations of the last couple of paragraphs reduce the semantic value problem for intermediate representations of the form $< ST \mid K >$ to the model-theoretic semantics for completed DRSs. Probably enough has been said here about this latter problem (in particular in Section 3.6). Fully explicit accounts of the model theory for somewhat different DRS languages than the one developed here can be found in (Kamp & Reyle 1993) and (Kamp, van Genabith & Reyle 2011).

3.11.5 The Use of λ Notation as Syntactic Sugar

[N.B. This subsection relates to work that has been done within the project B4 of the SFB 732, where representations do not simply consist of stores and DRSs but also make use of λ-binding.]

[76] A complication for this characterization of the value of $< ST \mid K >$ in M are stores in subordinate positions, such as we have seen in Sections 3.9 and 3.10. The contributions that are made,for instance, by complex conditions that contain subordinate stores will depend on who these stores will eventually be resolved. A proper solution to these complications depends on matters that can be dealt with only using tools not yet at our disposal. We will return to this matter in Section 4.

There is an extension of the representation formalism developed in Section 3 and summarized in the last three subsections in which representations can have λ *prefixes*. In this formalism we do not only have representations of the form $< ST \mid K >$ but also representations in which a structure of type $< ST \mid K >$ is preceded by a λ *prefix*, consisting of one or more lambda binders, as in (3.326).

(3.326)$\lambda\delta_1 ... \lambda\delta_n . < ST \mid K >$, where n is some number ≤ 1 and $\delta_1, ..., \delta_n$ are discourse referents of various sorts and types.

The model-theoretic semantics of such terms can be defined along the same lines as we did for representations of the form $< ST \mid K >$ in the last section. It is just that the λ prefix adds another layer of arguments: the semantic value of $\lambda\delta_1 ... \lambda\delta_n . < ST \mid K >$ is a 'three-level function' (rather than a 'two-level one) – a function which maps n-tuples of possible values for the drefs $\delta_1 ... \lambda\delta_n$ to functions that are the values of $< ST \mid K >$ (that we get when these values are assigned to the free occurrences of $\delta_1 ... \lambda\delta_n$ within K).

The point of introducing this form of λ notation into our formalism is to make certain aspects of linguistic composition more easy to understand, and also to make the syntax-semantics interface more constrained. The use of the symbol 'λ' is of course not accidental – there is an intended reference here to the λ calculus, and many aspects of the use we make of λ's here closely resemble aspects of the syntax and semantics of the λ calculus. The syntactic specification in (3.326) and the semantics described in the last paragraph are cases in point. But the differences are at least as important as the similarities and it is on the differences that are in focus in this section.

The most important difference is in how terms of the λ calculus are used in the λ calculus and how terms of the form (3.326) are used in the extension we are discussing of the formalism of Section 3. A term $\lambda\alpha.\Phi(\alpha)$ of the λ calculus can be combined with another term in only one of two ways, and both of these are instances of functional application: either (i) $\lambda\alpha.\Phi(\alpha)$ plays the role of function and the other term, which has to be of the same type as α, serves as argument, or (ii) $\lambda\alpha.\Phi(\alpha)$ plays the role of argument and the other term plays the part of function (in which case $\lambda\alpha.\Phi(\alpha)$ must be of the right type to be able to serve as argument for this second term). The syntactic execution of functional application in the λ calculus, we have seen, takes the form of applications of the rule of β-Conversion, in which, roughly speaking, the λ-prefix of the term that plays the part of function is removed from it and the other term is substituted for the remaining occurrences of

the variable from this λ-prefix. (It follows from the syntax and semantics of the λ calculus that the semantic value of the term that results from these operations is always the result of applying the semantic value of the function term to the semantic value of the argument term.)
The combinatory options of λ terms of our extended representation formalism – those of the general form (3.325) are similar. In fact, for these representations there is only one possibility, in which $\lambda\delta_1...\lambda\delta_n. < ST \mid K >$ plays the part of 'function term'. The representation that it combines with has to be of the form$< \delta_{1,ref}, ... \mid K'' >$, in which the referential argument is the same dref as the one that is bound by the outer λ of $\lambda\delta_1...\lambda\delta_n. < ST \mid K >$. The result of this operation is the following representation:

(3.327)$\lambda\delta_2...\lambda\delta_n. < ST, \delta_{1,ref}, ... \mid K' \bigoplus K''>$,

> where K' is the result of replacing all occurrences in K of argument slot symbols bearing the index i by β.

It is easy to see that the semantic value of (3.327) is not the result of applying functional application to the semantic value of the representations from which (3.327) has been obtained via Argument Insertion. (The formal presuppositions for applying functional application to these semantic values aren't even satisfied.) In fact, to describe in plain English how the semantic value of $\lambda\delta_2...\lambda\delta_n. < ST, \delta_{1,ref}, ... \mid K' \bigoplus K''>$ relates to those of $\lambda\delta_1...\lambda\delta_n. < ST \mid K>$ and $< \delta_{1,ref}, ... \mid K''>$ is not all that easy, and it wouldn't be very helpful. However, hat we have said about the semantics of representations of our formalism in 3.12.3 and about the semantics of representations with λ prefixes above should make it clear enough what this semantic connection is. And it should make it clear in particular that this isn't plain functional application, and in how it differs from that.

Consider the special case of this operation where the λ prefix of the first representation consists of a single variable binding, as in $\lambda\delta_1. < ST \mid K >$. In this case combination of the representation with a representation of the form $< \delta_{1,ref}, ... \mid K'' >$ is the representation $< ST, \delta_{1,ref}, ... \mid K' \bigoplus K'' >$. This representation could also be obtained by applying the rule of Argument Insertion, as defined in Section 3.12.2, to the representations $< ST \mid K >$ and $< \beta_{ref,i}, ... \mid K''>$. And much the same is true of representations with longer λ-prefixes. For instance, a representation of the form $\lambda\delta_1.\lambda\delta_2. < ST \mid K >$ will have to combine with two representations $< \delta_{1,ref}, ... \mid K''>$ and $< \delta_{2,ref}, ... \mid K^3>$ with the same result as we get from two applications of

Argument Insertion when we start with the representation $< ST \mid K^4 >$, where $K^4 >$ is like K except that its has argument slots $\underline{\delta_i}$ where K has occurrences of the drefs δ_i (i = 1,2).

More generally, the use of λ-prefixed representations in semantic composition can be simulated by Arguemnt Insertion, so long as the compositional system is set up to provide representations with argument slots instead of the λ prefixes; and up to a point the converse is true as well. So, as far as this is concerned, adding λ-prefixed representations to our formalism is a kind of syntactic sugar. Why do it?

The point is precisely that the options for integrating representations into larger representations is so very limited for λ-prefixed representations. To specify a certain λ-prefixed representation as representation of a certain syntactic constituent is therefore a way of making clear that the semantic representations for such constituents can be integrated into larger representations in only the one way that λ-prefixed representations permit. In this regard λ-prefixed representations are quite different from representations without λ prefixes. These have no restrictions built into their form as to how they may be integrated into larger representations; by and large, any of the rules described in 3.12.2 could in principle apply to them so long as the syntactic structure would demand that. Making use of λ-prefixed representations in the formulation of a syntax-semantics interface is therefore a way of reintroducing into a system that is quite liberal in the compositions rules it admits some of the constraints that many have seen as one of the great virtues of the use of the λ calculus (some version of it, such as Montague's Higher Order Intensional Logic) in natural language semantics: the compositional operations are in principle very limited – functional application, and usually certain cases of λ abstraction, as in the set up of Heim and Kratzer discussed in Part I, and sometimes also some further operations that can be defined interns of these two.

So the introduction of λ prefixes into the formalism, which will at face value look like a further extension of representational and compositional options, can in fact be exploited to impose tighter constraints on the parts that semantic composition may follow. So far, the main use that has been made of this option is in the analysis of 'sub-lexical' syntactic and semantic structure, of the kind that has been the central focus of the work in the B4 project (see in particular (Rossdeutscher 2013)). But there is no reason why it could not be used with equal effect at higher levels of compositional combination.

Chapter 4

Presupposition

4.0.6 Fast-forward Run through the History of Theorizing about Presupposition, from Frege to Van Der Sandt

4.0.6.1 The Logicians and the Philosophers

Frege is usually credited with introducing the topic of presupposition into logic and semantics. In his essay 'On Sense and Reference (Frege 1892) we find the following passage:

(4.1) The sense of the sentence: "After Schleswig-Holstein was separated from Denmark, Prussia and Austria quarreled" can also be rendered in the form "After the separation of Schleswig-Holstein from Denmark, Prussia and Austria quarreled". In this version it is surely sufficiently clear that the sense is not to be taken as having as a part the thought that Schleswig-Holstein was once separated from Denmark, but that this is the necessary presupposition in order for the expression "after the separation of Schleswig-Holstein from Denmark" to have a reference at all.

According to Frege definite descriptions (among them mathematical descriptions like *the smallest prime*, *the largest prime*, *the largest prime pair* and so on) come with a presupposition that their descriptive content has a unique satisfier. For Frege, who saw failure of reference of a descriptive term in a sentence as entailing a failure of the sentence to determine a truth value, definite descriptions whose unique satisfaction presuppositions are not satisfied constituted a threat to logic, which he saw as grounded in the principle of bivalence – the principle that every declarative sentence is either true or false.

In his famous 'On Denoting' (Russell 1905) and elsewhere Russell argued that Frege's worries about the threat of presuppositions to formal logic were unnecessary. The bivalence of sentences with definite descriptions could be salvaged by assuming that the determiner which heads such noun phrases – in English the definite article *the* – can be teated as a quantifier, just as Frege himself had proposed for the determiners *every/all* and *a/some*. The semantics of *the* differs from those of *every/all* and *a/some* in being a combination of an existential and a universal condition – that there is a satisfier of the descriptive content and that every satisfier of this content is identical to this one. For example, when applied to what has served as a kind of paradigm example in the debates over presupposition for more than a century now, the sentence *The King of France is bald*, the logical form attributed to it by Russell's proposal can be given in the form of the formula of Predicate Logic shown in (4.2.b)

(4.2) a. The King of France is bald.

b. $(\exists x)(\text{King-of-France}(x)\ \&\ (\forall y)(\text{King-of-France}(y) \rightarrow y = x))$

Russell's proposal to analyze definite descriptions as quantificational complexes is known as his 'Theory of Descriptions'. We will refer to it here sometimes as 'RTD'. On the face of it RTD has many attractive features. It restores bivalence for all sentences with 'non-denoting' definite descriptions, while making use of nothing more than the quantificational devices that Frege himself had already introduced in his Begriffsschrift. Moreover, Russell argued, treating *the* as a quantifying determiner makes it possible to assign to a sentence like 'The Golden Mountain doesn't exist', whose definite description 'The Golden Mountain' is also non-denoting in Frege's sense – there is no Golden Mountain – a logical form according to which it is true. This is in accordance with our intuitions (that the sentence *is* true) and thus is superior to the Fregean account, on which it is without truth value.

But not all that glitters is gold. For a linguist Russell's Theory of Descriptions is problematic. The theory says far too little about how the syntactic forms of natural language sentences determine their logical forms. For instance, Russell's discussion of the Golden Mountain example is tacit on what it is about the English sentence 'The Golden Mountain doesn't exist.' that licenses a logical form in which the negation has wide scope over the definite article. For all Russell tells us about this example, a similar logical form, in which *the* has narrow scope vis-á-vis the negation, would be equally acceptable possible for a sentence like 'The Golden Mountain is not as tall as

Aconcagua'. On such a logical form this second sentence would also have a reading on which it is true, again for the sole reason that there is no Golden Mountain. But in this case the predication, and the reason for it, seem much less obvious. In fact, RTD makes a different predication for the sentence 'Aconcagua is taller than the Golden Mountain'. This sentence is for all practical purposes equivalent to the previous one. But for this last sentence RTD seems to make available just one logical form, which predicts the sentence to be false.

It took some considerable time before the world was ready for this type of criticism of Russell's proposal. The first overt and reverberating criticism of RTD had to do with a different aspect of it. In his paper 'On referring' ((Strawson 1950)) Peter Strawson drew a distinction of which there is no clear trace in RTD – that between (i) the question whether a sentence (a syntactically well-formed sequence of words) is meaningful and (ii) the question whether its actual use on some particular occasion succeeds in expressing a proposition, and thereby determines a truth value. For instance, in connection with Russell's example (4.2) Strawson observes that it is a sentence with a perfectly clear meaning, which it has irrespective of whether it is uttered in 1950 or in 1750 (or whatever other time). But utterances of the sentence at 1750 and 1950 differ in that an utterance in 1750 would have expressed a definite proposition, viz. that the person who was the King of France at that time (Louis XV) was bald, whereas an utterance in 1950 does not express any proposition because at this time there was no King of France.

Strawson's distinction naturally led to the notion that utterances of sentences containing definite descriptions carry 'unique satisfaction presuppositions': such utterances express propositions only when these presuppositions are satisfied. That one utterance of a sentence like (4.2) can express a proposition while another utterance, made at a different time, does not, has to do with the fact that the extensions of descriptions may vary with time, so that the same description may have a unique satisfier at one time, but no satisfier (or more than one) at some other time.

On this account of presupposition it is *utterances* of definite descriptions, that come with presuppositions. And when definite descriptions occur as parts of sentences, it is the utterances of those sentences, rather than the sentences qua linguistic expressions, that come with presuppositions, associated with the utterances of the definite descriptions that are part of these utterances. This makes presupposition a *pragmatic* phenomenon: Presupposing is something that is done by those who produce utterances, as part of the complex acts that making an utterance constitutes. Presuppositions are thus not properties of linguistic expressions as such, even if there may be systematic connections between the presuppositions that speakers make as part of their utterances and the forms of words that their utterance make use of.

A strong case for such a pragmatic view of presuppositions can be found in the work of Stalnaker, see in particular (Stalnaker 1972), (Stalnaker 1974), (Stalnaker 1979). Stalnaker has argued that presupposing is something that is done by speakers. When a speaker makes an utterance she can, and typically will, make certain assumptions about what is taken for granted in the context in which her utterance is made and in which it is to be understood. These presuppositions may have something to do with the form of the expressions she uses, but they can also be motivated in other ways. For instance, a speaker may assume that her addressee is familiar with certain facts without which the information that her utterance overtly expresses would be useless to him, as in a case of the following sort: if you (the speaker) answer my question how far it is to Amsterdam by answering that Amsterdam is about as far as Paris, then you are likely presupposing that I know what the distance to Paris is. For without that your answer won't get me much nearer to what I want to know.

Missing from this characterization of Stalnaker's account is a distinction that is crucial for linguistics. Certain expressions carry certain presuppositions just by virtue of their form. When a speaker uses such an expression, she cannot help but make the presupposition or presuppositions associated with it. No matter what she may think or intend, by using the expression she will *count* as making these presuppositions. Definite descriptions provide a good illustration of this (though the point that they illustrate is one that equally applies to many other presupposition-carrying expressions, as we will see momentarily). As Frege perceived correctly, the unique satisfaction presupposition carried by singular definite descriptions are presuppositions that are associated with them by virtue of their morpho-syntax – as singular noun phrases beginning with the determiner *the*. Whether the presupposition associated with a given definite description is satisfied may depend on the time at which the description is used. But the presupposition itself is associated with the description once and for all.

For the linguist it has been such form-determined presuppositions that are the primary objects of interest. More about such presuppositions in Section 4.1.2. We conclude the present section with a couple of methodological remarks.

The first of these is a general comment on the relationship between the 'semantic' view of presuppositions, according to which they are part of the semantics of certain expressions, and the 'pragmatic' view, according to which presupposing is something that is done by utterance producers. Contrary to what has sometimes been implied, there is no real opposition between these two views. Rather, they can be seen as two sides of a single coin: Because certain expressions that have certain presuppositions associated with them in the sense of the semantic view speakers and authors who produce sentences containing those expressions are bound to make the corresponding pragmatic presuppositions. This doesn't prevent speakers or authors from making presuppositions other than those to which they are committed by their use of presupposition-triggering expressions. But when a presupposition-carrying-expression is used, then it is the coin as a whole that is brought into play, with its semantic and its pragmatic side.

The second remark relates to a point that will be relevant to much that is going to be said about presupposition in what follows. Among the observations that Strawson made about the presuppositions of definite descriptions is that such a presupposition is 'entailed' both by a simple sentence containing it and by the negation of that sentence. For instance, the presupposition that there is a unique King of France, associated with the definite description *the King of France* can, in some sense, be inferred both from the sentence 'The King of France is not bald' ((4.2.a) above) and from its negation 'The King of France is not bald'. Neither of these two sentences can be true unless the presupposition is true; and so, in other words, both (4.2.a) and its negation entail the proposition that there is a King of France. More formally, and abbreviating 'The King of France is bald' as p and 'There is a unique King of France' as q, we have that q is entailed by both p and $\neg p$. In classical logic this would entail that q is a tautology. But of course the proposition that there is a unique King of France is not a tautology.

The conclusion would seem to be that in the presence of presuppositions classical logic breaks down – supporting the fears Frege had expressed in relation to non-denoting descriptions like *the largest prime*. Many efforts have been made to cope with this apparent collapse of classical logic by trying to develop an alternative 'formal logic of presuppositions', which coincides with classical logic for sentences that are presupposition-free, but in which the apparent conclusion reached above – that when q follows from p and also from $\neg p$, then q is a logical truth – does not follow when presuppositions are involved. I will argue later that these attempts rest on a misconception of the relation between language and logic. But note well, it is important to distinguish between the observation that both (4.2.a) and its negation entail the proposition that there is a King of France and the need for a presupposition logic of the sort just hinted at. While, as we will see that there is no need for such a logic, the observation is true, and any account of presuppositions must be able to deal with it in some way.

The principle that a presupposition that comes with a sentence S is also carried by the negation of S can be seen as one of a number of general facts about presuppositions. In Section 4.1.3 we will see that it is one from a range of such principles, which are characteristic of presuppositions and which can be used as tests for determining whether a given proposition q does or does not stand in a presuppositional relation to a sentence S. But before we turn to such tests for presuppositionality, we will first, in the next section, give an impression of the variety of words and grammatical constructions that have been identified as presupposition carriers – or as*presupposition triggers*, we we will say from now on, complying with what has become the standard term for this within linguistics.

4.0.6.2 The Linguistic Turn: Many different kinds of Presuppositions and Presupposition Triggers

For about three quarters of a century the concerns over presupposition were the concerns of philosophers and logicians, and they were concerns over the presuppositionality of definite descriptions. Within that narrowly confined context, and given the rather cavalier attitudes of those involved in the debate with regard to details of how logical form is determined by natural language syntax, it is not surprising that many went along with RTD, which dismissed the worries about presupposition failure as an outgrowth of what they considered a mistaken view from the start, and one that Russell had debunked. But all that changed when in the late sixties and early seventies linguists began to realize that what had come to be identified as the central traits of presuppositional behavior wasn't restricted to definite descriptions. It gradually dawned on the linguistic community that many expressions and constructions have properties that are much like those which Frege and Strawson had seen as characteristic of definite descriptions. These expressions and constructions have conditions associated with them that must be satisfied lest the expressions and constructions with which they are associated fail to denote or are infelicitous in some other way. It seemed natural, therefore, to see these conditions also as 'presuppositions' of the expressions or constructions with which they are associated. And with this widening of the notion of presupposition definite descriptions came to be viewed as specimens of a kind – as instances of a phenomenon that had proved to be far more common than had been perceived until then.

Note well that the broad (and still quite unspecific) description we have just given of presuppositions does not match up with what we said in our discussion of the views of Frege and Strawson. is not quite the same definition that we discussed when talking about the views of presuppositionality that can be found in Frege and Strawson. On their view of presupposition failure was that a definite description whose presupposition fails does not denote and that an utterance involving this kind of failure fails to express a proposition. But on the revised view that is only one effect that presupposition failure can have. In some cases an utterance that involves presupposition failure can still succeed in expressing a proposition. But the way that proposition is worded is inappropriate. And it is inappropriate because of the presupposition failure.

What you find below is an illustrative list of English words and grammatical constructions that have been identified as presupposition triggers and of the presuppositions they trigger.

(4.3) a. **Factive verbs.**

(i) emotive: *regret, be happy, be sorry, ..*
(ii) epistemic: *know, realize, discover,..*

Examples:

(ad i) 'John regrets/doesn't regret that he went to the concert.'

Presupposes that John went to the concert.

(ad ii) 'She has/hasn't discovered that he is having an affair.'

Presupposes that he is having an affair.

b. **Aspectual verbs**: *stop, start, begin, continue, carry on, remain,..*

'John has/hasn't stopped smoking.'

Presupposes that John smoked at some time in the past.

c. **Additive particles**: *too, also, as well, even,..*

'Nixon is/isn't guilty too.'

Default interpretation: Presupposes that someone other than Nixon is guilty.

'She also gave him/did not also give him a kiss.

Possible interpretations:

(i) Presupposes that she gave him something else as well;

(ii) Presupposes that she gave someone else a kiss too;

(iii) Presupposes that someone other than her also gave him a kiss.

d. *again*

'Mary closed/did not close the window again.'

Possible interpretations:

(i) Presupposes that Mary closed the window before (*repetitive* reading of *again*)

(ii) Presupposes that the window was previously closed (*restitutive* reading of *again*)

e. *still*

'Mary is still here.'

Presupposes that Mary was here at some time before now and that this state of affairs has continued up to the present time.[1]

[1] All previous examples involved the presentation of an unnegated sentence together with its negation, to make plain that the presuppositions of unnegated sentences are also presuppositions of their negations. For the sentence 'Mary is still here.' this doesn't quite work: 'Mary isn't still here.' is not a particularly natural sentence; the natural way to express the negation of 'Mary is still here.' is something like 'Mary isn't here any more' or 'Mary isn't here any longer.' But note that these last two sentences also presuppose that Mary was here at some earlier time.

This situation – that there is no straightforward negation for a presupposition-triggering sentence – arises quite often, and this is so in particular for many of the examples below. It should be possible at this point to recognize that the examples do generate the specified presuppositions.

f. **Clefts**

‘It was/wasn’t Fred who solved the problem.’

Presupposes that someone solved the problem.

g. **Why-questions and how-questions**

‘Why did you take my piece of cake?’

Presupposes that the addressee took the speaker’s piece of cake.

‘How did you manage to get into the safe?’

Presupposes that the addressee managed to get into the safe.

h. **Past tenses**:

Examples:

‘Mary turned off/ didn’t turn off the stove.’

Presupposes that there was some particular, independently identifiable time in the past; the sentence asserts/denies thatMary turned off the stove at that time.

‘Mary had been/hadn’t been surprised ’

Presupposes that there was some particular, independently identifiable time in the past; the sentence asserts/denies that Mary was furious at that time.

i. **Pre-state verbs**: *go to Paris*, *die*, *promote to full professor*, ..

Examples:

‘On Tuesday Mary went/didn’t go to Paris.’

Presupposes that at some time on Tuesday (or just before Tuesday) Mary wasn’t in Paris.

Note that the impossibility to form negations of many presupposition-triggering sentences means that the ‘negation test’ for presupposition – a condition associated with a sentence is a presupposition iff it is likewise associated with its negation – has only limited applicability. We return to this point in the next subsection.

j. **Argument selection restrictions of lexical predicates**

Many lexical predicates presuppose that one or more of their argument positions must be filled with arguments of certain ontological sorts.[2]

dead, *die*, *kill*, *murder*, *assassinate*, ..; *prevent*, ..; *believe*, ..

(i) argument of *dead*: something for which the distinction between being dead and being alive is a meaningful distinction (organism);

(ii) subject of *die*/direct object of *kill*: same as for the argument of *dead*;

(iii) subject of *murder*: human being;

(iv) direct object of *murder*: human being;

(v) subject of *assassinate*: human being (like subject of *kill*);

(vi) direct object of *assassinate*: An individual holding a position of power or having considerable public political, social or cultural status;.

(vii) direct object of *prevent*:

(a) an eventuality type; in combination with a phrase or clause that denotes an eventuality type *prevent* says that an eventuality of this type did not occur over some explicit or understood period of time;

(b) an agent; in this case *prevent* has an additional argument which takes the form of the preposition *from* followed by a gerundive eventuality description.

(vi) subject of *believe*: agent capable of entertaining propositional attitudes like belief;

(vii) direct object of *believe*:

(a) a propositional content or entity that determines a propositional content (e.g. a statement or the speaker or author producing it);

[2]Questions of argument selection arise in connection with the vast majority of natural language predicate words. In fact, selection restrictions are a topic of investigation in its own right. (For extensive discussion see (Asher 2011).) Here is a small, more or less arbitrarily chosen sample.

(b) an agent (capable of entertaining beliefs)

k. **Focus**

'FRED solved the problem.' (Capitalization indicates focal prosodic stress.)

Presupposes that someone solved the problem.

l. **Definite Noun Phrases**:

Presuppose a way of identifying their referents on the basis of their descriptive content and, often, additional information that is presumed to be available to the interpreter.

(N.B. Definite descriptions are a special case of this, with identification presuppositions guided by their own specific principles for reference identification.)

What justification is there for classifying all the 'presupposed' conditions described in (4.3) as presuppositions? So far our only criterion for categorizing conditions that are implied by expressions as presuppositions of those expressions is that they are not only implied by the sentence itself but also by its negation. But we already saw (fn 1) that this criterion cannot always be straightforwardly applied. Either a sreaightforward negation cannot be formed (as for the sentence 'Mary is still here' that prompted fn 1); or else what may look like the straightforward negation of the sentence is clearly not the one we want as confirmation that a condition associated with the non-negated is a presupposition of it. For instance, the negation 'Why didn't you take my piece of cake' of the *why*-question 'Why did you take my piece of cake' in (4.3.g), has nothing to do with the presuppositional status of the condition that the addressee took the speaker's piece of cake. Rather, this sentence has its own presupposition, viz. that the addressee did not take the speaker's piece of cake.

In fact, from a linguistic point of view the so-called negation test for presuppositionality is a fragile one at best. For a range of sentences the test seems to work the way it should: there is a simple method for turning such a sentence S into its negation – put a negation onto its main verb! – which preserves the presuppositions of S. But far too often this method doesn't work in the way it is supposed to. Either it cannot be applied (as we saw for our sentence with *still*), or else it produces a sentence 'not-S' that doesn't stand the right semantic relation to S. Putting the problem this way has a kind of circular ring to it. But it is hard not to state the problem in non-circular terms, and that in fact is part of what the problem is.[3]

In the next section we are going to look a little more closely at the problem of finding tests for when expressions or constructions are presupposition triggers and what their presuppositions are.

4.0.6.3 Tests for Presuppositionality

If the negation test were the only test we had for determining what presupposes what, then many of the claims about presuppositions that can be found in the literature would be on thin ice. But fortunately there are other tests as well. Perhaps the most important one is a test that nowadays is commonly referred to as the 'Hey, wait a minute' test. The 'Hey, wait a minute' test is one that directly applies only to questions – both polar questions and various types of wh-questions. Suppose for instance that A puts the *why*-question in (4.3.g) to B, but that B didn't take A's piece of cake. What answer can B give to A's question? Evidently no answer of the kind that one is expected to give to a *why*-question will do. For the fact that one is asked to account for doesn't exist – the question demands an explanation for something, but in fact there is no such thing. So the only thing B can do in response to A's question is to signal that the question is an improper one (and that the reason for that is precisely that its presupposition – that B took A's piece of cake – isn't satisfied). One way to signal this is to respond

[3]Another problem with the negation test (and arguably an even more serious one) is that negation can often be understood as targeting the presupposition, or as including it within its target. Readings in which negations are interpreted in this way are prominent when the failure of a presupposition is given as reason why a sentence is true provided that its negation is interpreted in such a way. A classical illustration is given in (4.4.

(4.4) The King of France didn't open the exhibition. There is no King of France.

The prominent interpretation of this sentence combination is that the first sentence is true and that it is because the presupposition carried by the phrase *the King of France*: What is wrong with the unnegated sentence 'The King of France opened the exhibition' is that its presupposition – that there is a unique King of France – is not true. That makes it right to deny this sentence and why the result of negating it (viz. the first sentence of (4.4) can be regarded as true.
The capacity of negation to target a presupposition of the sentence it negates is just part of a much larger and essentially open-ended spectrum of aspects in which a sentence may be inadequate, from propositional content all the way to questions of orthography and pronunciation. This flexibility of negation has found an able, careful and thorough reporter in Larry Horn, as witnessed by his signal contribution (**?**). This problem with the negation test is one about which we will have nothing to say in these notes.

with ‘Hey, wait a minute’ and then go on with the reason for this response, as in: ‘Hey, wait a minute, I never took your piece of cake.’ (There are other such signals too, such as ‘What do you mean?. But ‘Hey, wait a minute’ has become a kind of icon for the range of the natural responses in such a situation. ((Von Fintel 2004)).)

The ‘Hey, wait a minute’ test for presuppositions of questions is a very intuitive one: when you know that what the question you have been asked presupposes isn’t true, then you simply cannot answer the question in a standard way. The only thing you can do is to make it clear to the asker that the preconditions for a proper answer are not given. The need for ways to signal this kind of impossibility is particularly clear for polar questions and for *why*- and *how*-questions. The standard way to answer a polar question is with either ‘yes’ or ‘no’. But when the question contains an unsatisfied presupposition, then neither answer will be right, because either will come across as endorsing the false presupposition. Awareness of this goes back to Greek antiquity, where asking polar questions of which the asker knew that they carried false presuppositions was a device used in cross-examinations to drive a witness or defendant into a corner where they did not deserve to be. If you are asked ‘Have you stopped stealing money from your mother?’ and you never stole money from your mother at any time, then neither the answer ‘yes’ nor the answer ‘no’ will do you much good. If you say, ‘yes’, you are wrongly endorsing that you did steal money from your mother, even if you also (and truly) assert that you do not steal money from your mother now and haven’t done so for some time. And saying ‘No’ is worse, for it not only endorses the false presupposition that you did steal money from your mother, but adds to this the further false claim that you haven’t stopped and thus carried on stealing money from your mother up to the present time.

For *why*- and *how*-questions the ‘Hey, wait a minute’ test works equally well and for similarly good reasons. Any regular answer to a *why*-question, which supplies some kind of reason, counts as an acknowledgment that that for which a reason is being given actually happened, or is/was the case. Therefore, if A puts a *why*-question to B and B knows that that for which it requests an explanation didn’t happen, then the response must make clear that the question is inappropriate: B should respond with ‘Hey, wait a minute’, or something in this spirit. The same problem can arise for *how*-, *when*- and *where*-questions. When there is nothing that the question requests details about - a recipe, or some kind of explanation or a location in time or space –, then an answer that complies with the request cannot be given. Such questions also call for a response which signals that they cannot be answered in the normal way.

(Parenthetical remark: For other *wh*-questions, in which the *wh*-word is *who*, *which* 0r *what*, it is less clear how the 'Hey, wait a minute' test can be applied. Consider for instance the question 'Who stole my piece of cake?'. Suppose that the person A how puts this question to B is a glutton prone to amnesia, who has eaten his piece of cake himself, but whose craving for another piece has rendered him oblivious to this fact. How should B, who knows that A himself ate his piece of cake, react to this question? Two kinds of reactions come to mind. One is 'Nobody did' and the other 'What do you mean? Nobody did.' These answers convey the same information, of course: that the set of thieves of A's piece of cake is empty. But they do so in subtly different ways. The first – 'Nobody did' – sounds like a regular response to A's question: that the set of cake-stealers is empty is one of the possible alternatives between which the question invites the addressee to choose. The second answer – 'What do you mean? Nobody did.' – rather sounds like the state of affairs it describes – that nobody stole the piece of cake – is not among the possible alternatives that the question presents. This second response seems to suggest that the question is understood as presupposing that the piece of cake was stolen.

Whether A's question does carry this presupposition is complicated by the fact that someone who utters this question will almost certainly assume that someone did steal their piece of cake – that, as Stalnaker would put it, he, as producer off the utterance, is presupposing this. That would be enough to explain why the second response is justified: By reacting in this way the addressee conveys to the questioner that he rejects the assumption which he takes the questioner to be making. But we need to distinguish here between the presuppositions that speakers make and the presuppositions that they are forced to make because of the expressions they utter. The fact that the first response – 'Nobody did' – to the piece of cake question is acceptable suggests that the question need not be understood as carrying the implication that someone stole the piece of cake.The interpreter is at liberty to treat the question as not carrying this presupposition, and therefore also to ignore the assumption that the questioner is making – that the piece was stolen –, even if it is plain that the questioner is making that assumption. Which *who*-, *which*- or *what*-questions come with built-in non-emptiness presuppositions and which do not remains a topic of debate. End parenthetical remark.)
The 'Hey, wait a minute' test is not only applicable to questions but also to sentences of other types, including in particular declarative sentences. The reason is fairly obvious. Suppose that S is a declarative sentence and that

S? is the interrogative corresponding to S. Then if the proposition q is a presupposition of S, q will equally be a presupposition of S?. So if q is false, then 'Hey, wait a minute' will be the appropriate reaction to S?. For both answering S? with 'yes' and with 'no' will convey that q is true. By the same token 'Hey, wait a minute' will also be the right reaction to an assertion of S. For objecting to S, e.g. by saying something like 'No, that isn't true.', will imply that q is true. And confirming S – even the implicit kind of conformation that is conveyed by not challenging S – will also count as a confirmation of q. Thus 'Hey, wait a minute' should be a proper reaction to an assertion of S in case q is false just as it is an appropriate reaction to the question S?. And indeed this is the correspondence we observe: 'Hey, wait a minute' feels like it is the right reaction to the assertion just in case it feels like the right reaction to the question.

For the same reason there also is a close connection between the 'Hey, wait a minute' test and the negation test in those cases where the latter test is applicable. For suppose that S has the presupposition q and that q is false. Then 'Hey, wait a minute' is the right reaction to the assertion of S, because both confirming and denying S will count as confirmations of q. The proper reaction must signal that something is wrong with the assertion and not just that that it is false. But these are of course precisely those cases when something is wrong both with the assertion of S and the assertion of $not - S$: neither can be (felicitous and) true in situations where q fails.

The upshot of these considerations is that both the 'Hey, wait a minute' test for declarative sentences and the negation test (to the extent that that test can be applied at all) can be derived from the 'Hey, wait a minute' test for questions. That leaves us with a single test. If it is a really good test, always applicable and always making the right predictions, then that is arguably all we need. And the 'Hey, wait a minute' test does appear to do a pretty good job on both counts. In the presupposition literature a further tests have been proposed, however. One is known as the 'modality test', or alternatively as the 'possibility test'; I will use the term 'possibility test' in what follows. The possibility test is like the negation test one that involves a general recipe for transforming sentences into grammatically related sentences. The transforms involved in the possibility test are *modalizations* of the sentences they transform– sentences that are obtained by either prefixing a given sentence S with a modal operator like 'it is possible that' or by 'modalizing' its finite verb by adding a modal auxiliary like *may*, *might* or *could*. For example, the sentence 'John regrets that he went to the concert.' can be 'modalized' by turning it into 'It is possible that John regrets that he went to the concert.'

or into 'John may/might regret that he went to the concert.' GThe possibility test then amounts to this: In order that a candidate presupposition q for a sentence S passes the possibility test it must be the case that the modalizations of S entail q (and not only the corresponding moralizations of q). For example, the proposition that John went to the concert passes the possibility test for being a presupposition associated with the sentence 'John regrets that he went to the concert.' iff the modalizations of S – such as 'It is possible that John regrets that he went to the concert' or 'John may/might regret that he went to the concert' –, also entail that John went to the concert. (In other words, 'It is possible that John regrets that he went to the concert' and 'John may/might regret that he went to the concert' do not just that it is possible that John went to the concert (or that John might go to the concert). Note in this connection that if the relation between 'John regrets that he went to the concert.' and the proposition that he went to the concert were one of ordinary entailment, then the principles of modal logic would only predict that 'It is possible that John regrets that he went to the concert' entails that it is possible that John went to the concert. In general, if p entails q, then 'it is possible that p' entails 'it is possible that q', but not the plain, unmodalized q.)

But is the entailment of 'John went to the concert' by 'It is possible that John regrets that he went to the concert' an indication that the relation between the two is that of presupposition? The following consideration makes this plausible. When q is a presupposition of S, then when S is uttered in a context in which q fails, then the failure of q impedes the normal truth-conditional evaluation of this utterance, either because no proposition is being expressed (the Strawson cases) or because the utterance is defective in some other way. This impediment, however, is not restricted to evaluating whether the sentence is true in actuality. It also extends to evaluating the sentence for possible truth in various counterfactual worlds or situations. But evaluating an utterance of S for possible truth comes to the same as evaluating 'It is possible that S' (or any other possibility modalization of S) for truth in the actual world. So when q fails, a truth evaluation of 'It is possible that S' will be impeded no less than the truth evaluation of S itself. Conversely, when 'It is possible that S' is either asserted or denied, then that comes with the implication that q is supposed to hold.

The connections between the three tests for presuppositions of declarative sentences that we have discussed strongly suggest that they are tests for the same relation and that this is a relation that holds between a sentence S and a condition q iff q is a precondition for proper use of S. 'Presupposition' has been used standardly as the name for this relation. But even if there is a single relation that the term 'presupposition' denotes, that doesn't mean that there couldn't be different ways in which the relation cant manifest itself; there could be a range of connected but distinct relations that are all plausibly subsumed under the single overarching concept. One difference between different presupposition types is between those presuppositions whose failure prevents a sentence from expressing a proposition and those types whose failure makes use of the sentence inappropriate in other ways. (A striking example of this second type are the presuppositions triggered *again*, about which a good deal will be said in Section 4.2.) But this difference is just one from a range of distinctions between types of presuppositions. Another important distinction is that between presuppositions that are *anaphoric* (in a sense to be explained later on) that are triggered by definite noun phrases and also by certain tenses and presuppositions that are not anaphoric in this sense. That distinction will be crucial in what follows, since it is in the first instance for a better treatment of tenses and definites that we will need presuppositions in what remains.

From the perspective of intuitive motivation the 'Hey, wait a minute' test for questions appears to be the most immediately compelling of those we have discussed. When you are asked a question, you are put under extreme pressure to react to it in some appropriate way. The optimal way to react is of course to answer the question, in as complete and direct a way that you can. If you cannot answer the question, then the next best thing you can do is to give a reason why you can't. There can be various reasons for this. You may not know the answer, or you may not feel at liberty to divulge the answer to the asker, or you may refuse to answer the question because you feel that it is improper, invading your or somebody else's privacy. But you may also recognize the question as one that cannot be answered directly because a presupposition of it isn't fulfilled. This may either make the question unanswerable because no clear question is expressed by it; or else a direct answer, though possible, would imply your assent to the presupposition, but you consider the presupposition as false and do not want to go on record that consider it to be true.

When this is the reason why you are unable or do not want to give a direct answer, then you have to signal this to the lone who put the question to you. 'Hey, wait a minute' is one locution that you can use. It may not always be the one you will want to use. For one thing it is rather colloquial and it is also too straightforward to come across as really polite. Still, jot is good example of what we can and sometimes do say, when we want to convey that tho is the reason why a direct answer is not forthcoming.

We have seen how the other tests we discussed are in some sense derivative from the ‘Hey, wait a minute’ test for questions. The ‘Hey, wait a minute’ test for presuppositions of assertions (and also of other kinds of speech acts, such as promises, permissions or directives) is also a good an natural test, but not quite as compelling, because speech acts of these other types do not as a rule impose quite as much pressure on the addressee as questions. So it isn’t quite right that these cases of the ‘Hey, wait a minute’ test ‘derive’ form the ‘Hey, wait a minute’ test for questions; rather, they are other generally somewhat less compelling applications of the same test. When it is possible to turn any such speech act into a corresponding question, then applying the test to the question may be expected to provoke the stronger judgments.

The negation test and the possibility test, on the other hand, can be derived from the ‘Hey, wait a minute’ test, along the lines we argued. I will therefore refer to these tests as ‘indirect presupposition tests’.

The ‘Hey, wait a minute’ test for questions is one that can be extended to other types of speech acts, such as promises, permissions, directives – requests, pieces of advice, recipes, commands, orders (as in ‘mail order’) and so on – offers and more. For each of these speech act types there are corresponding questions. For instance, corresponding to the speech act type of promising there are questions like (4.5.a).

(4.5) a. Will you promise to send me an electronic copy of your book on semantics?

b. Have you published your book on semantics yet?

To such questions the ‘Hey, wait a minute’ test can be applied just as it can be to questions that correspond to assertions, such as (4.5.b). For instance, when you address me with either (4.5.a) or (4.5.b) and I haven’t neither written a book on semantics nor ever seriously considered writing one, then something like ‘Hey, wait a minute’ would be an appropriate reaction, and I would feel compelled to say something of the sort.

In the presupposition literature this generalization of the ‘Hey, wait a minute’ test has so far hardly been considered. The main reason for this is no doubt that there has been in formal semantics an almost exclusive focus on assertion – i.e. on declarative sentences as used for the purpose of making statements. This focus has included a great deal of work on questions corresponding to assertions. But it has largely ignored other speech act types and by implication, the questions that correspond to those types of speech acts. In these Notes we are towing this general line: We are also focused just on the semantics of declarative sentences. And almost all we are doing here ignores questions of any kind – questions have been mentioned in this and the previous subsection, but only because of the key role they play in connection with presupposition.

This focus on declarative sentences and statements will continue to be our guideline for the remainder of these Notes. An account of questions (of any kind) won't be given in what is still to come any more than it has been so far. Because of this there is little more we can do in conclusion of this subsection than to summarize our observations about presupposition tests at the same informal level in which we have discussed them. The summary lists the three tests that can be stated as tests involving certain sentence forms: the question test, the negation test and the possibility test. Note that in a way all these are 'Hey, wait a minute' tests. For in each case the presence of an unsatisfied presupposition manifests itself as the appropriateness of something like 'Hey, wait a minute' as response by an addressee who perceives the sentence form in question as coming with a presupposition that isn't satisfied. A difference between the first of these tests and the latter two is that the question test involves single items: questions that come with one or more presuppositions. Applications of the negation test and the possibility test always involve pairs of sentences: a sentence S and its negation or S and a sentences of the form 'It is possible that S'. And in these tests the point is that a 'Hey, wait a minute' reaction is appropriate in response to one sentence of the pair if and only it is appropriate in response to the other.

To conclude, here once more statements of the three form-related tests for presuppositions of declarative sentences to which the discussions in this subsection have led us.

(4.6) (presupposition tests)

a. **Negation Test**: p implied by both S and $\neg S$

'John regrets that he went to the party.' implies 'John went to the party' and

'John doesn't regret that he went to the party.' implies 'John went to the party'

b. **Question Test**: p implied by S?:

'Does John regret that he went to the party?' implies 'John went to the party'

c. **Possibility Test**: p implied by 'It is possible that S' etc:

'It is possible that John regrets that he went to the party.', 'John may/might regret that he went to the party.'etc all imply that 'John went to the party'.

4.0.6.4 Presupposition Tests and the Projection Problem

One aspect to the general phenomenon of presupposition that was destined to become the central focus of interest to linguists almost as soon as presupposition became an active linguistic concern, is the *projection behavior* of presuppositions. Suppose that q is a presupposition of a sentence S. Suppose that S is integrated as a constituent of some larger or more complex sentence S'. Will q be a presupposition of S'? it turns out that the answer to this question is not simple. it depends in part of what kind of sentence S' is and on what kind of constituent S is of S'. But it also depends on other factors. The challenge posed by projection is to state in terms that are both maximally general and precise, when projections of simple sentences project and when they don't.

The projection problem has close ties with the presupposition tests listed in (4.6), especially with the negation test and the possibility test. The negation test requires that a presupposition q of a sentence S also be a presupposition of the negation of S. (If both S and $\neg S$ imply q, then it is also the case that $\neg S$ and the negation of $\neg S$ both imply q, since the negation of $\neg S$ is equivalent to S.) So the negation test embodies the claim that presuppositions are preserved by negation. Likewise, the possibility test can be seen as embodying the claim that presuppositions are preserved by possibility modalization. These observation are partial answers to the general question that the projection problem poses.

To get an idea of what is involved in presupposition projection. consider the following two sets of examples in (4.7) and (4.8). In the first set the presupposition trigger is the plural definite description *John's/his children*, which presupposes that John has two or more children. The second set is a partial replay of this set, but now with the factive presupposition trigger *regret*.

(4.7) a. If John has children, then his children are bald.

b. If John is bald, then his children are bald.

c. If John's children are bald, then John is bald.

d. If John's children are bald, then he has children.

e. If John has sons, then at least some of his children are bald

f. If at least some of his children are bald, then John has sons.

g. If John doesn't have children, then his children are bald. (???)

h. If John didn't have children, then his children would have been bald. (???)

i. John has children and, moreover, his children are bald.

j. John is bald and, moreover, his children are bald.

k. John's children are bald and, moreover, he is bald.

l. John's children are bald and, moreover, he has children. (???)

(4.8) a. If John went to the party, then he regrets that he went to the party.

b. If John is bald, then he regrets that he went to the party.

c. If John regrets that he went to the party, then he went to the party.

d. If Mary didn't want to talk to John at the party, then he regrets that he went to the party.

e. If John didn't go to the party, then he regrets that he went to the party. (???)

f. John went to the party and he regrets it.

All sentences in (4.7) and (4.8) have the form of conditionals, in which the relevant presupposition trigger occurs either in the antecedent or in the consequent. Here are two morals that can be drawn from these examples:

(i) The presupposition or presuppositions generated by a presupposition trigger in the antecedent of a conditional always project (i.e. they are always understood as presuppositions of the conditional as a whole).

(ii) When a presupposition trigger occurs in the consequent of a conditional, the matter is more complicated. Consider the proposition trigger *John's/his children* and the presupposition that John has at least two children. For occurrences of the trigger within the consequent of a conditional the question whether the presupposition projects depends on further factors. For instance, in (4.7.b), the proposition hat John has two or more children is perceived as a presupposition of the entire conditional, but the proposition does not come across in this way in (4.7.a).

Note that this explanation of the difference between (4.7.a) and (4.7.b) is consistent with the observation that presuppositions of the antecedents of conditionals standardly project: As we noted in our discussion of the treatment of donkey pronouns DRT, the antecedent of a conditional is 'accessible' from its consequent, in the sense that it provides a local context which can support the interpretation of the consequent. But the converse does not hold. The consequent of a conditional does not provide a local context for its antecedent. This means that when q is a presupposition of the antecedent of a conditional, then there is no local context within the conditional that could verify q (by entailing it) and thereby prevent it from projecting.

It is not hard to see why there should be this difference between (4.7.a) and (4.7.b): the proposition that John has two or more children is entailed by the antecedent of (4.7.a). In fact, in this example the presupposition and the proposition expressed by the antecedent are one and the same proposition. But in (4.7.b) there is no such entailment; it is hard to see how – John's having children could have anything to do with his being bald.

Why should the entailment matter? Here is a story – it is the one that is now widely accepted and one that may seem obvious once you think of it: Presuppositions must be verifiable in the contexts in which their triggers are used. But the antecedent of a conditional can serve as a context for the conditional's consequent. So if a presupposition triggered in the consequent of a conditional is entailed by the antecedent, then this shows that the 'local' context provided by the antecedent verifies the presupposition; so there is no need for further contextual information to verify it. Such 'locally verified' presuppositions of a complex sentence do not impose any constraints on its use. Such constraints arise only when presuppositions cannot be verified sentence-internally; in that case the sentence can be used felicitously only in contexts that make verification possible.

This account of the difference between (4.7.a) and (4.7.b) also explains why a presupposition triggered in the antecedent of a conditional, as for instance in (4.7.c), always imposes a constraint on the felicitous use of the conditional. For when a presupposition is generated in a conditional's antecedent, there is no other part of the sentence that could provide a local verification for it – the antecedent of a conditional naturally qualifies as context for the interpretation for its consequent, but not the other way round. (For someone familiar with the dynamic approach that we have been pursuing this principle should hold no surprises. But in any case, we will come back to it below.)

The observation that presuppositions generated within the antecedents of conditionals always project can also be stated differently, in terms of presupposition preservation: when a sentence S is embedded as antecedent in a conditional S', then the presuppositions of S become presuppositions of S'. The concept of presupposition preservation is closely related to the indirect presupposition tests. Recall the negation test: when q is a presupposition of S, then q is also a presupposition of the negation of S (provided a suitable negation can be formed). Likewise, the possibility test can be formulated as the principle that presuppositions are preserved from sentences S to sentences like 'It is possible that S'. In the light of our observations in tho section about presupposition projection in conditionals we can now see more clearly why this should be so. It has to do with the fact that negation and 'possibilization' are 1-place operators. The presuppositions of the consequent of a conditional, we saw, sometimes project to become presuppositions of the conditional as a whole and sometimes they don't. And whether they do or don't depends on whether they are verified by the local context provided by the conditional's antecedent. It is the fact that conditional formation is a 2-place operation, which puts together two clauses in such a way that one can provide a context for the other. 1-place operators are different. They do not create a local context for the one clause or sentence on which they operate.

Here are some comments on the other sentences in (4.7) than (4.7.a,b). (4.7.d) may look like it refutes the generalization that was illustrated at the hand of (4.7.c), viz. that presuppositions generated in the antecedents of conditionals always project. (4.7.d) doesn't come across as presupposing that John has children. But if we attend closely to this sentence, we can see that there is something slightly peculiar about the message that it appears to be conveying. The sentence has the flavor of a tautology. and there is also something 'metalinguistic' about it: It seems to convey the necessary truth that if its antecedent is true (and felicitously used), then (of course) its consequent cannot fail to be true (and felicitously used) also. (In fact, it is the kind of sentence that seems to invite the addition of something like 'of course'.) If this is right, then (4.7.d) doesn't refute the general claim that presuppositions from the antecedents of conditionals always project. Rather, it shows an escape route from that claim: There are certain ways of reinterpreting the utterances of certain sentences that suspend the constraints which presuppositions impose on the felicitous use of those sentences. Cases where such presuppositional constraints are overruled we will refer to as cases of *presupposition cancelation*. It is important to keep in mind that presupposition cancelation is a different kind of mechanism than the local verification of presuppositions that we find for instance in (4.7.a). We will encounter other examples of cancelation below.

The contrast between (4.7.e) and (4.7.f) is on the face of it a straightforward illustration of our two morals: in (4.7.f) the presupposition that John has children is generated in the antecedent and projects; in (4.7.e) it is generated in the consequent, and it does not project because it is locally verified by the antecedent. But there is a twist that sets this last example apart from the earlier ones. First, (4.7.e) is a case in which the presupposition generated by *his children* in the consequent stands in a relation of asymmetric entailment to the antecedent: that John has sons entails that he has children, but not the other way round. Second, it may be felt that the presupposition in this example does project after all.

Second, for some speakers it is not as obvious as one the above remarks would seem to suggest that there is no projection in this example of the presupposition that John has children. There seems to be a pull towards reading this conditional as one that is most naturally used in a context in which it is given that John has children and in which the possibility is contemplated that some or all of those might be sons. In such a context the conditional could be used to claim that when its antecedent is true (i.e. when this possibility obtains), then some of his children (perhaps the male ones) are bald. If this is the reason, however, why (4.7.e) suggests use in contexts in which it is given that John has children, then it isn't much of a threat against the generalization expressed above under (ii). But it does show that the generalization needs a more circumspect formulation. There can be cases where the local entailment condition for a presupposition is fulfilled, but where there are independent reasons, having to do with the possible discourse functions for which the sentence seems earmarked because of its particular content, nevertheless favor contexts in which the presupposition in question is assumed. When a presupposition that is generated somewhere inside a sentence S is locally verified within S, then this creates the prima facie possibility that S can be used felicitously in contexts in which the presupposition does not hold. But this possibility may be foreclosed by other considerations that favor or require contexts in which the presupposition does hold.

(4.7.g) can be construed as a confirmation of the line on presupposition verification and projection that we are taking. This sentence seems irremediably weird, and our account makes it understandable why that should be. First, the antecedent of the conditional does not entail the presupposition that John has children. Therefore the conditional can only be used felicitously in a context of which the information that John has children is already part. But if that is so, then the antecedent of the conditional contradicts the context in which it is made, and thus describes a state of affairs that is impossible within the limits set by the context. This incompatibility is in conflict with the indicative mood of the conditional (realized through the use of the Simple Present tense in both its antecedent and its consequent) which implies that its antecedent is compatible with the context. So the conclusion is that (4.7.g) cannot be used felicitously in any context. On the one hand a suitable context would have to be one in which the presupposition that John has children is entailed. But on the other hand, if the context satisfies this condition , then the antecedent of the conditional will be incompatible with it.

(4.7.h) shows that this cannot be the whole story, however. Intuitively this sentence is no better than (4.7.g). At first sight this could seem surprising, for one part of the explanation given in the last paragraph was that the antecedent of (4.7.g) must be compatible with the context in which the conditional is used, in virtue of its indicative mood. The mood of (4.7.h) is subjunctive and its interpretation could be counterfactual; that is, its antecedent might be interpreted as false and more particularly as contradicting the context in which the conditional is uttered. So the compatibility argument does not apply in this case. It is not hard to see, however, that this difference between (4.7.g) and (4.7.h) is not enough to salvage the latter. The verifiability requirement for the presupposition that John has children applies to (4.7.h) just as it does to (4.7.g). Since the presupposition cannot be verified by the antecedent of (4.7.h) (any more than it can be verified by the antecedent of (4.7.g)), the verification will have to come from the global context. But how could it be, even in this case? Admittedly the present case is somewhat more complicated. Suppose that (4.7.h) is used in a context C. It is widely accepted that evaluation of a conditional 'If it were/had been the case that A, then it would have been the case that B' in a context C involves modifying C, in a minimal sort of way, so that it makes A true and then to evaluate B in this second context. If B is true in this modified context C', then that verifies the conditional as a whole in C; if B is false in C', then the conditional is false in C.[4]

[4]The idea that a counterfactual conditional 'If it were/had been the case that A, then it

On such an account of the semantics of counterfactuals (4.7.h) comes out as necessarily incoherent. If (4.7.h) is counterfactual in C, then that means that its antecedent is false in C, i.e. that John has children according to C. Let C' be a minimal modification of C according to which John does not have children. To find out whether (4.7.h) is true in C we have to evaluate its consequent in C'. But evaluation of the consequent in any context is possible only when its presuppositions are entailed by that context. So in order that evaluation of the consequent of (4.7.h) in C' is possible C' must entail it presupposition that John has children. But we just saw that C' is a context which entails that John does not have children. So C' must entail contradictory information. There cannot be such contexts. So (4.7.h) cannot be coherently evaluated in any context. That accounts for why (4.7.h) strikes as as weird and as no better than (4.7.g).

The last four examples of (4.7) show that as far as presupposition verification is concerned, conjunctions are similar to conditionals, with the first conjunct of a conjunction playing the part played by the antecedents of conditionals and the second conjunct that which is played by their consequents: presuppositions generated in the first conjunct project, presuppositions in the second conjunct project unless verified locally, by the first conjunct. Thus, in analogy with what we have just seen for the case of the conditionals, the presupposition that John has children projects in (4.7.j) and (4.7.k), but not in (4.7.i). Like (4.7.d), (4.7.m) is an odd sentence. And here the oddity is if anything even more extreme: Since the first conjunct presupposes that John has children and since that presupposition necessarily projects, (4.7.l) can be uttered felicitously only in a context that entails this presupposition. But then its second conjunct is supposed to make a contribution that will be known to be part of the context (assuming that (4.7.l) is used felicitously in it) at the point when the second conjunct is to be interpreted. This is in clear violation of the 'Principle of Informativeness': do not present information that is already familiar as if it were new information.

The sentences in (4.8) have been given in order to show that the aspects of presupposition generation, verification and projection illustrated by the

would have been the case that B' is true in a given context C if the consequent B is true in a minimal modification C' of which accommodates A can be found in the literature in various forms. In philosophy the by now most familiar form is the theory of conditionals developed in closely similar wye by Lewis ((Lewis 1973)), Sobel ((Johansson, Österberg, Sliwinski & Sobel 2009)) and Stalnaker ((Stalnaker 1968)). But we find it in other guises as well, for instance in theories of belief revision ((Alchourron, Gärdenfors & Makinson 1985)).

examples in (4.7) are not specific to the presuppositions triggered by definite descriptions. Presuppositions triggered by factive verbs like *regret* show much the same behavior. Thus (4.8.a), (4.8.b) and (4.8.c) are direct analogues of (4.7.a), (4.7.b) and (4.7.d). (4.8.d) is a case of local verification in which the presupposition is 'asymmetrically' entailed by the antecedent (i.e. the antecedent is not entailed by the presupposition). (4.8.e) is incoherent in the same way and for the same reason as (4.7.g). And (4.8.f) is another case of local verification of a presupposition generated in the second conjunct of a conjunction by the conjunction's first conjunct.

4.0.7 The Dynamics of Presupposition Management

The sketch we have given for dealing with presupposition projection rely crucially on the notions of 'local' and 'global' context: When presuppositions generated in embedded clauses or sentences do not project – i.e. if their verification is not the task of the global context in which the entire sentence is uttered – this is usually because there is a local context, established by some other part of the sentence, that takes care of their verification. In a systematic reconstruction of how projection works we need to have a systematic way of identifying global and local contexts and the structural relations between them (which among other things determine which is the global context and which are the local ones). Importantly, the structure we need to be able to track is a dynamic one, which evolves in the course of discourse interpretation, as more and more content is extracted from the successive sentences of which the discourse is made up. In the first instance this evolution concerns the global context, in which more information gets accumulated as more and more sentences of the discourse get processed. But the growth of the global context also affects the new local contexts that come into play when the interpreter encounters logically complex sentences. For as we will see presently, these local contexts interact with the global context, in the sense that local presupposition verification will in general be a matter of verification by local and global context combined. So the contexts that are 'locally' available for presupposition verification are in fact always combinations of the local context (in the narrow sense of this term in which we have been using it) and the 'global' context, in which we take the complex sentence as a whole to have been used. (For instance, the context locally available for the verification of presuppositions triggered in the consequent of a conditional consists of the content of the antecedent – the local context in t narrow sense – in combination with the context in which the conditional is taken to have been used (the global context).)[5] So the locally available context is affected by the dynamics of discourse interpretation for two reasons: first, the global context unfolds as more of the discourse is processed; and second, because the interpretation of a sentence part that that goes into the locally available context may be affected by the global context that it makes use of: on the basis of the initial global context (or an earlier one in the chain of successive global contexts that become available as the discourse processing goes on) the interpretation might have been different, and with that the local context in the narrow sense.

Often global and local contexts can, we just stated, cooperate in the verification of a presupposition. This is shown by the following variant of examples discussed by Kripke's (see in particular (Kripke 2009)).[6]

(4.9) We shouldn't forget that Mary's birthday is before Bill's. If we are having pizza on Mary's birthday, then we are not going to have pizza again on Bill's birthday.

Intuitively it seems plain that in this example the local context established by the antecedent of the second sentence serves to justify the *again*-presupposition generated within the consequent – the proposition that there will have been an occasion when the people involved had pizza that preceded Bill's birthday. But strictly speaking the verification of this presupposition by what is said in the antecedent is possible only on the assumption that birthday comes after Mary's. That information is provided by the first sentence. If we assume that this information is part of the global context for the second sentence of (4.9) – something that may be plausible enough as things stand, but that will become even clearer as we go along – then the verification of the *again*-presupposition in this case depends on a combination of information,

[5] The locally available context can be even more complex, when the sentence part whose presuppositions are to be verified are more deeply embedded, for instance when the part is the consequent of a condition that is itself the consequent of another conditional. Here the locally available context is composed of (i) the global context, (ii) the antecedent of the outer conditional, (iii) the consequent of the outer conditional and (iv) the consequent of the embedded conditional. For now we set these complications aside.

[6] I am strongly convinced that these very examples (or close variants of them) are due to Kripke. They do not appear as such in (Kripke 2009), the eventual publication of a lecture that he delivered in the early nineties and of which an unofficial transcript had been in wide circulation since that time. But in any case, the methodological significance of these examples is very similar to that of many of the examples discussed in Kripke's paper. What will be said about these examples here does not agree in every respect with the analyses that Kripke offers for similar examples; but that is another matter.

some from the global and some from the local context.

I trust that even without further examples it will be plausible that this is only the tip of an iceberg. Sentences can have much more complex logical structures than simple conditionals. Such sentences can give rise to structured hierarchies of local contexts and in such cases presupposition verification can depend on several 'nested' local contexts as well as on the global context. A further complication has to do with the availability of the information from 'higher' contexts at the level of a 'lower' context. In what we said about (4.9) we assumed that information from the global context is available as support for the information of the local context. And in this case, where the local context is established by the antecedent of an indicative conditional, that is generally correct. But we have already seen, in our discussion of (4.7.h), that this need not be so in general. The antecedent of a counterfactual conditional can overwrite all or part of the global context, and so a further complication arises in such cases: How to determine which information from the global context is available for verification of a presupposition generated in the consequent. This problem arises whenever we turn to from talk about what is the case here and now to talk about other times or other worlds. The counterfactual is just one such construction, from a large and thus far only partially charted variety. In what follows this complication will be ignored.

Before we continue our explorations of context structure, let us briefly return to the historical dimension of this overview. In Section 4.1.2 we noted that the understanding of what kind of phenomenon presupposition is, and of the range of its manifestations, changed radically when linguists observed that the issue wasn't limited to definite descriptions: the tests suggesting that unique satisfaction is a presupposition of the proper use of definite descriptions apply to many other cases as well – see the list in (4.3). With the discoveries of new presupposition triggers came a more systematic reflection on tests for presuppositionality (cf. (4.6)), and as we have seen, these tests lead more or less directly into the problems of presupposition projection. As we noted earlier, many linguists whose interest in presupposition phenomena had been roused by the discovery of how diverse and ubiquitous those phenomena are soon came to see presupposition projection as the central challenge. Three contributions from those days stand out: (Langendoen & Savin 1971) explicitly recognized the central importance of the notion of projection for an account of presupposition, (Karttunen 1973) recognized the complexity of projection phenomena, and more specifically that the projection from embedded positions (such as that of the consequents of conditionals) depends on additional factors, and offered a machinery (which we will not discuss here) for dealing with these complexities. Then, in the following year, (Karttunen 1974) offered a revised account of those complexities, which in essence is the one we have rehearsed here: Parts of sentences can serve as local contexts for the verification of presuppositions generated in other parts. And that, we have seen, also introduces the dynamic dimension into presupposition theory. In retrospect it is clear, I think, that Karttunen's second paper was the seminal paper – the principal source of the mainstream of presupposition theory as we know it today.

Combining these last observations with those about Stalnaker's work in Section 4.1.1 leads to the following picture of the state of play by the middle of the nineteen seventies. On the one hand the so-called 'semantic' approach, which found its first crystallization in the work of Karttunen and on the other the 'pragmatic' viewpoint represented most prominently in Stalnaker's work, according to which presuppositions are assumptions made by speakers, which reflect what information they assume to be available to their addressees when these interpret what they are saying and on which she relies in choosing the words for what she wants to say. These two perspectives should not be seen as mutually exclusive alternatives, but as focusing on different aspects of what presuppositions are and how presupposition works. The two perspectives meet there where speakers make use of expressions that carry presuppositions by virtue of what they mean and how they function in the language. When a speaker makes use of such an expression, she is committed to the linguistic presupposition or presuppositions generated by that expression. In other words she is, as a competent and responsible speaker of her language, obliged to include among the presuppositions she is making in connection with her utterance all those that are carried by the expressions she is using. This dual picture of presupposition, as something that speakers do, but that in some cases they cannot help doing because the language forces them, is the one we also adopt in these Notes. (I hasten to add that in what follows here this duality will not be visible. For the presuppositions that we will consider will all be linguistic presuppositions, triggered by certain lexical items. Vis-a-vis these presuppositions of this kind the speaker has no choice – she has a commitment to include them among her pragmatic presuppositions Non-linguistic presuppositions won't be considered.)

Before turning to those applications let me briefly mention the three principal links in the chain that connects the treatment of presuppositions I will present here with Karttunen's seminal 1974 paper. (Heim 1983) formulates central aspects of presupposition projection in an overtly dynamic setting. Heim's formulation states in admirably concise terms the net effects of local

verification on the truth conditions of complex sentences with presupposition triggers. Heim's account was conceived in conjunction with the File Change semantics developed in (Heim 1982,1988), one of the first versions of dynamic semantics, and with close affinities to DRT. (FCS and DRT were developed independently and at roughly the same time – on my understanding an early version of FCS anteceded the first explicit version of DRT, as it appeared in (Kamp 1981*b*).) The formulation of presupposition projection in (Heim 1983), however, abstracts away from details in which the alternative version of dynamic semantics differ from each other. This is one of its virtues. But on the other hand it entails that a good deal of work is left to be done when one want to 'localize' the account to any one of these versions. In particular, turning Heim's insights into part of an explicit formulation of DRT requires additional ideas as well as a good deal of further labor.
The next crucial step in this direction was taken in (Van Der Sandt 1992) and (Van Der Sandt & Geurts 1991). These papers show how explicitly representing the presuppositions generated by presupposition triggers in a sentence, as part of a DRS-like sentence representation, gives us a formal handle on the question when these presuppositions project and when they don't, and also on the possibilities of *presupposition accommodation*. An important source of inspiration for the proposals of Van Der Sandt and Geurts was the observation (already implicit in (Heim 1983)) that the principles which govern what counts as a local context for the verification of a presupposition are the very same as those that in DRT determine which discourse referents are available as anaphoric antecedents for a given third person pronoun – accessibility of local context information available for presupposition verification is determined by the same configurational principles as accessibility of pronominal antecedents in DRT. In fact, the slogan that has often been employed to characterize the central contribution of this work is that there is no significant difference between presupposition and anaphora, or even that, as the title 'Presupposition Projection as Anaphora Resoluton' of (Van Der Sandt 1992) implies, presupposition is a kind of anaphora. In what follows we will diverge from this aspect of Van Der Sandt's account: anaphora will be treated as a special kind of presupposition (and not the other way round) and the representations of anaphoric presuppositions triggered by third person pronouns and other anaphoric sentence elements will be distinct from the presuppositions that we have been discussing so far in Section 4. Of the important contributions to presupposition theory that have further developed the approach first presented in (Van Der Sandt 1992) and (Van Der Sandt & Geurts 1991) the one that should be mentioned first and foremost is Geurts' monograph (Geurts 1999).

A second aspect of the Van Der Sandt-Geurts theory that we will not follow is what Van Der Sandt has to say in (Van Der Sandt 1992) about the mechanisms of what we will from now on refer to as *presupposition justification*. In order that an utterance can be felicitous every one of its presuppositions must be justified. But justification can take different forms. One of these is presupposition verification, the only kind of presupposition justification that has been considered in the examples we have been discussing so far. However, utterances with presuppositions that cannot be justified through straight verification may still be saved from uninterpretability through accommodation. Accommodation is something that interpreters resort to when canonical strategies for presupposition verification fail. It involves some form of adjustment of the information that the interpreter takes to be available for the verification of one or more presuppositions, in such a way that that verification is now possible. There appears to be wide agreement among presupposition theorists that accommodation plays an important part in presupposition management. But there is far less agreement over how accommodation works, or what the principles are that govern it. The view of accommodation we will adopt here is one that is most clearly articulated in the work of Beaver. (See in particular (Beaver 1992) and (Beaver 2001) as well as the eminently useful survey article (Beaver 1997) for a much more detailed and broader overview of the subject of presupposition than the bird's eye view presented here.) According to Beaver, accommodation is what the interpreter of a sentence with presuppositions can do when he is unable to verify one or more of its presuppositions on the basis of the contextual information available to him: Assuming that the speaker will be aware of the presuppositions to which she is committed by the choice of her words, and that she has therefore made sure that all those presuppositions can be verified on the basis of what she takes the context to be, the addressee will take his inability to verify one or more of these presuppositions in what up to this point he assumed to be the utterance context as an indication that his conception of the context must differ from the speaker's, and adjust his own conception in such a way that the verification of these presuppositions now becomes possible.

Some discussions of accommodation often create the impression that it is a whole-sale alternative to verification: if a presupposition cannot be verified in the context as is, then, as a last resort, the presupposition can be accommodated by adding it as a new proposition. But as we will see later on, that is an adequate picture only in some cases, and those are not necessarily typical. Beaver notes that presupposition accommodation is an attempt on the part of the interpreter to adjust his conception of the context to what he can plausibly assume to be the context that the speaker has in mind, given that that context must verify the presuppositions in question. Sometimes this adjustment will amount simply to the addition of the presupposition itself. But often something stronger than this is added to the context, so that the new context asymmetrically entails the one that would have been obtained by just adding the presupposition itself. For after all, what an accommodating interpreter does according to Beaver's perspective is engage in a kind of abduction: find the most plausible explanation for why the speaker took the presupposition to be satisfied. All sorts of considerations, about what information could have been available to the speaker, or about her general ways of thinking and prejudices, may be involved in these abductions. As a result, the contextual assumptions that the interpreter will attribute to the speaker in his effort to account for why she takes her presuppositions to be true may well contain information that exceeds what is minimally necessary for entailing the presupposition or presuppositions at issue.

Presupposition accommodation may also diverge from the simple notion of presupposition addition in the opposite direction. When presuppositions are generated in embedded positions, their local contexts may provide some of the information that is needed for their verification, but not all of it. In such cases accommodation will often take the form of adding to the global context just the information that yields the desired verification when combined with the local context. In such cases the accommodated information may be less than the presupposition, not more.

To conclude, the information that gets accommodated for the sake of verifying a presupposition can be either more or less than the presupposition itself. The only general constraint is that the presupposition is entailed when the accommodated information is combined with the global context, and with the relevant local contexts, in those cases where local contexts are relevant.

It is a default of linguistic communication that presuppositions are verifiable without recourse to accommodation. But nevertheless the need for accommodation is common enough. Here are two examples which may give a flavor of how common accommodation is and perhaps also a certain sense of why this could be. The first example is an instance of the often observed fact that definite descriptions can be used to communicate the fact that something of the kind described exists. Suppose that I ask you if we can meet on Friday at two and you reply with the sentence in (4.10).

(4.10) No that won't work. I have to fetch the boy friend of my daughter from the airport.

Suppose that I didn't know you had a daughter, or that I knew that but not that she had a boy friend. Then I lack the means to verify one or both of the presuppositions carried by the definite description *the girl friend of my son*. But this will hardly bother me. Perhaps I will quietly say to myself: 'I didn't know he had a daughter', or 'I didn't know that his daughter had a boy friend' and adopt this as part of the information your utterance conveys to me – that you can't make two o'clock on Friday and that fetching your daughter's boy friend is the reason. Accommodation more or less comes for free here, and you may have counted on this if you knew or had good reasons to assume that I didn't know you had a daughter or that she had a boy friend, in which case I get that information too out of what you are telling me.

The second example is a variant of (4.9). (This variant, given in (4.11), is actually closer to Kripke's original examples than (4.9). (4.11) differs from (4.9) only in that the first sentence is missing.

(4.11) If we are having pizza on Mary's birthday, then we are not going to have pizza again on Bill's birthday.

When discussing (4.9) we noted that the information that Bill's birthday will come after Mary's birthday is needed to verify the presupposition that 'we' had pizza before Bill's birthday. In (4.11) this information is not explicitly present. But when the second sentence of (4.9) is presented out of the blue to you, as it is in (4.11), then you are likely to accommodate this information, perhaps even unthinkingly. It almost feels like this accommodation is forced upon you, since it appears to be the only way in which you can make sense of the way (4.11) is worded; it is almost as if (4.11) entails that Bill's birthday is later than Mary's.

Especially the first of these two illustrations suggests that often interpreters aren't even aware that their interpretations involve accommodation or at least that they pay no particular attention to the fact that they accommodate, and also that speakers can exploit this by using presupposition-carrying expressions which force their interlocutors to accommodate the presuppositions carried, thereby adopting them as new information. And because accommodation often comes as easily as it does, distinguishing between presupposition verification without and presupposition verification with accommodation is often hard, and may sometime be impossible. Suppose that the content p is needed to verify the presupposition q that is generated by a sentence S that has been uttered in some context C and that the interpreter makes use of p in his verification of q. Is p information that is genuinely new to him, and that he obtained through abductive reasoning? Or might p be information that he had acquired before but that lay buried in his memory and that the presupposition is now reminding him of?

However, even if the dividing line between verification with and verification without accommodation may often be hard to draw, the distinction is crucial. It is for the sake of making this distinction explicit that I have introduced the pair of terms 'presupposition verification' and 'presupposition justification'. Presupposition verification refers to the process which establishes that the presupposition is entailed by the context as the inter peter assumes it to be. Presuppositon justification is a more inclusive concept, which not only covers cases of presupposition verification (in which accommodation is not involved) but also processes in which accommodation does play a part. Presupposition justification is thus a task that faces any interpreter of any sentence when dealing with a presupposition triggered by some constituent of that sentence.

We have reached the point where we can start with the integration of our own account of presuppositional phenomena into the DRT-based semantics that we are developing in these notes. But one last general remark before we start on the real work. Nearly all formal studies on presupposition adopt a methodology that has long been standard practice within linguistic pragmatics: What propositions are expressed by which sentences is taken for granted. The part of the interpretation process in which this is determined is supposed to have taken its course at the point where pragmatic mechanisms come into play; and these operations operate only on the propositional contents that have been assigned, not on the sentences and clauses themselves. Furthermore, the assignment of propositional contents to sentences and clauses is regarded as the province of 'semantics', another component of the over-all theory of language interpretation. Pragmatic accounts of presupposition phenomena have been in keeping with this general perspective insofar as they assume that the presuppositions one is dealing with are given as propositions, which have already been identified, presumably as part of the 'semantic' processes for which pragmatics bears no responsibility.

In a general set-up like the one we have been pursuing this modus operandi isn't acceptable – doing so would be incompatible with our commitment to making the form-meaning relation as explicit as we are trying to do throughout these Notes and which has guided our efforts in Section 3. What we need in order to maintain that commitment is to spell out in detail how presupposition triggers determine semantic representations of the presuppositions to which they give rise, in the various syntactic environments in which they occur. We need, that is, as a kind of preamble to the questions of presupposition justification on which so much of the presupposition literature has focused, also an account of the rules according to which the representations of presuppositions are constructed, just as we have been concerned in Section 3 with the construction of the representations of non-presuppositional content. We will see that this 'preliminary' part of an account of linguistically generated presuppositions is not without its challenges. One reason why the construction of the representations of presuppositions is a non-trivial matter is that it includes all that is needed to account for the construction of non-presuppositional representations. One reason for this is that among the presupposition triggers we will deal with is the definite article *the*. The presuppositions triggered by occurrences of *the* obviously depend on, and thus vary in accordance with, the nominal phrases with which *the* combines into the definite descriptions with which those presuppositions are associated. Often these nominal phrases contain relative clauses and these relative clauses can be syntactically as complex as you like; any challenge to the construction of non-presuppositional representations can thus be reproduced as a challenge for the representation construction of presuppositions. And *the* is by no means the only trigger that gives rise to this general problem. Factive verbs are perhaps the most obvious example, since there are evidently no limits to the variety and complexity of their sentential complements. Other examples are presupposition-triggerng words like *again*, *still*, and the 'additive particles' *also*, *too*, *as well*. We will encounter some of these representation-related issues in an exploration of presuppositions generated by *again* in the next could of sections. There we will see that the construction of presupposition representations is a major part of presupposition management. In fact, in what follows much of the emphasis will be on the representation construction for presuppositions, rather than on the justification processes that come into action only when the construction process has been completed.

In what is coming we will develop a construction algorithm that constructs representations for sentences with presupposition triggers which will include representations for all the presuppositions these presupposition triggers generate in their respective syntactic locations. Furthermore, these presupposition representations are inserted locally, in positions that are left-adjoined to the representations of the non-presuppositional contents of the clauses that contain their triggers. The resulting sentence representations will be called *preliminary representations* (or *preliminary DRSs*) ((Van Der Sandt 1992)). A preliminary representation is not yet the final representation of the input (the sentence, utterance, discourse or text for which a representation is being constructed). Rather, it is itself the input to a further interpretation module, the task of which is to justify the various presuppositions represented within the preliminary DRS. When and only when this module succeeds will the final representation have been reached. This final representation will then have the form of a DRS, of the sorts that we have been constructing in Section 3. The construction procedure we will develop in this section is thus going to be two-stage, in a way that the construction procedure of Section 3 was not: upon the construction of the preliminary representation, which may be thought of as part of the syntax-semantics interface in the strict sense, follows the justification of presuppositions, which can be thought of as part of linguistic pragmatics. Note well that it is only at this point that contexts (and thus, more particularly, when we are dealing with multi-sentence texts and conversations, the discourse contexts) are coming into play.

The two-stage procedure of constructing semantic representations has an important methodological implication. Unlike DRT in its original form, and also unlike most other versions of dynamic semantics, the treatment we will present restores a sharply defined division between semantics and pragmatics: preliminary DRS construction belongs to semantics (and for us here this is all that semantics, in the strict formal sense of the word, consists of), and conversion of preliminary DRSs into DRSs via presupposition justification belongs to pragmatics.

4.1 Presuppositional phenomena as we treat them here

4.1.0.1 Computation and Justification of Presuppositions triggered by *again*

Consider example (4.12).

(4.12)It rained yesterday. It rained again today.

In Section 3 we discussed almost all that is needed to construct the DRS for the first sentence of (4.12). The one thing that wasn't discussed explicitly is the lexical semantics of the verb *rain*. We treat *rain* as a '0-place verb', i.e.

as a verb which has a referential argument but no non-referential arguments. Thus the semantics of its lexical entry will have the form in (4.13).

$$(4.13) \qquad < e \mid \begin{array}{|c|}\hline \quad \\ \hline \text{e: rain'} \\ \hline \end{array} >.$$

But otherwise all the pieces for the construction of this DRS are in place, so its construction is left as an exercise. The resulting DRS is displayed in (4.14).

$$(4.14) \begin{array}{|c|}\hline e \quad t \quad d \quad d' \\ \hline t \prec n \quad e \subseteq t \\ \text{day}(d) \quad \text{day}(d') \quad n \subseteq d' \quad d \supset\subset d' \quad e \subseteq d \\ e : \text{rain'} \\ \hline \end{array}$$

(4.14) will serve as discourse context for the DRS construction of the second sentence of (4.12). However, the point at which and the way in which the discourse context comes into play will from now on differ from what we have seen in PART I. There we introduced the LF for the new sentence directly into the context DRS and the construction operations applying to it stepwise integrated pieces from the LF into this DRS, in the form of new drefs and DRS conditions. From now on we split this process into two distinct stages, (i) the construction of a *preliminary representation* for the new sentence, a construction in which the discourse context does not yet play any part; this preliminary representation will contain representations of the presuppositions generated by presupposition triggers from the new sentence. In stage (ii) these presuppositions will be 'resolved' using the information provided by the discourse context, after which the now presupposition-free representation for the new sentence will be put together with the context DRS.[7]

In order to get started we need an LF for the second sentence of (4.12). There are three things about the syntax of this sentence that we haven't yet dealt with. The first is the sentence-final adverb *today*. We assume that this adverb is adjoined to TP, just as the sentence-initial adverbs to which our attention has been confined up to now. The second is the dummy subject *it* of the verb *rain*. The fact that *it* doesn't occupy an argument position

[7]For discussion and references see Section 4.1.5.

in the semantics of the verb becomes manifest upon lexical insertion for the verb, in that there will be no coindexation at this point between the subject DP and a non-referential argument position. (There could not be such a coindexation, since the verb *rain* has no non-referential argument positions; see (4.13).) The absence of an index on the DP will serve as indication that it is to be simply ignored at the point where DP and T' are brought together. The third is the adverb *again*. As we will see below, the syntactic position of *again* is a delicate matter. We assume, for reasons that will become clearer as we go along, that *again* is treated as a right adjunct to VP in this example. (4.15) gives the LF after lexical insertion for the verb.

(4.15)

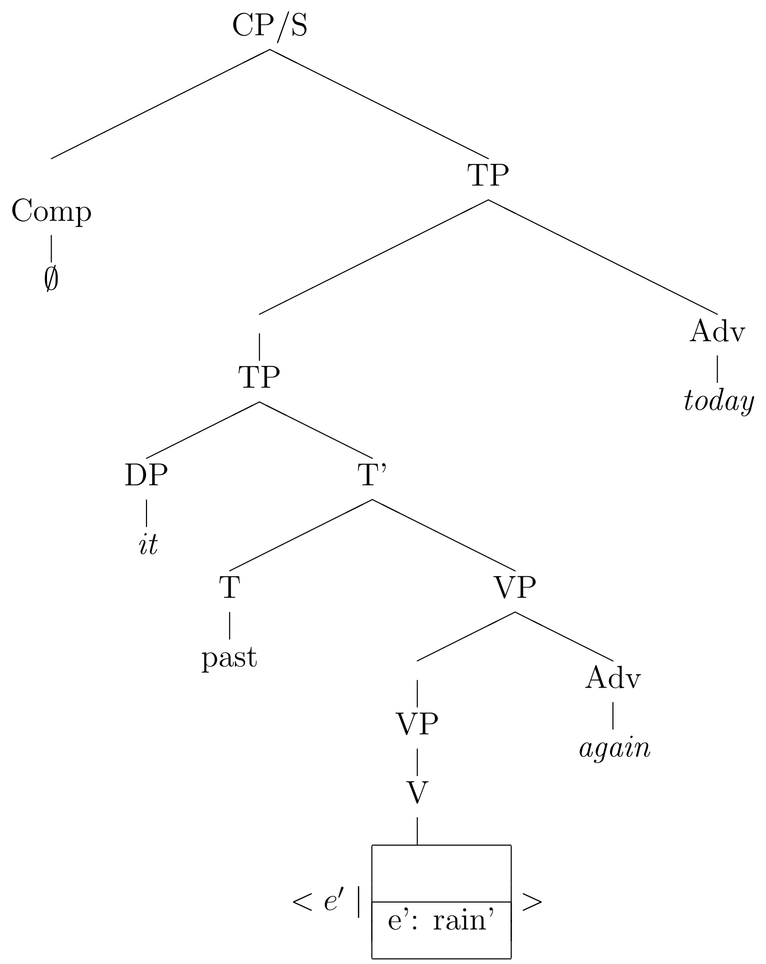

The first of the next construction steps – and the one that matters most right now – is that which combines the lower VP representation (which is identical with that of the V-node) with the semantics of *again*. This requires a lexical entry for *again*. In dealing with this and the next few examples I will just show the result of combining this entry with the semantic representation of *again*'s adjunction site. A lexical entry for *again* will be presented and discussed in Section 4.2.4.

(4.16)

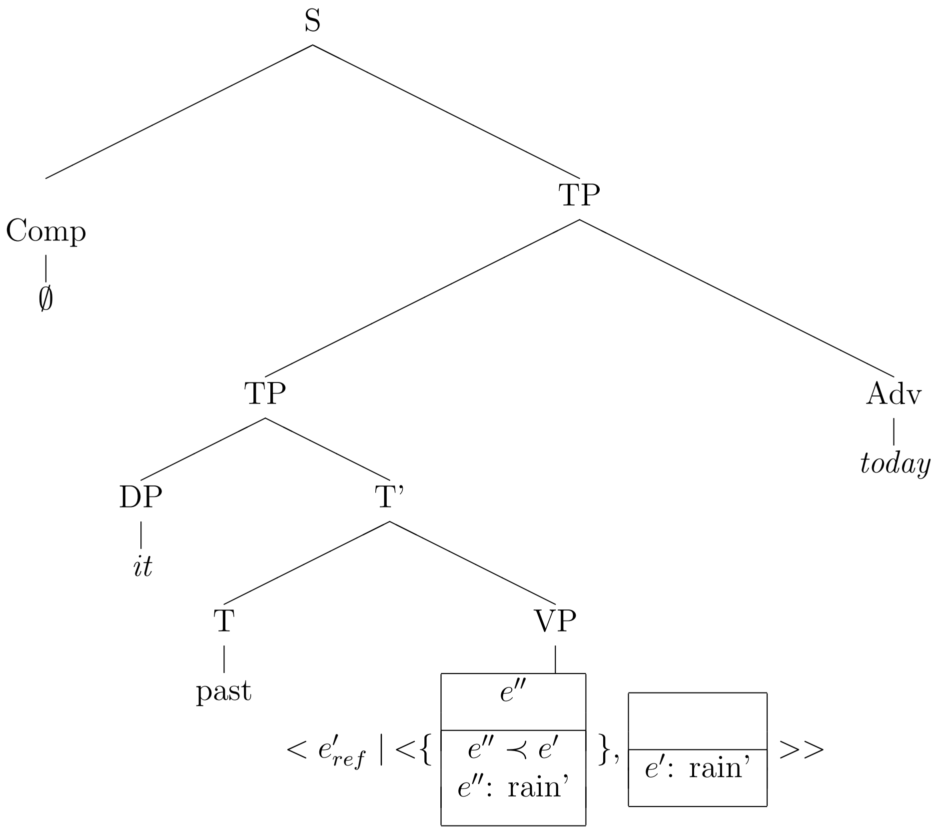

Note the form of the semantic representation of the upper VP node in (4.16). It consists of a store (with as its only element the referential argument e') followed by an ordered pair consisting of (i) a non-empty presupposition set (in this case a singleton set) and(ii) the non-presuppositional representation.

The remaining construction steps that are needed to obtain the upper TP-representation, which incorporate the semantic contributions made by the tense feature past and the adverb *today*, are covered by established construction principles. But there is one new twist, which also applies to other construction operations that we have encountered in Section 3. All these operations are to be performed on the non-presuppositional part of the representation only – they leave the presupposition set invariant. (This is the formal reflection of the idea that when a presupposition has already been triggered and represented before such an operation is executed, its representation will establish how it is related to the non-presuppositional content, even if the representation of that content is not yet completed. Subsequent compositional operations that develop the non-presuppositional content further should leave the representations of the presuppositions untouched.

Applying the operations required to reach the representation of the upper TP node in this spirit leads to the structure shown in (4.17).

(4.17)

S

Comp　　　　TP

$\emptyset$

$$< e'_{ref}, t', d'' \mid <\{ \begin{array}{|c|}\hline e'' \\ \hline e'' \prec e' \\ e''\text{: rain'} \\ \hline\end{array} \}, \begin{array}{|c|}\hline \ \\ \hline t' \prec n \quad e' \subseteq t' \\ \text{day}(d'') \quad n \subseteq d'' \quad e' \subseteq d'' \\ e'\text{: rain'} \\ \hline \end{array} >>$$

The one remaining operation is the transfer of the drefs in the store to a suitable DRS Universe. Here the principle mentioned above applies once more: the operation leaves the presuppositions untouched; that is, the drefs are to be transferred to the Universe of the non-presupposition DRS. This leads to the CP representation shown in (4.18).

(4.18)
$$<\{ \begin{array}{|c|}\hline e'' \\ \hline e'' \prec e' \\ e''\text{: rain'} \\ \hline\end{array} \}, \begin{array}{|c|}\hline e' \ t' \ d'' \\ \hline t' \prec n \quad e' \subseteq t' \\ \text{day}(d'') \quad n \subseteq d'' \quad e' \subseteq d'' \\ e'\text{: rain'} \\ \hline \end{array} >$$

(4.12) is an example in which the presupposition of the second sentence is justified because it is verified by the discourse context that is provided by the first sentence; and that is so, intuitively, because the discourse context entails the presupposition. But note that strictly speaking this last claim isn't true as stated, when we identify the discourse context for the second sentence with the DRS in (4.14), which represents the first sentence. The presupposition representation in (4.18) isn't entailed by that DRS on its own, but only (a) in conjunction with information from the non-presuppositional DRS of (4.18) (viz. information to the effect that the event represented by the dref e' (which has a free occurrence in the presupposition DRS) is one occurring today) and (b) given certain general assumptions about the temporal relations $\prec$, $\subseteq$ and $\supset\subset$ and the calendar predicate 'day' (viz. that an event which happened yesterday preceded an event that happened today). (4.19) gives an explicit statement of this entailment relation. It makes all information from the non-presuppositional DRS of (4.18) available to the context from which the *again*-presupposition must follow in order that it count as resolved. This is achieved by merging this DRS with the context-DRS (4.14) provided by the first sentence of (4.12). This DRS merge gives us the premise DRS K_{pr} of the entailment relation stated in (4.19). The conclusion DRS K_{con} of this relation is the representation of the *again*-presupposition that occurs as sole member of the presupposition set in (4.18).) $\models$ is the entailment relation, which relates the DRS'K_{pr} in (4.19) as premise to the 'conclusion DRS' K_{con}.

(4.19)
$$\begin{array}{|c|}\hline e \;\; t \;\; d \;\; d' \;\; e' \;\; t' \;\; d'' \\ \hline t \prec n \quad e \subseteq t \\ \text{day}(d) \quad \text{day}(d') \quad n \subseteq d' \quad d \supset\subset d' \quad e \subseteq d \\ e\text{: rain'} \\ t' \prec n \quad e' \subseteq t' \quad \text{day}(d'') \quad n \subseteq d'' \quad e' \subseteq d'' \\ e'\text{: rain'} \\ \hline \end{array} \quad \models \quad \begin{array}{|c|}\hline e'' \\ \hline e'' \prec e' \\ e''\text{: rain'} \\ \hline \end{array}$$

$$\Uparrow \qquad\qquad\qquad \Uparrow$$

$$K_{pr} \qquad\qquad\qquad K_{con}$$

Two points are worth observing about (4.19). The first is that the entailment relation in (4.19) doesn't hold as it stands. We also need the general assumptions hinted at in (b) above. These assumptions could also have been added to K_{pr}, thereby rendering the statement in (4.19) true without further qualification. Another strategy, adopted here, is to count for the role of such general assumptions about the nature of the worlds (in the present instance: about the structure of time) to the model theory. That is, we assume that the set of admissible models is restricted to those in which the general assumptions hold.

The second point is that the conclusion DRS K_{con} contains a dref that is bound in K_{pr} but free in K_{con}. This is a general feature of entailment relations between DRSs. It is compatible with the semantic entailment relation expressed by $\models$, which is defined as follows: $K_1 \models K_2$, iff for every admissible model M and every function f that verifies K_1 in M, there is an extension

$g \supseteq f$ which verifies K_2 in M. For the case at hand that comes to this:

The entailment relationship stated in (4.19) holds iff for every model satisfying M the relevant assumptions about the relation $\prec_M$ and every embedding function f into M that is defined for the drefs $e, t, d, d'.e', t', d''$ and verifies in M f the Conditions of the premise DRS K_{pr} can be extended to an embedding function g that is also defined for e'' which verifies the Conditions of the conclusion DRS K_{con} in M.

Exercise

State general principles involving $\prec$, $\subseteq$ and $\supset\subset$ and 'day' such that the statement in the last sentence holds for all models M that satisfy these principles.

For the presupposition of (4.12) the entailment relation formulated above is intuitively the one that has to hold for the presupposition to be verified. The recipe it illustrates – specify the premise of the relation as the merge of the discourse context and the non-presuppositional DRS to which the presupposition is left-adjoined – is right not only for this case but for *again*-presuppositions generally and likewise for presuppositions triggered by *still*, additive particles like *also* and a range of other cases. But it won't work for all presuppositions. It doesn't for instance for the 'identification presuppositions' of definite descriptions and other definite noun phrases, as we will see in Section 4.3. In the remainder of the present Section 4.2, however, we only consider presuppositions whose verification criterion is like that for the *again*-presupposition of (4.12).

Once the validity of the entailment relation in (4.19) has been established and the presupposition has thus been justified, the presupposition can be removed from the presupposition set in (4.18) and the now empty presupposition set can be eliminated. What remains is the non-presuppositional DRS, which can now be merged with the old discourse context (4.14) to form the new discourse context (which may be used to justify presuppositions of the next sentence, in case there is a next sentence). In this case (and others in which presupposition verification takes the same form as in (4.19)) the new context DRS that results from this incorporation happens to coincide with the premise DRS K_{pr} of the entailment relation in (4.19). But we will see later that this is by no means always so.

4.1.1 *again* and negation

Consider the two sentence pairs in (4.20).

(4.20)a. It rained yesterday. (But) it didn't rain again today.

b. It didn't rain yesterday. It didn't rain again today.

That both (4.20.a) and (4.20.b) are acceptable deserves a comment. For the first sentences of these two discourses are contradictories. So if they can both justify the presupposition generated by *again* in the second sentence, then that would suggest that this presupposition is a tautology. For there is a proposition A (the content of the first sentence of (4.20)), such that the presupposition follows both from A and from not-A. The solution to this little puzzle is of course that there isn't just a single presupposition that we can take *again* to generate in the second sentence of (4.20.a,b), but two different ones, one to the effect that there was an earlier period of rain and one to the effect that there was an earlier period of no rain. Which of these presuppositions will emerge in the course of DRS construction will depend on whether the negation is taken to be inside or outside the scope of *again*. The second sentence can be interpreted either way, and the interpreter of either (4.20.a) and (4.20.b) will more or less automatically and unconsciously interpret the sentence in such a way that its presupposition 'matches' the first sentence (in the sense of being verified by it).

In order that the second sentence of (4.20.a,b) can be interpreted in these two different ways, it must allow for two different LFs, which display the different scope delations between *again* and negation that lead to the different presuppositions.[8] First consider (4.20.a). In the light of what we have assumed in Section 3.10.1 about the position of negation and in the last subsection about possible positions of *again*, a plausible LF for the sentence is the one shown in (4.21).

[8] The structural ambiguity in virtue of which the second sentence of (4.20.a,b), 'It didn't rain again today', allows for these two construals with the two different presuppositions is closely connected with its exact word order. For instance, when *again* is placed in sentence-initial position, as in 'Again it didn't rain today', the only possible interpretation is one on which the presupposition is that there was an earlier 'no rain' episode. To get a sense of how the surface order of a sentence may constrain the range of possible LFs for it go through all possible permutations of 'It didn't rain again today' and decide for each of those permutations that is a well-formed sentence, what possible LFs this permutation can have. For sentences with one or more adverbials the question what LFs they have can be a quite complex one.

(4.21)

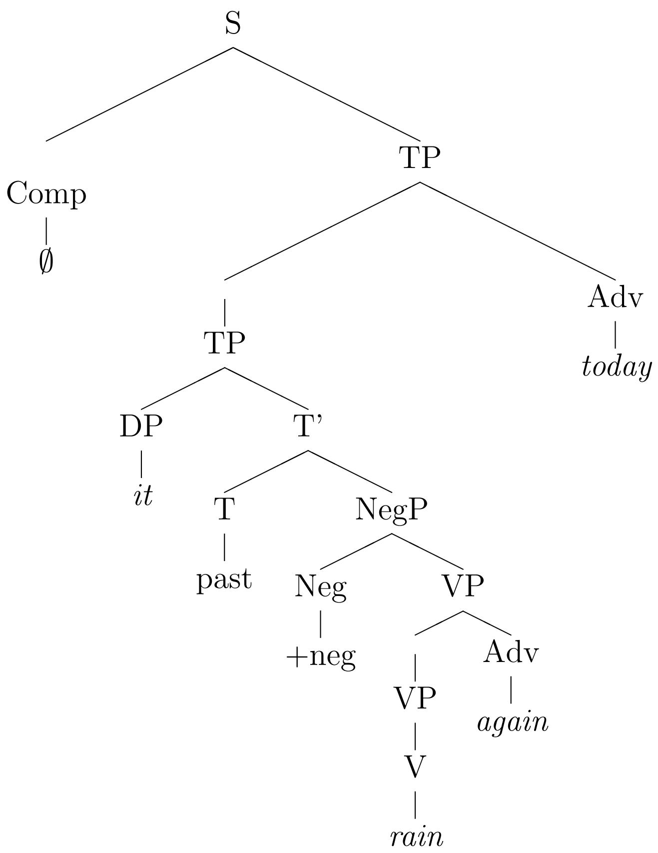

This LF generates the presupposition that there was an earlier raining event. The DRS construction that starts from this LF as input can be computed using construction rules that are already in place. In particular, the representation of the upper VP is obtained by performing the same steps as in (4.16). The result is given in (4.22).

(4.22)

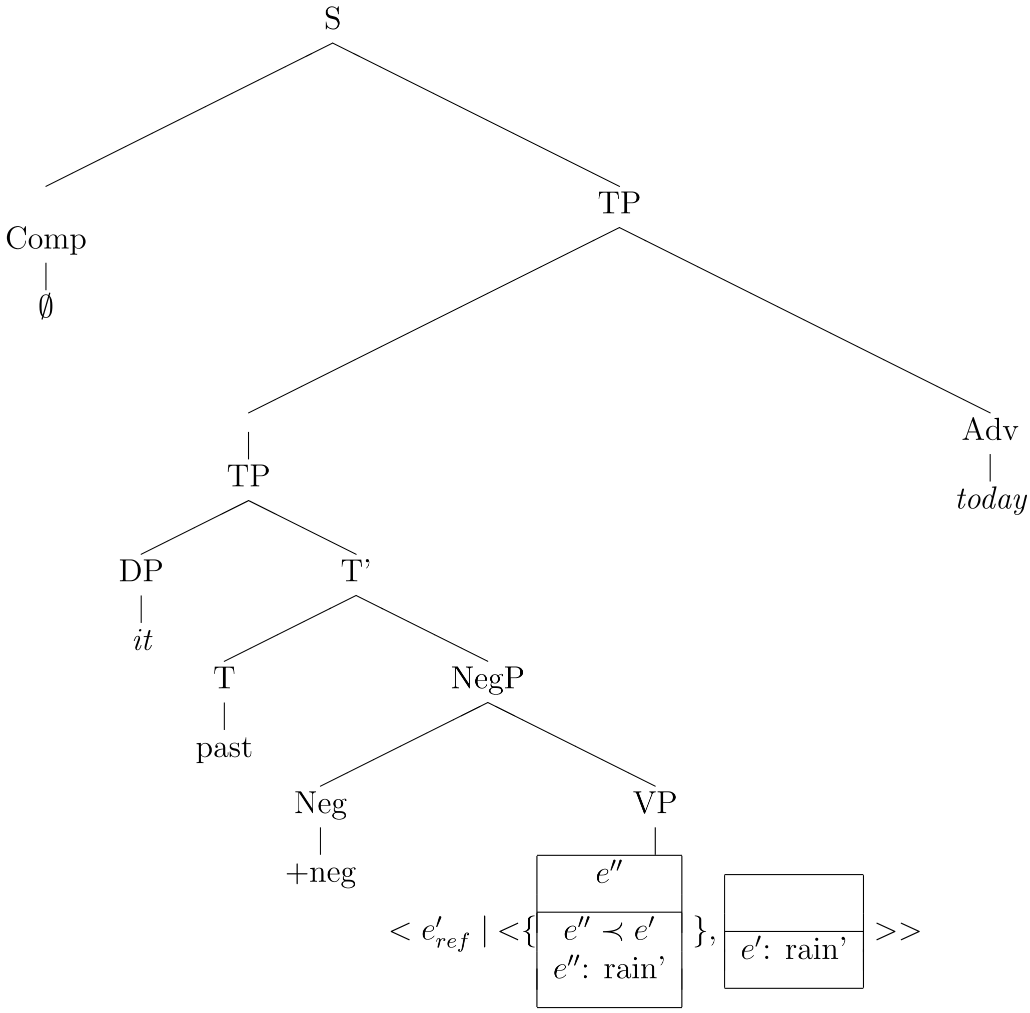

The next step is another illustration of what we saw in the previous section. When construction steps of the kinds we already encountered in Section 3 are performed on representations that involve non-empty presupposition sets, then the operations that are to be performed in these steps must be applied to the non-presuppositional part of the input representation, and leave the presuppositions as they were. The present instance of this more general principle is of special interest, however, since it involves negation. That negation is to be applied to the non-presuppositional part of the input and leave its presuppositions untouched has the effect that presuppositions *are preserved under negation*. This, you might say is part of the reason why the negation test we discussed in Section 4.1.3 can be used as a presupposition test. When a presupposition is generated below a negation operator in a sentence, then it will be there whether there is a negation higher up or not.

This explanation of why the negation test works indicates at the same time some of its limitations, to which we alluded when wee disused the negation test in Section 4.1.3. When a sentence S generates a presupposition q, then we may expect q to also be a presupposition of a sentence S' that we obtain by 'negating' S. But there is a complication. S' may look like the negation of S on the surface. But whether it can count as its negation depends on the relation between the logical forms that are assigned to the two sentences. For instance, the second sentence of (4.20.a,b) counts as negation of the second sentence of (4.12) only when its syntactic structure is analyzed as in (4.21), but not when the sentence is given the syntactic structure in (4.28) below. In this case there is at least one possible analysis of the 'negation' S' of the presupposition-generating sentence S which preserves the presupposition of S. But there are also cases where no construal of the 'negation' seems possible on which we get preservation. An example is the sentence (4.23) with the presupposition trigger *too*.

(4.23)FRED too solved the problem.

Here the focal stress on 'Fred' indicates that Fred is the 'focus associated with *too*, i.e. that the presupposition generated by *too* is that there was someone other than Fred who solved the problem. What is the negation of (4.23)? Probably anybody's spontaneous guess would be 'FRED too didn't solve the problem'. But the presupposition of this sentence would not be that someone else solved the problem, but that someone else didn't solve the problem. On the other hand there are sentences that might count as 'negations' of (4.23) in a more liberal sense, such as for instance 'FRED didn't solve the problem too', in which *too* has been moved to the end of the sentence. This sentence does have the presupposition that someone other than Fred solved the problem. The general moral of these observations is that in order for the negation test to works properly the negation that is added to a sentence S when 'its negation' is formed should not interfere with presupposition construction by taking narrow scope with respect to the presupposition trigger, with the result that this trigger now gives rise to a different presupposition.

We conclude this reflection on what the examples in (4.20) tell us about the negation test with a final, more general observation. We began our reflection by noting that the construction step triggered by a negation operator consists in a transformation of the input representation that only affects its non-presuppositional part. But we already saw in the previous section that this principle applies much more widely. It does not only apply to negation, but also to other operators, among them possibility operators such as the modal auxiliaries *may*, *might*, *could* or to periphrastics like *it may/might/could be the case that*. These too have a semantics that involves (a) scope and (b) an operator-specific content (roughly: that it is possible for what is within their scope to be true). Because these expressions also make their contributions to the non-presuppositional part of the representations that serve as their inputs, they too will preserve the presuppositions that are generated within the part of the sentence whose semantics has already been constructed, and that are therefore already represented within the input. But here too there is the same caveat: Depending on how these modalities are expressed at the surface different LFs may be possible and these may impose different scope relations between modal and presupposition triggers.

Back to the construction step that led to these reflections. I said that the negation in (4.22), represented by the feature +neg, operates only on the non-presuppositional part of the semantic representation of the VP node that is the sister of the head Neg. That is true in spirit, but it is actually not literally true. The reason is that the presupposition triggered by *again* is to the effect that there was an earlier eventuality of the same kind as the one described, where 'earlier' means 'before the time of the described eventuality'. But what is the 'time of the described eventuality' in the case of a negated sentence like the second sentence of (4.20.a)? In such cases it is the period of time during which the described eventuality is said not to have occurred. In a formal implementation of negation as proposed in Section 3.10.1, where the non-occurrence of an event e' of a certain kind is formalized as the occurrence of some event e''' which consists in there being no occurrence of en event of the type of e' within its duration, this means that the *again* presupposition should now be understood as requiring the occurrence of an event of the kind described before the time of e'''.[9]

The upshot of these last remarks should be clear. The effect of negation on the non-presuppositional part of the VP representation in (4.22) is as shown

[9]This is an awkward feature of the formalization we have adopted. It could be avoided if right away the representation that results from insertion of the lexical semantics for the verb contains a time that serves as the location time for the contribution made by the verb. This time could then become the occurrence time of the eventuality described by the verb in the absence of operators like negation and as location for the non-occurrence of such an event in case a negation higher up in the tree imposes such an 'absence' interpretation. Perhaps there is an elegant and intuitively justified of setting things up in this way. But right now I do not see how to do this. (And moreover, even if there were a good way of doing this, adopting it would require a complete overhaul of the formal implementation that we have been at great pains to put in place in these Notes – a daunting enterprise for which I would not be prepared at this point.

in Section 3.10.1 and the effect on the representation of the presupposition of the VP representation is to replace the old referential argument e' by the new referential argument e''', whose annotation qualifies it as a quantification event. The result is shown in (4.24).

(4.24)

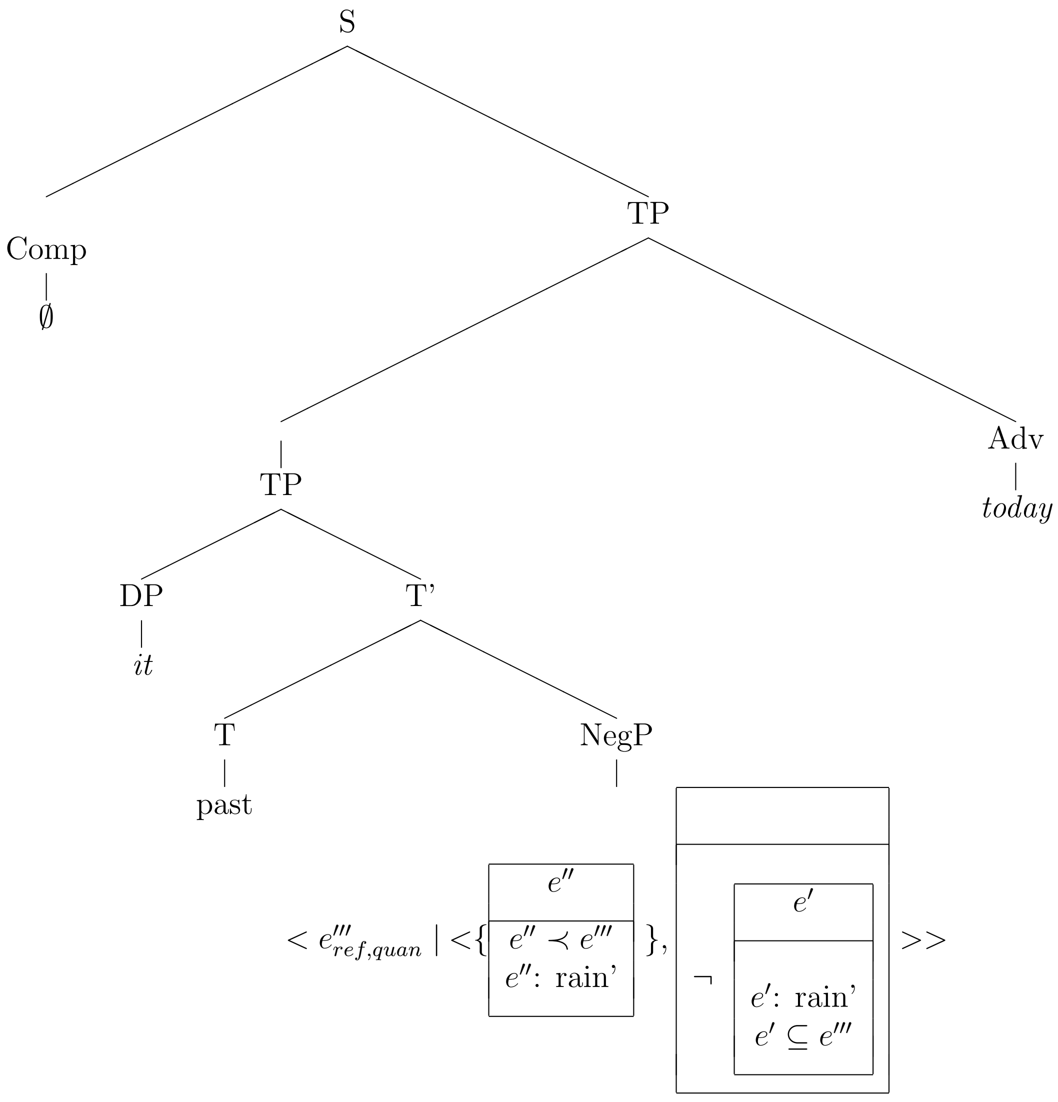

The remaining construction steps dictated by the LF structure in (4.24) lead to the preliminary representation for the S node in (4.25).

$$(4.25) < \{ \begin{array}{|c|} \hline e'' \\ \hline e'' \prec e''' \\ e'' \text{: rain'} \\ \hline \end{array} \}, \begin{array}{|c|} \hline e'''_{ref} \quad t' \quad d'' \\ \hline t' \prec n \quad dur(e''') = t' \\ \text{day}(d'') \quad n \subseteq d'' \quad e''' \subseteq d'' \\ \neg \begin{array}{|c|} \hline e' \\ \hline e' \text{: rain'} \\ e' \subseteq e''' \\ \hline \end{array} \\ \hline \end{array} >$$

As in the case of (4.12) it is clear that the presupposition of the preliminary representation in (4.25) is entailed by the discourse context. (Both presupposition and discourse context for (4.25) are the very same as for (4.12).) So the remaining steps – presupposition justification, elimination of the presupposition set, merge with the discourse context representation and store transfer – can be easily seen to lead to the DRS for the two sentence discourse (4.16.a) shown in (4.26).

$$(4.26) \quad \begin{array}{|c|} \hline e \quad t \quad d \quad d' \quad d'' \quad e''' \quad t' \\ \hline t \prec n \quad e \subseteq t \\ \text{day}(d) \quad \text{day}(d') \quad n \subseteq d' \quad d \supset\subset d' \quad e \subseteq d \\ e : \text{rain'} \\ \\ t' \prec n \quad dur(e''') = t' \\ \text{day}(d'') \quad n \subseteq d'' \quad e''' \subseteq d'' \\ \neg \begin{array}{|c|} \hline e' \\ \hline e' \text{: rain'} \\ e' \subseteq e''' \\ \hline \end{array} \\ \hline \end{array}$$

So much for (4.20.a). But how do we get a satisfactory account of (4.20.b)? Before we do anything else let me present, in (4.26.) the semantic representation of the first sentence of (4.20.b). There is no need to dwell on the construction of this DRS as all construction steps are familiar.

$$(4.27)\quad \begin{array}{|c|} \hline e\ \ t\ \ d\ \ d' \\ \hline t \prec n \quad dur(e) = t \\ \text{day}(d) \quad \text{day}(d') \quad n \subseteq d' \quad d \supset\subset d' \quad dur(e) = d \\ \neg \begin{array}{|c|} \hline e' \\ \hline e'\text{: rain'} \\ e' \subseteq e \\ \hline \end{array} \\ \hline \end{array}$$

The first question to be addressed about the second sentence of (4.20.b) concerns its syntactic structure. We already noted that (4.20.b) requires a different LF from the one we adopted in the semantic representation construction for (4.20.a). What we need now is an LF that makes it possible for the negation to be interpreted within the scope of *again*. The simplest (and perhaps only) way to achieve this is for the negation to occur within the syntactic scope of *again* in the LF. But what could make such an analysis possible? There are several options that could be considered. For instance, the LF in (4.21) might be taken as basic structure and some kind of movement - either lowering the negation in (4.21) in some way to a position below that of *again* or raising *again* to a position above negation – might give us what we need. As it is, I can see no good arguments for either of these movements. So I will assume a more neutral position, to the effect that the scopal ambiguity in the second sentence in (4.20.b) results from an indeterminacy of the attachment position of *again*: *again* can be adjoined to VP, as in (4.21), but it can also be adjoined to other constituents, and in particular to NegP. In other words, the second sentence is structurally ambiguous with regard to the question where in the LF *again* is adjoined. An interpreter of (4.20.b) will be more or less automatically led to the assumption that the LF for its second sentence is one in which *again* is attached above the negation, just as an interpreter of (4.20.a) will assume that its LF is the one in (4.21).

Let us assume, then, that the LF for the second sentence of (4.20.b) is the one given in (4.28).

(4.28)

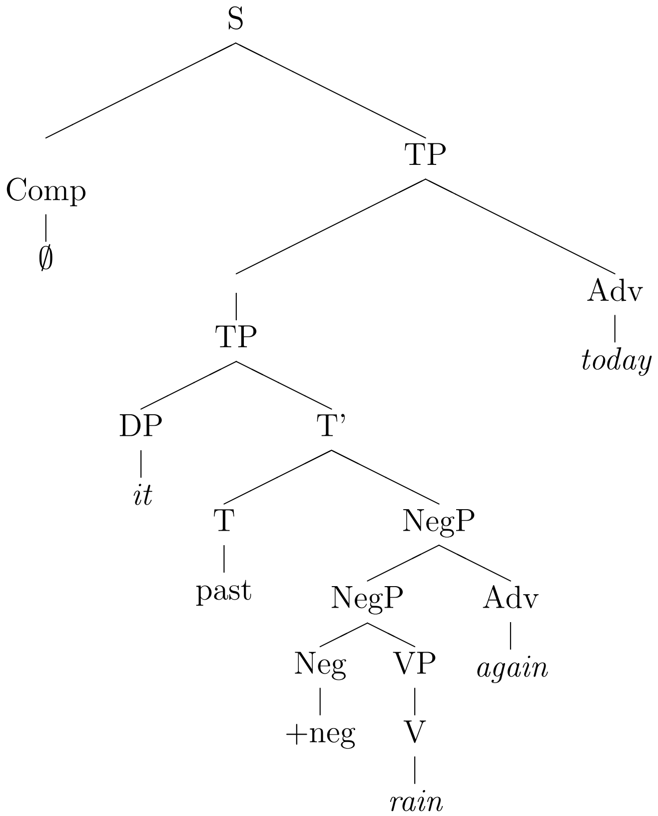

(4.29) shows the point at which the lower NegP representation has been established and (4.30) that at which *again* has made its contribution in the form of a presupposition that there was an earlier ‘no rain’ event.

(4.29)

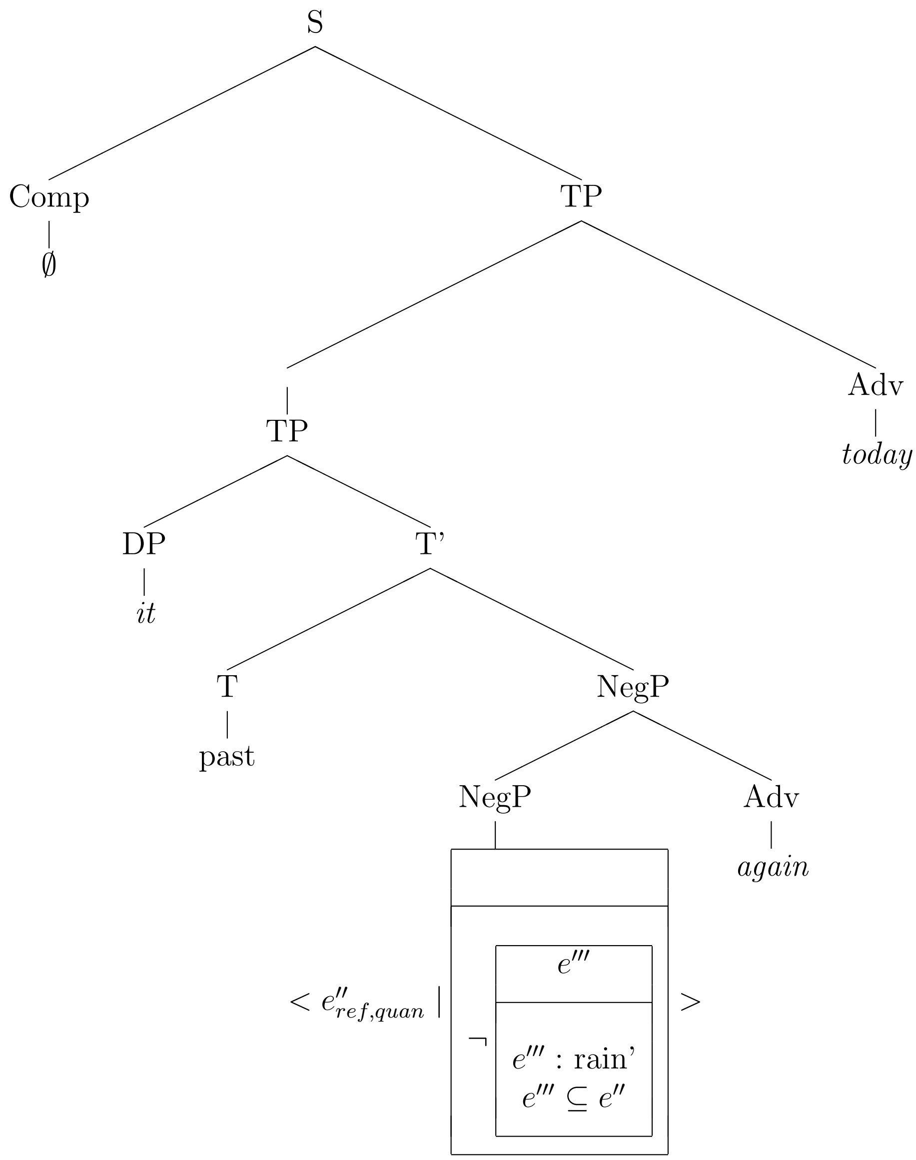
S
Comp
∅
TP
TP
Adv
today
DP
it
T'
T
past
NegP
NegP
Adv
again
$< e''_{ref,quan} |$
$\neg$
e'''
e''' : rain'
$e''' \subseteq e''$
$>$

(4.30)

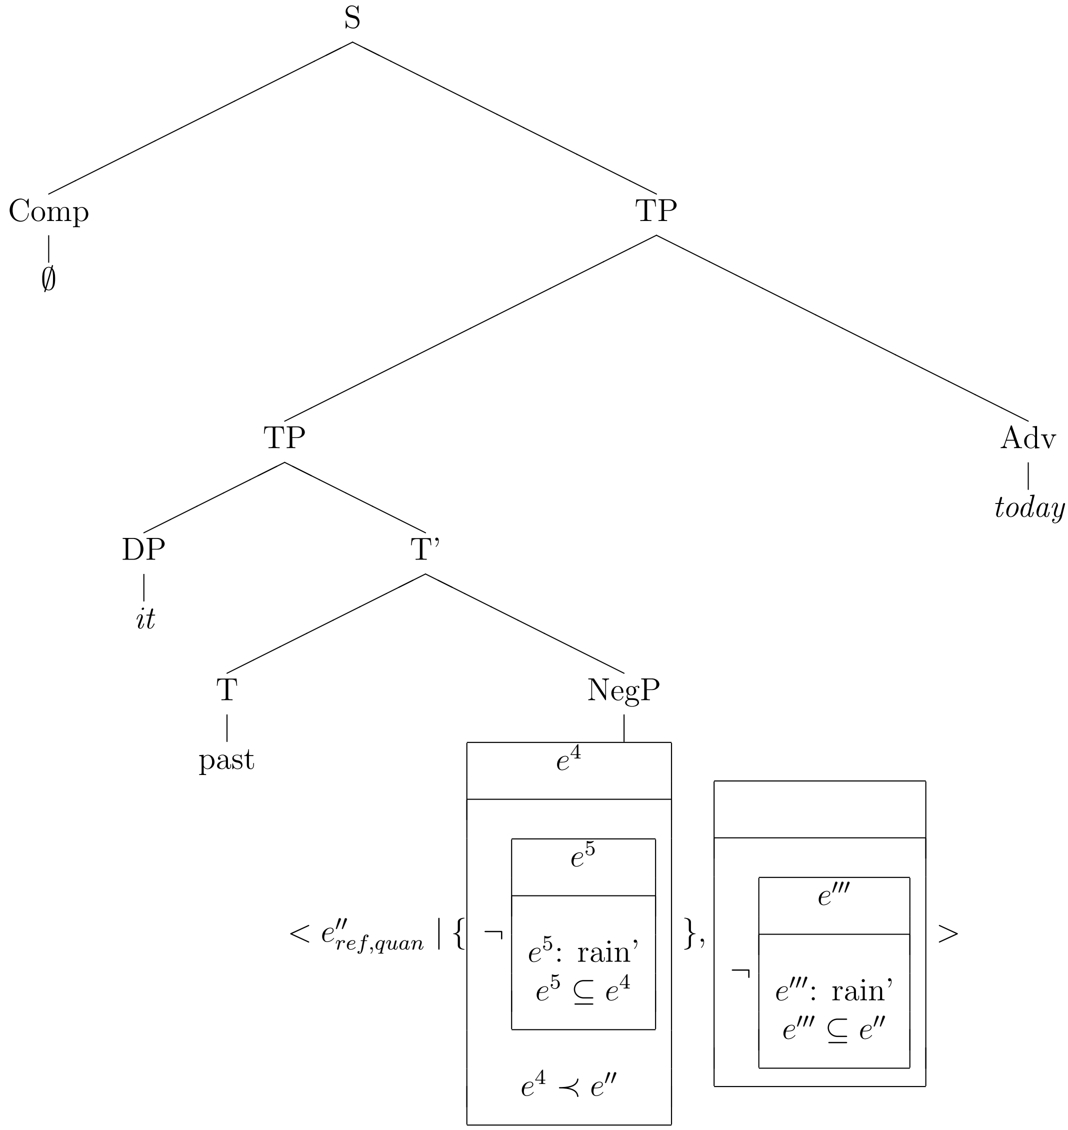

Completion of the representation construction of which (4.30) provides an intermediate record leads to the preliminary representation attached to the S node, shown in (4.31).

(4.31)

$$<\{\begin{array}{|c|}\hline e^4 \\ \hline \neg \begin{array}{|c|}\hline e^5 \\ \hline e^5\text{: rain'} \\ e^5 \subseteq e^4 \\ \hline \end{array} \\ e^4 \prec e'' \\ \hline \end{array}\}, \begin{array}{|c|}\hline e'' \quad t' \quad d'' \\ \hline t' \prec n \quad dur(e'') = t' \\ \text{day}(d'') \quad n \subseteq d'' \quad dur(e'') = d'' \\ \neg \begin{array}{|c|}\hline e''' \\ \hline e'''\text{: rain'} \\ e''' \subseteq e'' \\ \hline \end{array} \\ \hline \end{array} >$$

The presupposition of (4.31) is entailed by the discourse context (ref314.1). This can be verified in the same way as we did for (4.12) and (4.20.a). The DRS for the two sentences together is gown in (4.32).

(4.32)a.

$$\begin{array}{|c|}\hline e \quad t \quad d \quad d' \quad d'' \quad e'' \quad t' \\ \hline t \prec n \quad dur(e) = t \\ \text{day}(d) \quad \text{day}(d') \quad n \subseteq d' \quad d \supset\subset d' \quad dur(e) = d \\ \neg \begin{array}{|c|}\hline e' \\ \hline e'\text{: rain'} \\ e' \subseteq e \\ \hline \end{array} \\ t' \prec n \quad dur(e'') = t' \\ \text{day}(d'') \quad n \subseteq d'' \quad dur(e'') = d'' \\ \neg \begin{array}{|c|}\hline e''' \\ \hline e'''\text{: rain'} \\ e''' \subseteq e'' \\ \hline \end{array} \\ \hline \end{array}$$

Now that we have gone through the nitty-gritty of the semantics construction for the two discourses in (4.20) let us return briefly to the general question that such examples raise about interpretational procedure. To deal with these two examples we have been using different LFs for the sentence containing the presupposition trigger. That was necessary insofar as this is the way we have adopted to make sure that in each of the two cases the triggered presupposition can be justified in the given discourse context. But how should the parser that is responsible for delivering the right LF know what LF it should to come up with in each of these cases? The answer would seem to have to be: It doesn't know; it has no basis on which to make such a choice, for the choice should be guided by a complex process of interpretation and presupposition justification and that is information that a syntactic parser is not supposed to possess; it is information that can only be revealed by another processing module (the syntax-semantics interface; for us the DRS construction algorithm) which operates on the basis of what it gets from the parser.

For a model of sentence and discourse interpretation like that which our account provides this seems an embarrassment. For cases of the sort discussed in this section the best that a syntax-semantics interface of the kind developed here can offer is this: Whenever a sentence admits parsing into two or more alternative LFs, the Construction Algorithm will have to try out what interpretation can be constructed from each of the LFs that the parser can assign to the sentence and then discard those attempts that do not lead to a coherent interpretation. For instance, in the case of (4.20.a) the algorithm should construct preliminary representations for each of the two LFs for the second sentence of (4.20.a) that we have considered (when we dealt with (4.20.a) and with (4.20.b), respectively), establish that the presupposition of the preliminary representation that is constructed from the first LF can be justified in the context provided by the first sentence of (4.20.a), whereas the second LF leads to an incoherent interpretation, reject that one and retain the first. (When applied to (4.20.b) the same procedure would obviously lead to the opposite selection. Here the first LF would be rejected and the second retained.)

This looks like a highly unrealistic model of how human interpretation of discourse might actually work. For one thing, how many other LFs would qualify as possible LFs for the second sentence of (4.20.a,b)? And how many different LFs might there be for each of the first sentences of (4.20.a) and (4.20.b)? Going through semantic constructions for all possible combinations of LFs for the different sentences that make up a discourse or text is an enterprise of exploding complexity even when we limit attention to texts and discourses of quite moderate length.

For what must by now be three decades at least, syntactic ambiguity has been recognized as an immense obstacle to effective implementations of rule-based systems of natural language syntax and semantics. This is one reason why Computational Linguistics has turned away from rule-based approaches and is now giving nearly undivided attention to radically different methods. But is this also a reason for linguists, who want to understand the nature of linguistic knowledge, to turn away from such models as well? I do not think so. Finding rule-based grammars that can account for the vast array of possible and impossible interpretations that linguists have documented and that are now available as targets for linguistic explanation is a daunting challenge. When such a system manages to make the right predictions about such a vast array of partly highly subtle and complex data, that can be taken as evidence that the system captures essential aspects of human linguistic knowledge. Such a rule-based theory may not explain how human interpreters actually arrive at their interpretations with the speed and apparent ease with which they are able to do this. There can be little doubt that something very different must be involved in the human computation of discourses and text interpretations. But that need not mean that knowledge of grammatical rules and principles plays no part in the way we humans make sense of what we hear or read. For it could well be that rule-based grammars play an active part in interpretation, not as strict guidelines that are to be followed meticulously in the course of constructing interpretations, but as a kind of control mechanism – a filter on the interpretations that the interpreter generates by other means. In order for an interpretation to count as an interpretation of the input it would have to pass this filter; and filter checking is something that might be done very fast, even though discovering viable interpretations by a grammar-driven breadth first search hinted at above would be prohibitively labor-intensive and slow.

4.1.2 Presuppositions in Subordinate Positions: Projection or no Projection?

We noted in Section 4.1.4 that from the time when presupposition became an important concern for linguists, presupposition projection was perceived as one of the central challenges. When a sentence consists of a single clause, then the presuppositions generated by presupposition triggers occurring in it always project – in the sense that the sentence can be felicitously used only in contexts that verify those presuppositions. But when these same presupposition triggers occur in clauses that are part of a larger, multi-clausal sentence, then their presuppositions will not always project. And the reason for that, we observed, is that presuppositions can be 'locally' satisfied – by a local context all or some of which is provided by one or more other clauses that are part of the sentence.

We can account for cases of non-projection by relying on the same principles of presupposition verification that we used in the last two sections. All we need in addition is (i) to allow the semantic representations of clauses that are part of a larger sentence to play the part of discourse contexts and (ii) to define the rules that determine which clauses of a complex sentence can provide discourse contexts for the presuppositions of which other clauses. For a first, comparatively simple pair of examples that illustrate how this can work we consider the conditionals in (4.33), in which the first sentences of (4.20.a) and (4.20.b) are the antecedents and in both of which the second sentence of (4.20.a,b) is the consequent.

(4.33)a. If it rained yesterday, it didn't rain again today.

b. If it didn't rain yesterday, it didn't rain again today.

The presuppositions triggered by *again* in these conditionals – this much seems intuitively clear – are verified in essentially the same way as they are in (4.20). The only difference is that it is now the local discourse context provided by the antecedent that is responsible for the verification and not the discourse context that is provided by a preceding sentence.

We start with a closer look at (4.33.a). The first matter we have to deal with is the LF for this sentence. But what that LF should be is determined by decisions we have made at earlier points. The LF for the main clause should be that which we adopted in the last section for the second sentence of (4.20.a). And the way in which the LF for the *if*-clause and the main clause fit together was decided in Section 3.10.2 and the LF for the *if*-clause 'it rained yesterday' is covered repeatedly in Section 3. Putting these earlier decisions together we are led to the LF in (4.34).

(4.34)

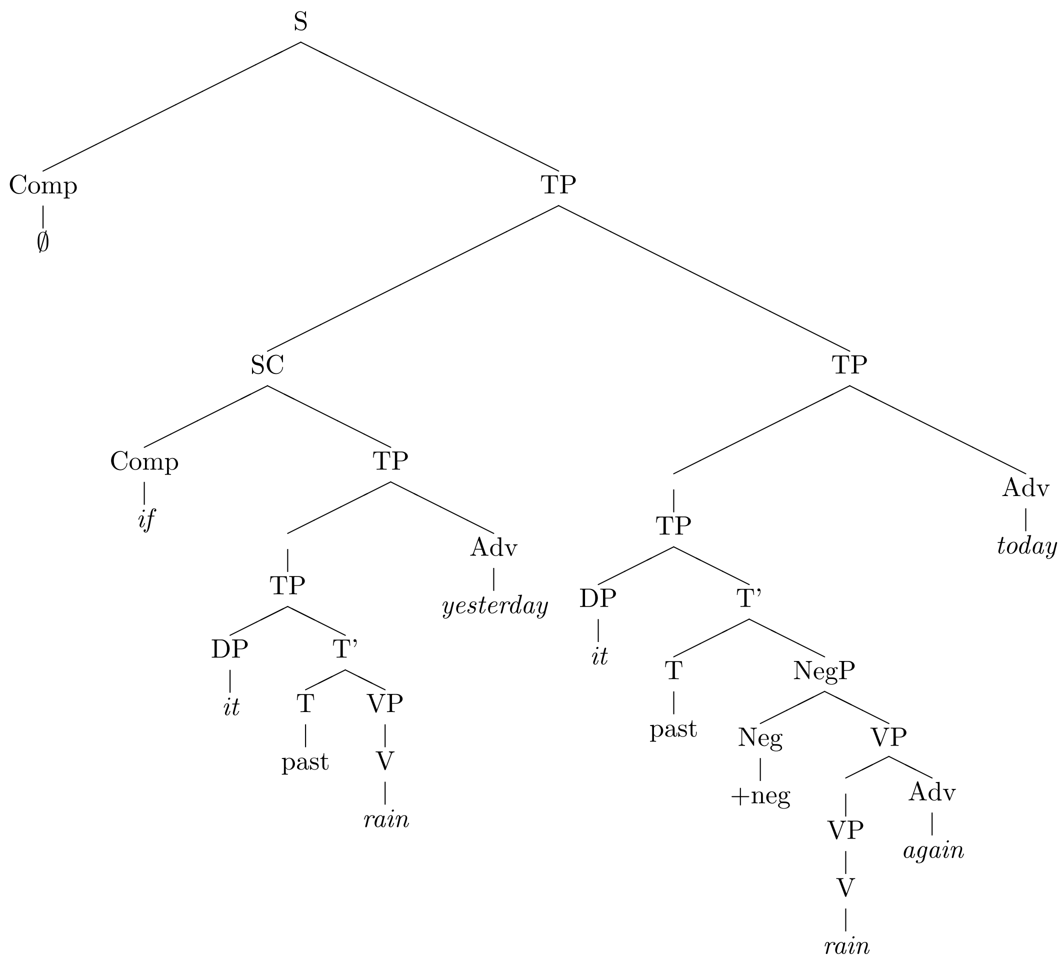

The interpretation of this LF involves steps that are familiar from earlier constructions, on the one hand those for the first and second sentence of (4.20.a) and on the other the local dref transfers from the stores of antecedent and consequent of the $\Rightarrow$ condition to the Universes of the non-presuppositional DRSs following those stores that led in Section 3.10.2 from (3.154) to (3.155).

(4.35)

$$\begin{array}{|c|}\hline \\ \hline \begin{array}{|c|}\hline e\ \ t\ \ d\ \ d' \\ \hline t \prec n\ \ e \subseteq t \\ \text{day}(d)\ \ \text{day}(d') \\ n \subseteq d'\ \ d \supset\subset d' \\ e \subseteq d \\ e\text{: rain'} \\ \hline \end{array} \Rightarrow \\ \\ < \{ \begin{array}{|c|}\hline e''' \\ \hline e'''\text{: rain'} \\ e''' \prec e' \\ \hline \end{array} \}, \begin{array}{|c|}\hline e'\ \ t'\ \ d'' \\ \hline t' \prec n\ \ dur(e') = t' \\ \text{day}(d'')\ \ \ n \subseteq d'' \\ dur(e') = d'' \\ \neg \begin{array}{|c|}\hline e'' \\ \hline e''\text{: rain'} \\ e'' \subseteq e' \\ \hline \end{array} \\ \hline \end{array} > \\ \hline \end{array}$$

Since the antecedent of a conditional DRS Condition as accessible from its consequent, its representation can be used as context in the verification of presuppositions generated within the consequent. This means that the the very same entailment relation – that shown in (4.19) – will verify the presupposition in (4.35) that also verified the presupposition of the second sentence of (4.12) and the second sentence of (4.20.a). In this case the premise of the entailment relation is obtained as the merge of the antecedent DRS in (4.33) and the non-presuppositional DRS of the consequent.

The remainder of the construction is also as before. After the presupposition has been justified, the presupposition set can be dropped from the representation of the consequent, and that gives us the final representation of (4.33), shown in (4.36).

(4.36) (DRS for the upper TP node of (4.33.b))

$$
\begin{array}{|c|}
\hline
\\
\hline
\begin{array}{|c|}
\hline
e\ t\ d\ d' \\
\hline
t \prec n\ e_0 \subseteq t \\
\mathrm{day}(d)\ \ \mathrm{day}(d')\ \ n \subseteq d' \\
d \supset\subset d' \quad e \subseteq d \\
e\text{: rain'} \\
\hline
\end{array}
\Rightarrow
\begin{array}{|c|}
\hline
e'\ t'\ d'' \\
\hline
n \prec t'\ \ dur(e') = t' \\
\mathrm{day}(d'')\quad n \subseteq d''\ \ dur(e') = d'' \\
\neg \begin{array}{|c|}
\hline
e \\
\hline
e\text{: rain'} \\
e \subseteq e' \\
\hline
\end{array} \\
\hline
\end{array}
\\
\hline
\end{array}
$$

The story for (4.33.b) is much the same. This time the LF we want – one that yields the presupposition that there was an earlier ‘no rain’ period – is that in (4.37).

(4.37)(LF for (4.33.a))

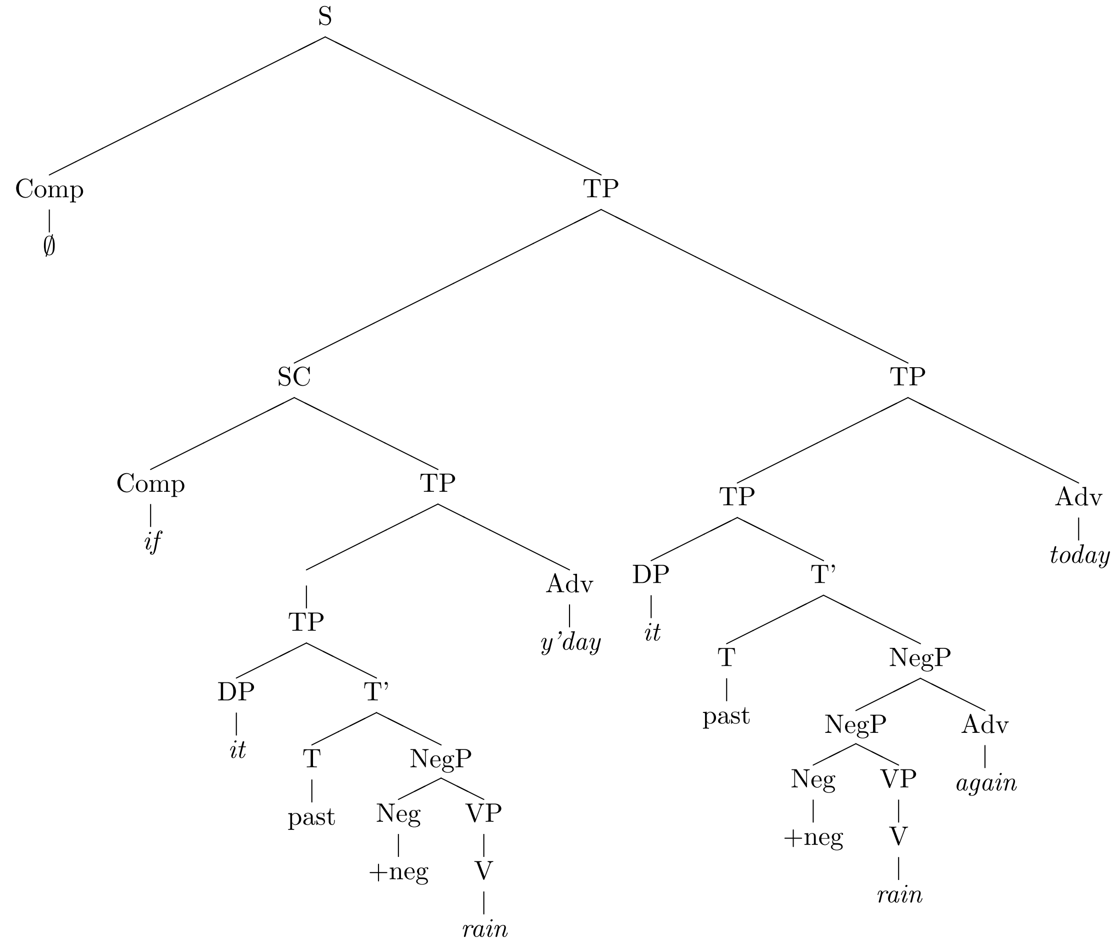

The steps leading to the preliminary semantic representation for this LF are also familiar from earlier constructions and the construction is left to the reader. The preliminary representation is shown in (4.38).

(4.38) (DRS for (4.33.b))

$$
\begin{array}{|c|}
\hline
\\
\hline
\begin{array}{|c|}
\hline
e \;\; t \;\; d \; d' \\
\hline
t \prec n \;\; dur(e) = t \\
\mathrm{day}(d) \;\; \mathrm{day}(d') \;\; n \subseteq d' \\
d \supset\subset d' \;\; dur(e) = d \\
\neg \begin{array}{|c|}
\hline
e' \\
\hline
e' : \mathrm{rain'} \\
e' \subseteq e \\
\hline
\end{array} \\
\hline
\end{array}
\Rightarrow \\
< \{
\begin{array}{|c|}
\hline
e^4 \\
\hline
e^4\text{: rain'} \\
e^4 \prec e''' \\
\neg \begin{array}{|c|}
\hline
e^5 \\
\hline
e^5 : \mathrm{rain'} \\
e^5 \subseteq e^4 \\
\hline
\end{array} \\
\hline
\end{array}
\},
\begin{array}{|c|}
\hline
e'' \;\; t' \;\; d'' \\
\hline
t' \prec n \;\; dur(e'') = t' \\
\mathrm{day}(d'') \;\; n \subseteq d'' \;\; dur(e'') = d'' \\
\neg \begin{array}{|c|}
\hline
e''' \\
\hline
e''' : \mathrm{rain'} \\
e''' \subseteq e'' \\
\hline
\end{array} \\
\hline
\end{array}
> \\
\hline
\end{array}
$$

Let me state once more the main point of this subsection. The presuppositions triggered by *again* in (4.33.a,b) are the same as the ones *again* triggers in (4.20) and the verification of these presuppositions is also the same in these cases, except the role off the global context in the case of (4.20.a,b) is in the cases (4.33.a,b) taken over by the local context provided by the conditional's antecedent. It should be emphasized in this connection, however, that all the cases of presupposition representation construction and presupposition justification we have considered so far are comparatively simple. First, in all these cases there was just one presupposition that needed to be dealt with. In general, sentence interpretations may have to deal with multiple presuppositions, and these presuppositions may be adjoined to the representations of different sentence parts. When this is so, a certain hierarchy will often be imposed on presupposition justification. For instance, when presuppositions are generated within both the antecedent and the consequent of a conditional, then the presupposition(s) of the antecedent will have to be justified before the non-presuppositional part of the antecedent can be used as local context in the justification of the presupposition(s) of the consequent. Furthermore, it is also possible for presuppositions to be generated inside other presuppositions. This too complicates the strategies of presupposition justification. (We will encounter examples of these possibilities in Section 4.3.)

Second, thus far justification has always taken the default form of verification without any need for accommodation. Third, verification was so far always either global (i.e. the context was just the global context) or strictly local (i.e. the only contextual information used was that provided by the local context). In Section 4.2.5 we will encounter examples where global and local context are both playing their part in the justification of a given presupposition.

4.1.3 A lexical entry for *again*

Before we turn to these examples, however, a question needs to be addressed that is now close to overdue. What could a lexical entry be like for a presupposition trigger such as *again*? It should be clear from the way in which the construction of presuppositions triggered by *again* was handled in the last three sections that within our semantic representation system *again* functions as an operator, which turns eventuality descriptions into other eventuality descriptions. In this respect it is like the tense and aspect operators triggered by such features as past, +prog, +perf, and disp. We adopt the general format for lexical entries of operators that we used to specify the entries for such operators in Section 3, as involving schematic descriptions of both the input representation to the operator and the representation that the operator returns when applied to this input. There are certain differences with the operator entries from Section 3, however, which reflect the fact that*again* is a (pure) presupposition trigger. *again* computes a presupposition from the input representation, which it adds to the presupposition set of the input representation (or, in case there was no presupposition set so far, it forms a new presupposition set with the just computed presupposition as its only member).

A second difference with earlier operator entries, implicit in what was just said about the first difference, is that both input and output representations are now of the more complex form that distinguishes the representations of Section 4 from those of Section 3. The general format with which we are working now is that of a representation that consists of i) a store, (ii) a presupposition set and (iii) a 'non-presuppositional' part, but where, moreover, this 'non-presuppositional' part may contain conditions that have stores and presupposition sets in their turn, as in some of the representations of the last section. More specifically: the input representation will be given in the form: $< \beta_1, .., \beta_n \, |< \{K_1, ..., K_m\}, \ K >>$ and the output representation will add a schematic description of the generated presupposition to this representation.[10]

In first approximation we can describe the difference between presupposition triggers such as *again* and the non-presuppositional operators from Section 3 as follows:

(4.39)a. The operator determined by a presupposition trigger will, when operating on an input representation $< \beta_1, .., \beta_n \, |< \{K_1, ..., K_m\}$, $K >>$, add one or more presuppositions to the set $\{K_1, ..., K_m\}$, while leaving the non-presuppositional part K of the input invariant.

b. A 1-place non-presuppositional operator will, when operating on an input representation $< \beta_1, .., \beta_n \, |< \{K_1, ..., K_m\}, K >>$ modify the non-presuppositional DRS K while leaving the presupposition set $\{K_1, ..., K_m\}$ untouched.

This way of characterizing the difference between presuppositional and non-presuppositional operators requires several comments. First, as we just saw when we looked at the conditional sentences (4.20.a,b) in the last section, the presuppositions triggered within the input representations of binary operators may end up in embedded positions in the output representation. (This is what enables the present approach to account for non-projection of presuppositions through local justification.) Second, it isn't true in general that presuppositional operators leave the presuppositions of the input representation untouched and only add new presuppositions to the otherwise

[10] What is strictly speaking called for at this point is a recursive definition of the set of these more complex representations, in which DRS Conditions are defined as structures consisting of a store, a presupposition set and ('non-presuppositional') DRS, where DRSs are defined as pairs consisting of a Universe and a set of DRS Conditions of this new more general kind and where the presupposition representations that make up the presupposition sets are in general also representations of the kind that is being defined. It is left as an exercise to the reader to give such a definition which should minimally cover the representations that have been discussed up to this point.

unaffected presupposition set. Consider the sentence 'John stopped smoking again'. This sentence asserts that there was an event to the effect that John stopped smoking while presupposing an earlier event of the same kind. But the event description 'stop smoking' comes itself with a presupposition to the effect that just prior to the occurrence of the event described by the phrase the subject did smoke. So the representation of the event description 'stop smoking' will consist of a presupposition (that there was a state of the subject smoking) and a non-presuppositional part asserting the occurrence of an event that terminated that state. When this representation becomes the input to *again*, the output ought to be a representation in which the presupposition – that there was an earlier occasion of the subject's stopping smoking – is a replica of the input structure, consisting in its turn of a presupposition (a state of smoking) and a non-presuppositional part (describing the transition from that state to a state of non-smoking).

This is just one example of a sentence in which a presupposition trigger occurs in the scope of another presupposition trigger. The interactions between different presupposition triggers within a single sentence are a source of considerable complexity. So far there hasn't been a great deal of work on this intriguing topic. (Within a DRT-based framework, an early discussion of nested presuppositions that are generated by complex definite noun phrases in which one definite is in the scope of another one, such as *his daughter*, can be found in (Van Der Sandt 1992). (The presupposition triggered by the DP *his daughter* as a whole contains another presupposition, which is triggered by the pronoun *his*.) For some further explorations of the topic of presupposition nesting see (Kamp 2001*a*).) In the lexical entry (4.40) that is shown below we account for the possibility that the input representation to *again* may have a non-empty presupposition set by assuming that the presupposition that *again* adds to the input representation is itself a combination of (a) a non-presuppositional part and (b) presuppositions, which are obtained by substituting the new eventuality argument of the *again* presupposition (ev' in (4.40)) for the referential argument ev of the input representation. See the entry (4.40), without which this last description must be hard to parse.

There is one other observation I want to make about the semantics of *again* before we wind up this section with the lexical entry that we will adopt. The interpretation of *again* to which the entry in (4.40) is restricted is its so-called *repetitive* interpretation. On its repetitive interpretation, the application of *again* to an eventuality description adds to this description a presupposition to the effect that an eventuality of the same description occurred earlier. But when *again* occurs in a suitable position of a clause which is used to describe events with target states, it also allows for a so-called *restitutive* interpretation, according to which the eventuality that must have obtained at an earlier time is note an event of the type described by the clause, but only a state of the corresponding target state type. Intuitively, the restitutive interpretation of *again* presents the described event e as the 'restitution' of the presupposed state, which reemerges in the form of e's own target state. An example is the sentence 'Mary closed the window again.', which can be interpreted as presupposing that the window was closed at some earlier time, but was then opened and then became closed once more through what Mary did. On this interpretation there is no implication that the window had been previously closed by Mary, or for that matter that it had been closed by anybody. For instance, it might have been the case that the window was prefabricated as closed and inserted in its closed condition into the wall of a newly built house. It was then opened only once, whereupon Mary closed it 'again'.

English *again* is not the only word that allows for restitutive as well as repetitive interpretations. Other Germanic languages, such as German or Dutch, have such words as well (*wieder* in German and *weer* in Dutch). But by no means al languages do. (For discussion see (Beck & Snyders 2001).) In these Notes we will have nothing more to say about restitutive interpretations.

We conclude this section with an explicit formulation of a lexical entry for *again* that only covers its repetitive use.

(4.40) (lexical entry for repetitive *again*)

again

Sel. Restr: eventuality description

Sem.Repr: $< ev_{ref}, \beta_2, ., \beta_n \mid < \{K_1, ., K_m\}, K >> \rightsquigarrow$

$$<ev_{ref}, \beta_2, ., \beta_n \mid <\{K_1, ., K_m, \boxed{\begin{array}{c} ev' \\ \hline <\{K_1[ev'/ev], ., K_m[ev'/ev]\}, K[ev'/ev]> \\ ev' \prec ev \end{array}}\}, K>>$$

In applications of this entry the eventuality dref ev' must always be of the same ontological type as ev – an event dref if ev is an event dref, a state dref if ev is a state dref.

To repeat, (4.40) describes the presupposition triggered by *again* as the result of forming a 'copy' of the non-presuppositional part of the input representation *together with all its presuppositions*. This is consistent with the observation above about *stop smoking*. Not all presupposition triggers work quite this way. In particular, we will see later on in Section 4 that the presuppositions triggered by tenses and by most definite noun phrases involve different construction principles.

4.1.4 Presupposition Justification by Local and Global Context, with and without Accommodation

We now come to some cases in which presupposition justification depends on both global and local context.

Before anything more is said I would like you to look at the following pair of sentences – on the surface the one and only difference between them is the presence of *again* in the second sentence and its absence in the first – and decide in what way, if any, they convey different information (in what way they differ in meaning', in an intuitive, pre-theoretical sense of 'meaning'). Try to arrive at your own assessment of these sentences in as unprejudiced a way as possible, as if the discussions of the preceding sections were entirely unknown to you. And take your time.

(4.41)a. We won't have pizza on Bill's birthday, if we are going to have pizza on Mary's birthday.

b. We won't have pizza again on Bill's birthday, if we are going to have pizza on Mary's birthday.

For those who still have Section 4.1.5 clearly in mind judging these sentences from a theoretically unbiased perspective will have been hardly possible, since the sentences in (4.41) are closely reminiscent of two examples that were mentioned in that section, viz. (4.9) and (4.11). (See footnote 6 on p. 554.) In particular, the second sentence of (4.41) differs from (4.11) only in having *if*-clause and main clause reversed, with the *if*-clause occurring in sentence-final rather than sentence-initial position. (I reversed *if*-clause and main clause in (4.41) because it is my impression that the contrast between presence and absence of *again* is somewhat sharper with sentence-final than sentence-initial *if*-clause.)

I have found that people who are asked to describe the difference between (4.41.a) and (4.41.b) come up with something like this: "(4.41.b) implies that Bill's birthday comes before Mary's birthday. There may some suggestion of this also in (4.41.a), but here the effect is much weaker."
But before we launch into a treatment of the sentences in (4.41), let me present the minimal pair in (4.42), to which we will turn when we will be done with (4.41). I want you to look at this example in as unprejudiced a way as you can and decide for yourself what the difference is between the two sentences of this pair.

(4.42)a. I know that Mary's birthday will be only after Bill's birthday. But even so, this is what I propose: We won't have pizza on Bill's birthday, if we are going to have pizza on Mary's birthday.

b. I know that Mary's birthday will be only after Bill's birthday. But even so, this is what I propose: We won't have pizza again on Bill's birthday, if we are going to have pizza on Mary's birthday.

Here is my own reaction to these sentences: (4.42.b) seems to imply that there is must be some third occasion of 'us' having pizza, which is before Bill's (as well as Mary's) birthday. (4.42.a) on the other hand carries no such implication.

In the light of this judgment of (4.42.b) the strength with which (4.41.b) appears to suggest the implication that Bill's birthday is after Mary's birthday is surprising. If there is any truth to the impression that (4.41.b) entails that Mary's birthday comes before Bill's, it is an entailment that can be defeated by providing a more explicit context. So long as we have not enough contextual information for the justification of certain presuppositions, and thus feel compelled to accommodate, the accommodation we should make may seem uniquely determined by the presupposition and by the context as we have it. In such cases what we feel compelled to accommodate comes across to us as entailed; from a phenomenological standpoint there isn't muh to give and take between dictating a proposition and entailing it.

We will come back to this aspect of presupposition justification towards the end of the section. But to get a grip on the formal treatment of sentences of the kind we have been exploring, it is best to begin with the least controversial case, in which no accommodation is required. So let us return to our first example of the kind we are discussing in this section, the one given as (4.9) in Section 4.1.5. We consider a slight variation of this example, which somewhat lightens the task of constructing a DRS for the first sentence, which sets the context for the second sentence.

(4.43) Bill's birthday will be two days after Mary's birthday. So if we are going to have pizza on Mary's birthday, we won't have pizza again on Bill's birthday.

The second sentence of this example resembles the examples of Section 4.2.3 in being a conditional whose consequent has an occurrence of *again*. But the difference with the sentences of Section 4.2.3 is that justification of the presupposition generated by this occurrence of *again*in the second sentence of (4.9.2) requires the global context provided by the first sentence as well as the local context provided by its own antecedent.

A formal treatment of this example, along the lines of what have seen in Sections 4.1.1 - 4.1.3, is not all that difficult. But let us, as a kind of joint exercise, have a look at its most important stages. (4.44) is a representation of the first sentence of (4.9.2). There are some details having to to with the construction of this sentence that I do not want to get bogged down in at this point. One of them is that noun phrases of the form 'x's birthday' can be used to refer to particular calendar days, and that that is the way in which *Bill's birthday* and *Mary's birthday* are used here. A further question is how to analyze these DPs into their components, *John*, *'s* and *birthday*. I haven't bothered to do tho here, and treated *Bill's birthday* and *Mary's birthday* as if they were proper names of particular days lying within the twelve month period starting at the imd of utterance. The analysis of the copular construction *be two days after Mary's birthday* also involves some non-trivial issues. I am assuming that the copular complement *two days after Mary's birthday* is a prepositional phrase with a tacit temporal preposition. This PP is a predicate of entities that have temporal location and says of them that their temporal location is at two days before Mary's birthday. Most of the ingredients of a formal compositional analysis can be found in Section 3.11. Here we just assume that all construction principles that are needed for the construction of a DRS for the first sentence of (4.9.2) are in place and present the result without further ado.

(4.44)

$d_1 \quad d_2 \quad d'$
day'(d_1) day'(d_2) day'(d') $n \prec d_1 \quad n \prec d_2$ 'J Bill's-birthday'(d_1) 'Mary's-birthday'(d_2) $d_2 \supset\subset d' \supset\subset d_1$

The construction of a preliminary representation for the second sentence can be carried out by using very much the same rules that we applied in the construction of representations for the conditionals in Section 4.2.3. Once more we need the assumption that *Bill's birthday* and *Mary's birthday* are being used to denote particular days, and in addition that their occurrences in the second sentence refer to the same days as their occurrences in the first sentence. (Coreference with an earlier occurrence of the same phrase is one of the possible interpretations of 'discourse-second occurrences' of definite descriptions. And it is a strongly preferred one: for the most part discourse-second occurrences are coreferential with the same-shaped occurrences that precede them in the given discourse or text. For more on this see Section 4.4.) In the preliminary representation of the second sentence in (4.45) the coreference of the two occurrences of *Bill's birthday* and that of the two occurrences of *Mary's birthday* have been captured in the same way that pronominal anaphora was treated in PART I, viz, by introducing new drefs d_1' and d_2' and stipulating these to be coreferential with the drefs d_1 and d_2 of (4.44) by adding equations to that effect. For more see Sections 4.3 and 4.4.

Two last points about the preliminary representation of the second sentence: (i) I haven't bothered to analyze *have pizza* into its parts, but treat it as if it was a lexicalized intransitive verb; (ii) the pronoun *we* is represented by the dref W. *we* is a kind of indexical, insofar as its denotation must always include the speaker; but its denotation must also contain at least one other element. The rules for how this remainder of the denotation is determined are more complex. In (4.45) I have assumed that the two occurrences of *we* refer to the same set. Strictly speaking that assumption is unwarranted – for all the sentence says, the guests on Bill's birthday (and therewith the set of potential pizza eaters on that occasion) and the guests on Mary's birthday (and therewith the set of potential pizza eaters then) need not coincide. But in (4.45) I have ignored this possibility and assumed that the two occurrences of *we* denote the same set. We do this by using two drefs W and W' for the two occurrences of *we* and setting them equal. The drefs are both placed in the main Universe, as if *we* was some kind of proper name.[11]

$$
(4.45)\quad
\begin{array}{|c|}
\hline
W \;\; W' \;\; d'_1 \;\; d'_2 \\
\hline
\text{`Mary's-birthday'}(d'_1) \;\; \text{`John's-birthday'}(d_{2'}) \;\; d'_1 = d_1 \;\; d'_2 = d_2 \;\; W' = W \\
\begin{array}{|c|}
\hline
e \;\; t \\
\hline
n \prec t \;\; e \subseteq t \;\; e \subseteq d'_1 \\
e\text{: `have-pizza'}(W) \\
\hline
\end{array}
\Rightarrow \\
<\{
\begin{array}{|c|}
\hline
e''' \\
\hline
e'''\text{: `have-pizza'}(W) \\
e''' \prec e' \\
\hline
\end{array}
\},
\begin{array}{|c|}
\hline
e' \;\; t' \\
\hline
n \prec t' \;\; dur(e') = t' \\
dur(e') = d'_2 \\
\neg
\begin{array}{|c|}
\hline
e'' \\
\hline
e''\text{: `have-pizza'}(W') \\
e'' \subseteq e' \\
\hline
\end{array} \\
\hline
\end{array}
> \\
\hline
\end{array}
$$

The presupposition of (4.45) is verified because it is entailed by its locally available context, in which the representation of the antecedent of the second sentence is combined with that of the global context (the representation of

[11]The possibility that the two occurrences of *we* might denote distinct sets is not without importance for what is our focus of interest here, viz. the justification of the *again* presupposition. If the sets are different, then it isn't clear from what has been said so far that the antecedent of the conditional can be used to verify the presupposition. But intuition suggests that the justification of the presupposition doesn't depend on this question: the presupposition appears to be justified whether or not the two *we*-sets coincide or not. Apparently events that can both be described as events of 'we have pizza' count as of the same type for the purposes of *again presuppositions* irrespective of precisely what the sets are that the *we*s of these descriptions refer to. At least this appears to be the case so long as a single speaker is involved and the two sets share this speaker as their most salient element. There is a more general problem here: I do not fully understand what the principles are that govern the identity of event types for the purpose of justifying *again* presuppositions.

the first sentence of (4.9), as given in (4.44)). Formally, verification follows because of there is entailment of the presupposition in (4.45) by a premise DRS that is the merge of (i) the global context (4.44), (ii) the antecedent DRS of the conditional DRS condition in (4.45) and (iii) the non-presuppositonal DRS to which the presupposition is left-adjoined in (4.45). In other words, the premise is the DRS given in (4.46).

(4.46)

$$\begin{array}{|c|}\hline d_1 \;\; d_2 \;\; d' \;\; W \;\; W' \;\; d'_1 \;\; d'_2 \;\; e' \;\; t' \\ \hline \\ \text{day'}(d_1) \quad \text{day'}(d_2) \quad \text{day'}(d') \\ n \prec d_1 \;\; n \prec d_2 \\ \text{John's-Birthday}(d_1) \quad \text{Mary's-Birthday}(d_2) \\ d_2 \supset\subset d' \supset\subset d_1 \\ n \prec t \;\; e \subseteq t \;\; e \subseteq d_1 \\ W' = W \\ e\text{: `have-pizza'}(W) \\ \\ \text{`Mary's-birthday'}(d'_1) \;\; \text{`John's-birthday'}(d_{2'}) \quad d'_1 = d_1 \;\; d'_2 = d_2 \\ \\ n \prec t' \quad dur(e') = t' \quad dur(e') = d'_2 \\ \\ \neg \begin{array}{|c|}\hline e'' \\ \hline e''\text{: `have-pizza'}(W') \\ e'' \subseteq e' \\ \\ \hline \end{array} \\ \\ \hline \end{array}$$

Making sure that (4.46) entails the presupposition of (4.45) is not difficult, though it involves careful checking, in particular, that the events represented by the drefs e and e' stand in the $\prec$ relation to each other. Here too that inference presupposes a number of assumptions about the relevant temporal relations. Again this checking is left to the reader.

This concludes our discussion of (4.9.2), the variant among our pizza examples in which the *again* -presupposition can be verified without the need to resort to accommodation. We now return to the examples in (4.41). They are repeated below, this time together with yet another pair of variants.

(4.47)a. We won't have pizza on Bill's birthday, if we are going to have pizza on Mary's birthday.

b. We won't have pizza again on Bill's birthday, if we are going to have pizza on Mary's birthday.

c. We have just had pizza on Freddy's birthday. So we won't have pizza again on Bill's birthday, if we are going to have pizza on Mary's birthday.

d. We didn't have pizza on Freddy's birthday. And we won't have pizza again on Bill's birthday, if we are going to have pizza on Mary's birthday.

A few pages ago I drew attention to the strength of the impression one has that (4.47.b) entails that Bill's birthday will cone after Mary's so long as that sentence is presented on its own (and in a 'neutral' context, which carries no information about when the the protagonists have their respective birthdays). And yet, this impression can evaporate when the sentences is presented in a context that provides an alternative verifier for its *again*-presupposition, as we see in (4.47.c). (4.47.c) has a natural interpretation in which the pizza eating event on Freddy's birthday, which has already taken place and therefore will be a past event at the time of Bill's birthday, which is still in the future, serves as presupposition verifier.

That apparent entailments of presupposition-carrying sentences can be so easily and completely overwritten when they are placed in the right context is one of the most remarkable phenomena in the presupposition literature and one of its most trying challenges. But the phenomenon also shows how much resemblance there is between presupposition and anaphora and thus provides further support for the view of these phenomena that we owe Van Der Sandt and Geurts and a version of which we have adopted. Compare the pair (4.47.b), (4.47.c) with the pair (4.48.a) and (4.48.b).

(4.48)a. The doctor said she had made a mistake.

b. Mary refused to believe it. The doctor said she had made a mistake.

When you are presented with sentence (4.48.a) on its own, as in (4.48.a), then the only interpretation that seems possible for*she* is as anaphoric to *the doctor*. Note well that that requires a little bit of an accommodation: *she* can be interpreted as anaphoric to *the doctor* only if it is assumed that the doctor is a woman. Some people may find this accommodation harder than others, and the ease is also a matter of what generation you belong to or what phase in history we are considering. But it is good bet that anyone who is confronted with (4.48.a) on its own will accommodate the assumption that the doctor of the first sentence is female.

Not so in (4.48.b). In the context provided by the first sentence of (4.48.b) *she* becomes ambiguous: the pronoun can now be interpreted either as anaphoric to *the doctor* or as anaphoric to *Mary*. The first option still requires the accommodation that the doctor is a woman. But this accommodation has lost its halo of inevitability, since there is an alternative resolution for the pronoun which doesn't depend on it. (Perhaps the more misogynous, the greater your inclination to resolve *she* to *Mary*.)

For pronouns the phenomenon illustrated by the pair (4.48.a,b) more contextual information may have the effect of widening the range of interpretational options rather than restricting it – is well-known. (and it is a directly predictable consequence of the way pronominal anaphors has been dealt with in DRT since its beginnings). But for presuppositions like the one the *again* presupposition is at issue in this section the matter doesn't seem to have been given much attention.

Nevertheless, the issues can be more complex here. Intuitively, accommodations like the one that imposes itself on us when we are confronted with (4.47.b) without further context – that J Bills birthday comes after Mary's birthday – are 'minimal' accommodations, which enable the interpreter to make optimal use of contextual information that *is* present in the locally available context (such as in (4.47.b) the information provided by the conditional's antecedent). But how can we give precise content to this notion of 'minimal accommodation needed to complete the explicit contextual information'? To the best of my knowledge this question still is wide open. If I am right, it is perhaps the main remaining challenges for a systematic account of presupposition accommodation, and one of the main challenges for presupposition theory in general.

(4.47.d) shows that the picture which has been emerging from the examples we have discussed in this section has to be refined even further. The first sentence of this discourse does not verify the presupposition that there was a pizza eating event before Bill's birthday. But (4.47.d) is nevertheless acceptable, because the second sentence can be interpreted as generating the presupposition that there was an earlier 'event' of no pizza eating – just as (4.20.b) in Section 4.2.2 gave rise to an interpretation involving the presupposition of an earlier 'event of no rain'. Since this is a possible interpretation of the second sentence of (4.47.d) in the context provided by its first sentence, it is curious that this possibility doesn't pop up when the sentence is offered on its own, as in (4.47.b). Earlier in this section I observed in connection with (4.42.b) that this sentence pair leaves the interpreter with the impression that the context must contain information about some further pizza eating event, which does precede Bill's birthday. The possibility suggested by (4.47.d), that the context contains information about some earlier 'event' of no pizza eating, doesn't seem prominent enough to come to mind. This apparent asymmetry between the two possible interpretations of the second sentence that (4.47.d) shares with (4.42.b) is another matter that needs to be looked at more closely.

This is a good point to recall what was said about accommodation in Section 4.1.5. There I endorsed Beaver's perspective according to which presupposition accommodation is something that an interpreter is called upon to do when a presupposition cannot be verified on the basis of the context that he takes to be available for this purpose. He concludes on the basis of this that his context cannot be the one the speaker is assuming and tries to reconstruct what the context that the speaker must be assuming could be like. That is a kind of abductive process: try to find the best explanation for why the speaker expressed herself in a way that carries the given presupposition. The accommodations that we have been talking about in this section can also be seen as abductions of a certain kind: find a 'minimal' enrichment of the explicitly given context which makes verification of the presupposition at issue possible. But these abductions seem to be driven more by formal, inference-related principles than by speculation about the state of information of the speaker (and perhaps about her the state of mind more generally). A general theory of presupposition accommodation will have to make room for both of these two kinds of abductive guidelines, the abstract information-theoretical principles and the more *ad feminam* ones. But saying this much and no more is just a way of saying that such a theory is still to be developed.

That accommodations prompted by the need for the justification of certain presuppositions can come across as entailments from the sentences in which these presuppositions are generated is not an isolated phenomenon, instantiated by a few cleverly thought out examples like the ones of this section. Kripke's lecture and the paper (Kripke 2009) contain enough examples to show that this isn't so. A somewhat different kind of illustration of the same point is the following example from (Kamp 2001*b*).

(4.49)John gave the workers a generous tip. One thanked him. The other one left without saying a word.

When you are asked how many workers this little discourse is talking about, the answer seems plain: Two. Trivial as it may be to draw this conclusion, reconstructing in detail what is involved in drawing it isn't altogether trivial. This is in large part because true inference is based on interacting presuppositions that ar due to different triggers: (i) the definite description *the workers* in the first sentence, (ii) *one* in the second sentence and (iii) the definite description *the other one* in the third sentence and its constituents (iv) *one* and (v) *other*. In fact, (4.49) is just one instance of a more general paradigm. Other instances can be obtained by varying its different presupposition triggers. For instance we could replace *one* in the second sentence by *two*, or we could make such a substitution for *one* in the third sentence, or we could that for both occurrences of *one* athe same time; or we could change the definite description *the other one*into the indefinite description *another one*. it is easy to verify that each of the three-sentence texts that can be obtained in these ways generates its own complete or partial answer to the question: 'How many workers?'. A correct identification of the presupposition that are generated by each of the triggers involved enables us to predict all these answers.

Examples of this kind also illustrate an observation about presupposition that has often been made: The considered use of presupposition carrying expressions is often an effective way of getting information across to your audience 'through the back door'. By making them accommodate presuppositions in certain ways you get them to adopt the accommodated information as information you want to get to them.

I conclude this section with another observation that goes back to Kripke. Perhaps the most famous example from his lecture and later paper is the statement (4.50).

(4.50)John too is going to have dinner in NewYork City tonight.

The word *too* in this sentence carries the presupposition that one or more persons other than John will have dinner tonight in New York. As Kripke rightly observes, we all know that millions of people are having dinner in New York every night, and so in particular that here will be lots and lots of people other than John who will have dinner tonight in New York. But this knowledge cannot be used to justify the *too* presupposition of (4.50) when the statement is made out of the blue. Some of the discussions of this example seem to suggest that it is the purely existential, non-specific nature of this general knowledge that prevents it from verifying the *too* presupposition in an out of the blue use of (4.50). But this cannot be right, for the discourse in (4.51) is perfectly acceptable.

(4.51) Every night lots and lots of people are having dinner in New York. And tomorrow John too is going to have dinner there.

There is no need for the context to have information about some particular person or persons other than John going to have dinner tomorrow night in NYC. What matters is that the context must contain explicit information from which it can be inferred that people other than John who are going to have dinner in NYC tomorrow. What matters is the 'contain explicit information in this last statement. This is the case when, as in (4.51), such information is introduced explicitly into the discourse context. Knowledge of the information as such isn't good enough, no matter how widely shred or how deeply entrenched.

4.2 Identification Presuppositions of Definite Noun Phrases

When we switched from Top Down to Bottom Up DRS construction, one important advantage was dropped. This was the treatment of anaphoric pronouns and more particularly of 'donkey pronouns'. One motive for the design of the Top Down construction algorithms was the way they can handle anaphoric pronouns: at the point where the algorithm reaches an anaphoric pronoun a dref for the intended antecedent of the pronoun must have been introduced into the DRS that is being built. If and only if a dref for this antecedent is available at that point – in an accessible position, so that interpretation of the pronoun can make use of it – is the pronoun interpretable as anaphoric to this antecedent.[12] When DRSs are built Bottom Up, the possibility for this kind of on-line processing account of pronominal anaphora is lost. When you build a DRS bottom up you will usually get to the pronoun before you get to its antecedent and so no dref for the antecedent will yet

[12] As we saw in our earlier discussion of different types of donkey sentences, the Top Down approach to donkey anaphora is not without its problems. But at least it seems a good start and a reasonably good approximation of the facts.

be available. Therefore pronominal anaphora now has to be dealt with in a quite different way.

It has been announced more than once what is to replace the on-line treatment of pronouns by Top Down construction algorithms: pronouns will form now on be treated, like other definite noun phrases, as triggers of so-called *identification presuppositions* – presuppositions whose resolution requires finding an antecedent for the pronouns that trigger them. It is part of this presuppositon-based strategy that resolution of these identification presuppositions may wait until the preliminary sentence representation has been fully constructed. At that point drefs representing the grammatically possible antecedents for given pronoun will have been introduced.

Time has come to make good on the promise that pronominal anaphora can be accounted for along these lines. But in fact, that promise wasn't just the promise of a viable treatment for anaphoric pronouns. The promise was made in relation to other definite noun phrases as well, and in particular in relation to proper names, the non-quantificational DPs which have occurred abundantly in the sample sentences and discourses that we have been using for illustrative purposes in Section 3 and the preceding subsections of Section 4. As I have stressed repeatedly, our treatment of proper names up to now has always been a provisional one. The only aspect of our treatment of proper names, both before and after the shift to Bottom Up construction, that has merit was the rule that the referential argument of a proper name should always be added to the Universe of the main DRS, giving proper names always scope over other scope bearing elements of the sentences in which they occur. We will see how our new presupposition-based treatment of proper names vouchsafes this principle.[13] That proper names outscope all other scope bearing elements of the sentences in which they occur is intuitively right. But it is right by virtue iof following from the intuitive principle that the reference of a proper name is fixed 'outside of the given discourse'. A sentence containing a name is understood as saying something about the referent of the name and which entity that referent is has been fixed by a naming convention adopted at some point within the speech community to which the speaker belongs. Unfortunately the general framework we have adopted in these Notes isn't rich enough to do full justice to this 'social' aspect of the introduction and use of names as 'labels' of their referents.

[13]It has been argued that this principle, according to which the representing dref for a proper name should always be added to the Universe of the main DRS, doesn't hold for all cases. See (Geurts 1997). More about such cases in footnote 18.

Our presupposition-based treatment of names will not change this. It will continue to the be the case that our sentence and discourse representations represent the contributions of proper names as wide scope existential quantifiers. For instance, the algorithm developed in Section 3 assigns a sentence like 'Mary slept' in (4.52.a) a DRS like (4.52.b) and that will also be the DRS assigned to (4.52.a) by the new construction method. (4.52.b) expresses the truth conditions that there is someone by the name 'Mary' who slept at some time before the utterance time. That is the best we can do about the truth-conditional contributions of proper names within our framework. [14]

(4.52)a. Mary slept.

b.
$$\begin{array}{|c|}\hline t \quad e \quad x \\ \hline t \prec n \quad e \subseteq t \\ \text{Named}(x, Mary) \\ e\text{: sleep'}(x) \\ \hline \end{array}$$

Giving rise to propositions that are about some particular thing or things – or to *singular propositions* as the philosopher's term has it – isn't a unique feature of sentences containing proper names. We find this for each of the five types of so-called *definite* noun phrases that are found in English:

(4.53)a. Proper names

b. Third person pronouns

[14] A proper account of the truth conditions of sentences with proper names, as expressing propositions about the entities to which the names occurring in the sentence refer, requires a richer setting, in which it is possible to talk in a formally explicit way about contexts in which the referents of proper names are fixed. In such a setting the dref representing a (DP which consists just of a) proper name can be linked via the context to some particular referent and the content of the DRS containing the dref as the attribution of a certain property to this referent. One setting in which utterance context can link drefs in logical forms to particular referents in the world or situation a sentence is used to talk about is a communication-theoretic setting in which speaker and interpreter each have a set of 'entity representations' in which names are linked to referents. The speaker's use of 'Mary' depends on her relying on an entity representation of hers which links the name 'Mary' to some particular individual of that name and true task of the interpreter is then to find a matching entity representation in his entity representation set (i.e. one which links 'Mary' to the same referent). For details see (Kamp 2015). More modest frameworks, in which fewer assumptions are made about the mental states of the communication participants, but which nevertheless allow for a more satisfactory account of the referential role of proper names could also be adopted.

c. First and second person pronouns

d. Simple and Complex demonstratives, i.e. phrases beginning with the determiners *this*, *that*, *these* and *those*, either followed by an NP (complex) or not followed by one (simple).

e. Definite descriptions, which in English always begin with the determiner *the*.

Third person pronouns can give rise to singular propositions either when they are anaphoric to some antecedent that refers to some particular thing, as when they are anaphoric to a proper name, or else when they are used *deictically*. Deictic uses of DPs are those where the DP points at some referent in the perceptually accessible environment. Such uses are especially common for complex demonstratives. You whisper to me: "That man near the fireplace is the cousin of my wife." looking discretely in the man's direction to make sure that I get who you mean.[15] In this example the deictically used DP is the demonstrative phrase *that man near the fireplace*. But pronouns too can be used deictically, as when you say to me, pointing at a retain man at a 'council members only' meeting: "He shouldn't be in here". And definite descriptions can be used in a similar manner too, as when you say to me: "Just put the book on the table", referring to the book in my hand and the only table in the room.

First and second person pronouns also can give rise to singular propositional content. But here different semantic mechanisms are involved. In fact, first and second person pronouns always give rise to singular propositional content. An elegant and convincing argumentation for this claim can be found in the work of David Kaplan, especially in his (Kaplan 1989). (This paper was written and circulated in 1970, but did not appear in print until 1989.) Moreover, first and second person pronouns are special not only in that they invariably lead to singular content, but also in the kind of singular content that they give rise to. When a speaker uses the first person pronoun *I*, the thought she expresses isn't just one whose content is a singular proposition about herself; the thought takes the form of attributing to yourself the property of having a headache. You can have thoughts whose contents are singular proposition about yourself, but without being aware that you yourself are the person that the proposition is about. For example, someone

[15]The simple demonstratives *this* and *that* can also be used deictically but less happily. "I want this", pointing at a particular toy in a toy shop, is an utterance common enough in young children. But a more 'adult' way to say this would be "I want/ would like to have this one".

may play recording of a voice to you and your reaction may be that ‘that person has a peculiar voice indeed,quite unlike anything you can remember ever having heard. That thought is plausibly construed as a thought about the person whose voice you are listening to, a thought whose content isa singular proposition about that person. As it turns out, it is you own voice that they have been playing to you. but you don’t know that, your thought is a singular thought about you, but it is not a *de se* thought. It turns into a *de se* thought only when they tell you that it was your own voice you were listening to and you believe them. Until that point you could only have expressed your reaction by saying something like ‘That person has a funny voice’. Only when you have come to realize that the voice is your own, can you express your opinion by saying ‘I have got a funny voice’, and add perhaps ’I never knew that’.

It is not hard to see that we can use the pronoun *I* to express a thought only when that thought is a thought *de se*. But it is thought *de se* only for ourselves. For the recipient of a sentence containing a first person pronoun the proposition expressed by is not *de se*, but rather *de re*. For him the speaker is an external entity, so he will have to interpret the sentence as expressing a content about this external entity. The content will still be a singular proposition and it will be about the same entity as the thought the speaker has expressed, But for the recipient the content is *de re* and not *de se*. With the second person pronoun the sides are reversed. The thought that a speaker expresses when she uses *you* is an attribution to the addressee, an external entity for the speaker, and thus a proposition *de re*. But the use of *you* is an instruction to the addressee to interpret the utterance as an attribution to his self: his interpretation should take the form of a proposition *de se*.

A these uses of definite DPs which give rise to singular propositions require a richer framework, in which it is possible to talk systematically about the communication-theoretic roles of the DPs in question and also about the causal relations that some of them bear to the mental representations of speaker and interpreter. For lack of such a framework none of them will be discussed in what follows here, except proper names and those only because what has been said about them up to now is badly in need of some correction.[16]

[16]For more about the communication-theoretic and mental representation aspects of these problems see, besides the already mentioned (Kamp 2015) also (Kamp 1990), (Kamp 2001,2011) and the forthcoming English original of (Kamp 2003).

This removes most items from the agenda for a systematic account of the uses of definite noun phrases. What remains are, besides proper names, the non-deictic uses of third person pronouns and of definite descriptions. (Demonstrative DPs can also be used non-deictically, but those uses ar subject to quite special constraints, which are hard to explain without taking their decitic uses into account as well. So it is better to set these aside as well here.) The next three subsections deal with proper names, anaphoric pronouns and some of the uses of definite descriptions.

4.2.1 Proper Names

We have already touched on the main problem that one is facing when trying to account in a framework like ours for the semantic contributions of proper names. We lack the formal means of linking uses of proper names t particular entities that are fixed by the context. This problem is not particular to our set-up, but inherent to more or less all current formal approaches to natural language semantics. The goal of these approaches it is to describe the possible interpretations of the expression of natural language expressions in an abstract setting like the one we have been assuming in these Notes – a setting in which the social foundations and psychological dimensions of language use are set aside. What enables a proper name to function as the name of some particular entity in the uses that is made of it by a given speech community is a correlation between name and referent that must have been established at some point in the history of the language – by some form of 'baptism' as the matter was first put by Kripke in (Kripke 1980). A baptism can be thought of as a kind of pact between some small group of members of the speech community to use the name to refer to the given referent. This pact can then spread to other members of the speech community through a process that might be described as one of transferring referential intentions from one speaker, who already had the capacity to entertain these intentions, to another, who until that point did not have that capacity. None of this can be adequately formalized in the framework we have adopted. But even so, let me say a little more about this picture of how names work, for that will make it easier to see how the formal treatment proposed here is inspired by it.

Here, then, is a little more of the story one would want to be able to tell in more formal terms about what is involved in the use of proper names in verbal communication. When a speaker S resolves to use a name N as part of what she wants to say in order to refer to an entity r she must have some kind of representation of **r** and know that N is used as a name of the entity of which this is the representation. The decision to choose N to refer to r constitutes S's *referential intention*. The task of the recipient H of S's utterance is to capture this referential intention of S. Here we have to distinguish between two cases. The first is that where H already has a representation of $\mathbf{r}$, knows that N is used as a name for the referent of that representation and uses it to interpret S's use of N on the given occasion. The second case is that where H does not have such a representation for $\mathbf{r}$ of which he can make use when interpreting S's utterance of N. In this case the best H can do is to accommodate, by adopting a new representation for r, as the entity that S has just referred to by using N.[17]

In our approach what matters is interpretation. Furthermore, since reference identification is now a matter of the resolution of identification presuppositions, it is in the resolution of the presuppositions triggered by proper names that the decisive connection between name and referent will have to be made. It is precisely at this point, where these identification presuppositions must be resolved, that our framework prohibits us from telling the story that really ought be told. All that it allows us to say is that there exists some entity that is named N and that the sentence is about. Formally speaking this is no improvement over the wide scope existential account of proper names with which we have been making do up to now. But at least the presupposition-based treatment shows more clearly how it could be made into a more satisfactory account when embedded within a richer framework than the one we are using.

In our presupposition-based approach to definite DPs, the differences between the various uses of definite DPs manifest themselves as differences in their identification presuppositions. There are two ways in which these differences can show up: in the representational form of the presupposition and in the principles that govern its resolution. The representational forms we will assume for the different definite DP uses will not look very different from each other, and they give no clear indication of the often very different ways in which they are to be resolved. The mode of resolution of an identification presupposition is a function of the DP use that triggers the presupposition. This is information that must be recoverable from the representation of the presupposition, and since it cannot be recovered form the presupposition's representational form, it has to be encoded in some way. We will assume that this is done in the form of a subscript to the presupposition's representation. In particular, the subscript we will use to indicate that the presupposition

[17] This is a much simplified version of the story. For more details see (**?**).

was triggered by a (standardly used) proper name will be $_{pr.na.}$.
Resolution of the identification presupposition of a definite DP always takes the form of identifying a referent or value for the discourse referent that represents the DP in the preliminary sentence representation. This dref must be a constituent of the presupposition representation and it must be singled out from any other drefs that may occur in the Universe of this representation. We do this by placing a question mark behind the dref. For instance, the presupposition triggered by the name *Mary* of (4.52.a) is as in (4.54).

(4.54) $$\begin{array}{|c|}\hline x? \\ \hline \text{Named}(x,N) \\ \hline \end{array}_{pr.na.}$$

Presuppositions whose representations contain such a question-marked dref we call *referential presuppositions*. The resolution of a referential presupposition always takes the form of finding a 'referent' for the question-marked dref which satisfies the Conditions in the Condition Set of the DRS that represents the presupposition. (I have placed 'referent' in scare quotes because in some cases the value actually plays the part of a bound variable in the larger DRS. Cases of this sort are common when the triggering DP is a third person pronoun; see the next section for examples.) Referential presuppositions must be contrasted with *propositional* presuppositions. The *again*-presuppositions we considered in Sections 4.2.1 and 4.2.2 were examples of propositional presuppositions. The resolution of a propositional presupposition consists in showing that the presupposition is entailed by the available contextual information.

The next question we need to address is how the identification presupposition triggered by a proper name DP is computed from the LF of the sentence containing that DP. This raises the preliminary question what internal syntactic structure a DP has (if, that is, it has any internal structure at all). Until now we have dealt with this last question in a rather pragmatic way. DPs with an overt structure, consisting of a Determiner and an NP complement, have been treated as constituents of such a complex form. So far this has been relevant only for non-definite DPS, quantificational DPs like those beginning with *every* and indefinite DPs beginning with *a* or *some*. But these DPs are not among the ones discussed in the present section. The only definite DPs with an openly Det + NP structure are the complex demonstratives and the definite descriptions. And of these it is only the definite descriptions that we will look at. The two definite DP types that we will discuss before we get to that – the proper name DPs that preoccupy us in the present section and the third person pronouns that will be discussed in the next one – have been treated as DPs without internal structure.

We could persist with this policy and some linguists would perhaps be happier with that. But I will pursue a different possibility, according to which proper name DPs also have an internal Det-NP structure. A well-known syntactic argument supporting this assumption is due to Longobardi ((Longobardi 1994)). Longobardi's argument applies directly to Italian, where the well-formedness of DPs like *Roma antica* vs. *l'antica Roma* is best explained by assuming that DPs containing proper names have a Det position, which is sometimes filled by an article and in other cases, when no overt determiner is present, gets filled by the name, which moves to Det from its base position inside NP.

Admittedly it is not obvious that the conclusion from this argument can be extended without further ado to other languages. But I see the argument nevertheless as at least suggestive in this wider setting and assume that English proper name DPs also have Det + NP structure. The Det of such a DP is always morphologically empty. The DP's NP is assumed to have the form of an empty head N with a kind of PP adjunct, whose 'preposition' is the 2-place predicate*Named* and whose DP consists just of the name. *Named* expresses a relation between its referential argument – the referent of the DP as a whole – and the referent of the DP it governs. The referent of the governed DP is the name *N* itself, as phonologically or orthographically identified expression. Thus as constituent of the DP '*N*' functions self-referentially (or 'autonymously', as technical terminology has it): the referent of '*N*' is '*N*'.

According to these assumptions the DP whose only overt constituent is the name '*N*' has the following form:

(4.55)

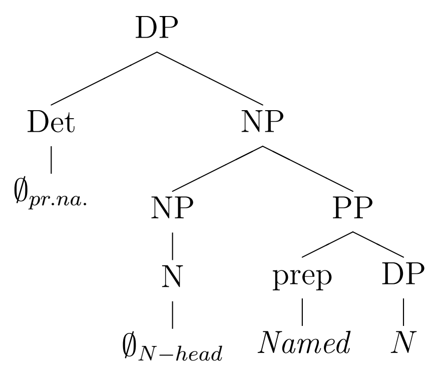

This may look like a proliferation of constituents where arguably there is only one. But let us not get sidetracked by this worry and see what this analysis commits us to by way of lexical entries. (4.56) gives all the entrees we need in order to build the semantic representation for (4.55). For compactness' sake only the semantic parts of those entries are shown.

(4.56)a. Lexical entry for the complement DP of *Named*:

$$< y_{ref} \mid \begin{array}{|c|}\hline \quad \\ \hline y = \text{`}N\text{'} \\ \hline \end{array} >$$

b. Lexical entry for the preposition *Named*:

$$< x_{ref} \mid \begin{array}{|c|}\hline \quad \\ \hline \text{Named}(x,\underline{y}) \\ \hline \end{array} >$$

c. Lexical entry for the empty N head $\emptyset_{N-head}$:

$$< z_{ref} \mid \begin{array}{|c|}\hline \\ \hline \\ \hline \end{array} >$$

d. Lexical entry for the empty Determiner $\emptyset_{pr.na.}$:

$$< x_{ref}, y \mid \begin{array}{|c|}\hline \quad \\ \hline \text{Named}(x,y) \\ y = \text{`}N' \\ \hline \end{array} > \rightsquigarrow < x_{ref}, y \mid < \{ \begin{array}{|c|}\hline x? \\ \hline \text{Named}(x,y) \\ y = \text{`}N' \\ \hline \end{array} \}, \begin{array}{|c|}\hline \\ \hline \\ \hline \end{array} >>$$

The only entry from this list that may need a comment is the entry in (4.56.d) for the Determiner feature $\emptyset_{pr.na.}$. Like for the determiners considered earlier, the semantics for $\emptyset_{pr.na.}$ is treated as an operator, which transforms the representation of the NP governed by the Determiner node, which is fed to the operator as input, into the representation of the DP. (Recall in this connection for instance the entry we gave for *every* in Section 3.9.1.) What is new about the entry for $\emptyset_{pr.na.}$ is that it transforms the non-presuppositional input it gets into an output in which that non-presuppositional input has been turned into a presupposition. In the course of this process the non-presuppositional input DRS is deprived of all its content, so that only its empty shell remains.

A second feature of this entry for $\emptyset_{pr.na.}$ is that while it has the general form of a schematic specification, with N as schematic letter, its specification of the input is very detailed. This is possible because Determiner nodes with the feature $\emptyset_{pr.na.}$ only occur as part of DPs whose NP constituent has the form of the input in (4.56.d). And it is desirable to specify the input in this detailed way, in order to make explicit that it is only to such inputs that the operator in (4.56.d) is ever applied.

Putting the pieces from (4.56) together by use of by now familiar composition rules gives us for the PP, the NP and the DP nodes of (4.55) the semantic representations in (4.57.a,b,c).

(4.57)a. PP:

$$< x_{ref}, y \mid \begin{array}{|c|}\hline \\ \hline \text{Named}(x,y) \\ \text{y} = \text{`}N\text{'} \\ \hline \end{array} >$$

b. Upper NP:

$$< z_{ref}, y \mid \begin{array}{|c|}\hline \\ \hline \text{Named}(z,y) \\ \text{y} = \text{`}N\text{'} \\ \hline \end{array} >$$

c. DP:

$$\left\langle z_{ref}, y \mid \left\langle \left\{ \begin{array}{|c|}\hline z? \\ \hline \text{Named}(z,y) \\ y = \text{`}N' \\ \hline \end{array} \right\}, \begin{array}{|c|}\hline \\ \hline \\ \hline \end{array} \right\rangle \right\rangle$$

The dref y is playing a redundant part in this last representation, which can be simplified to (4.58).

$$(4.58)\left\langle z_{ref} \mid \left\langle \left\{ \begin{array}{|c|}\hline z? \\ \hline \text{Named}(z,\text{`}N\text{'}) \\ \hline \end{array} \right\}, \begin{array}{|c|}\hline \\ \hline \\ \hline \end{array} \right\rangle \right\rangle$$

Those who see no reason for assuming that proper name DPs have internal syntactic structure can take (4.58). as the lexical entry for the proper name

N, which after lexical insertion gets passed up unchanged to become the semantic representation of the DP consisting just of N.

When a DP-representation like that in (4.58) is combined with the representation of its sister node in a predication configuration, its referential argument is inserted into the coindexed argument slot in the DP's sister representation. As with the representations of indefinite DPs the preferential argument is kept in the store. But the difference with indefinites is that the referential argument of a definite DP will be retained in their store at the point when the Comp node of the clause triggers transfer of the drefs that are still in store to appropriate DRS Universes. A dref that occurs with a question mark in a referential presupposition representation will be transferred from its store only when the presupposition of which it is the question-marked dref has been resolved.

To see how this works let us at long last go through the DRS construction for the sentence 'Mary slept', mentioned under (4.52) in the introduction to Section 4.3. The LF for this sentence is shown in (4.59).

(4.59)

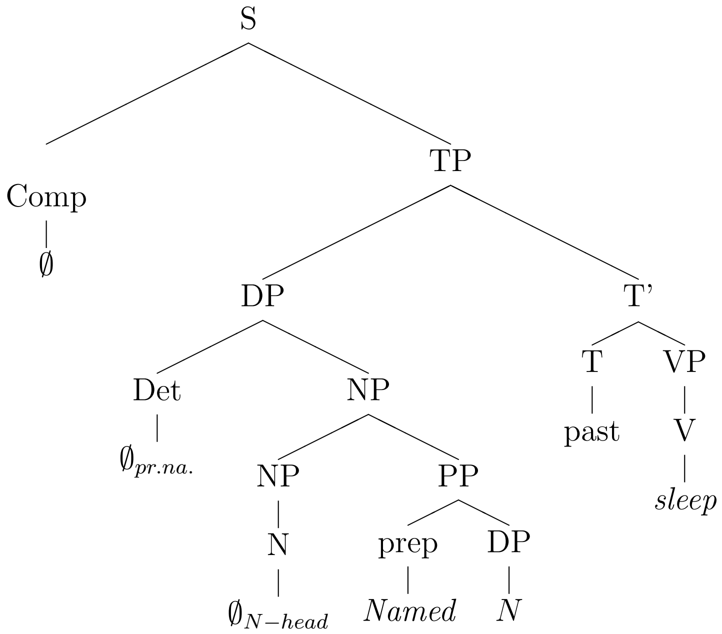

(4.60) is the result of computing the representation for the proper name DP *Mary* and that for the T' node (which involves construction steps all of which are familiar).

(4.60)

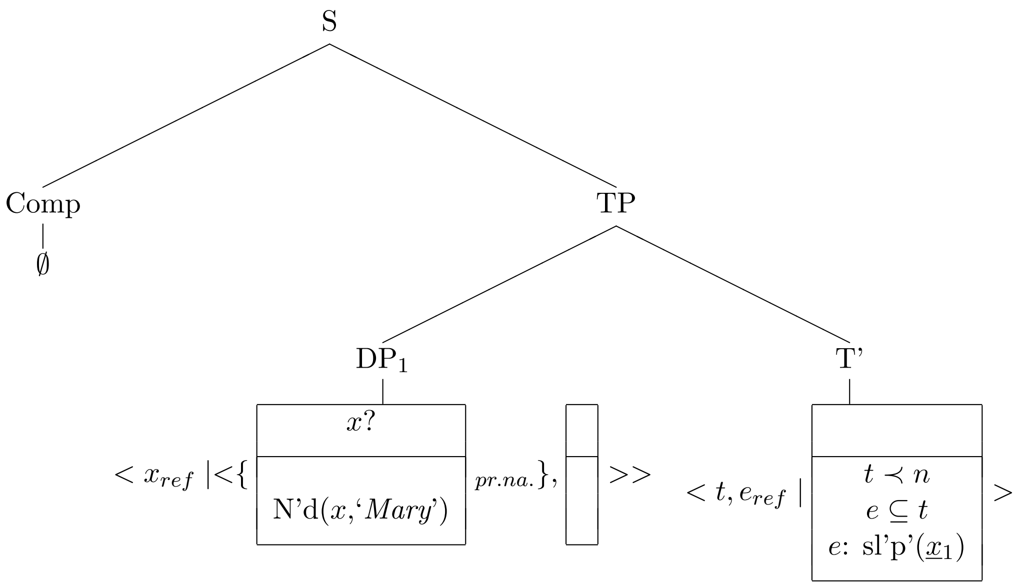

Passing from (4.60) to the construction of the TP representation follows largely the principles of argument insertion that are familiar from our treatment of quantifying and indefinite DPs. But note that the non-presuppositional part of the resulting representation is just the non-representational part of the T representation, except that the argument slot $\underline{x}$ has been replaced by the referential argument x of the argument phrase. This is because the non-presuppositional part of the DP is the empty DRS, so the merge of it with any other DRS is just that other DRS. Furthermore, the presupposition set is in this case just the singleton set containing the identification presupposition for the proper name *Mary* as its only member, since in the present case the presupposition set of the VP representation is empty.

(4.61)

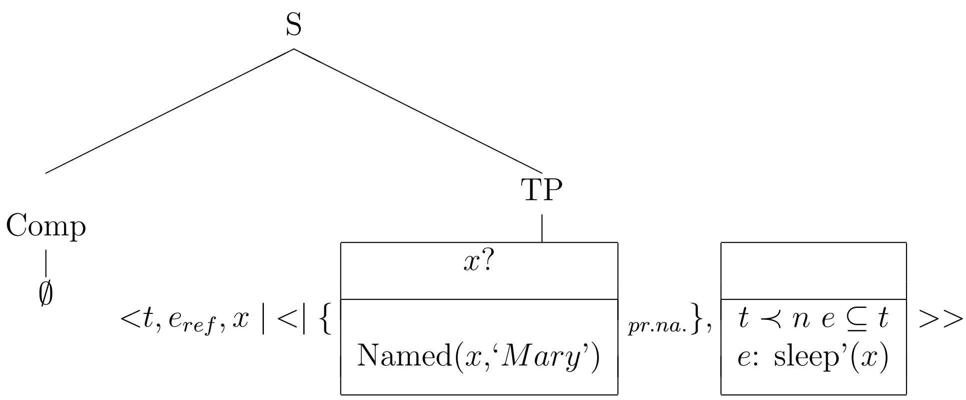

As always for main clause Comp nodes, the Comp node of (4.59) carries the instruction that drefs which are still occurring in a store should be transferred to a suitable DRS Universe. In the present instance this is once again straightforward since there is only one DRS Universe to which drefs can be moved. However, as already noted above, the global instruction for store clearing that have been associating with main clause Comp nodes must now be qualified: The instruction does *not* apply to drefs that occur question-marked in referential presuppositions that have not yet been justified. For the time being these are kept in the stores where they were. Thus the effect of store clearing in the case of (4.61) leads to the representation in (4.62).

(4.62)

$$< x_{ref} \mid \{ \begin{array}{|c|}\hline x? \\ \hline \text{Named}(x, Mary) \\ \hline \end{array}_{pr.na.} \}, \begin{array}{|c|}\hline t \quad e \\ \hline t \prec n \quad e \subseteq t \\ e\text{: sleep'}(x) \\ \hline \end{array} >$$

(4.62) is the preliminary representation for the sentence 'Mary slept'. But how do we get from this to the final representation of the sentence? This is where our framework doesn't enable us to say very much. It offers no way to represent the information that interpreters typically make use of to resolve proper name presuppositions. We have no way of distinguishing between cases where the presupposition is directly verified by the interpreter's antecedent knowledge about what the name refers to and cases where the interpreter lacks this knowledge and can do no better than accommodate the presupposition. The other, arguably more serious shortcoming of our framework is that it doesn't enable use to represent the content of a sentence like 'Mary slept' as a singular proposition about Mary. The best we can do, it was argued above, is to move the material from the resolved or accommodated presupposition to the main DRS – the dref to its Universe and the 'Named'-Condition to its Condition Set.

$$(4.63) \quad \begin{array}{|c|}\hline t \quad e \quad x \\ \hline t \prec n \quad e \subseteq t \\ \text{Named}(x, Mary) \\ e\text{: sleep'}(x) \\ \hline \end{array}$$

The present example doesn't provide a telling illustration of te principle that the material from a proper name presupposition must be transferred to the main DRS, since there is only one non-presuppositional DRS that the material can be moved to. For the next example does provide such an illustration. Consider the sentence in (4.64). The LF for (4.64) is given in (4.65).

(4.64) If Mary slept, John didn't sleep.

(4.65)

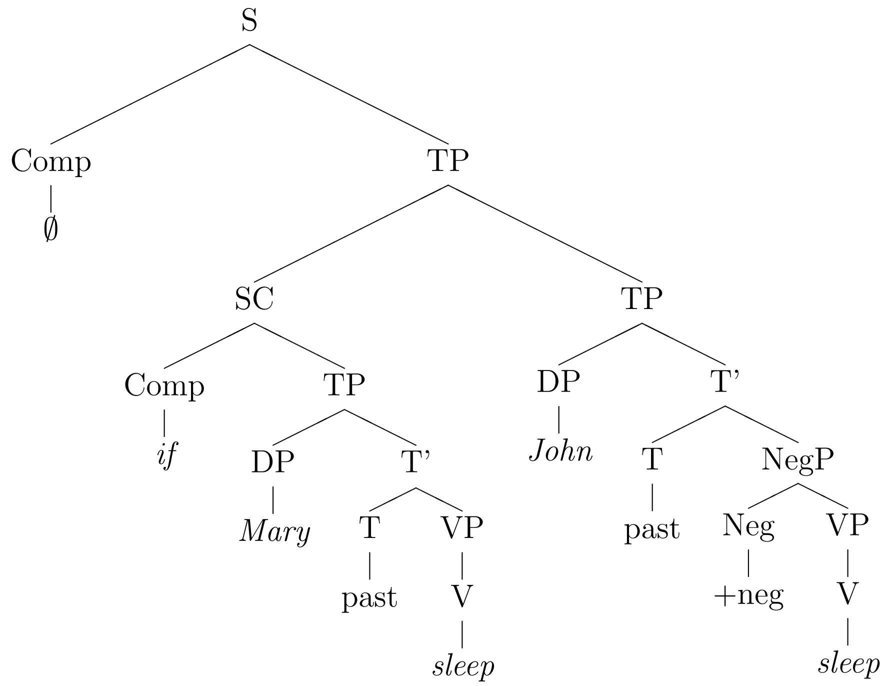

The representations of the TP of the *if*-clause and the lower TP of the main clause are given in (4.66) and (4.67).

(4.66) $<t, e_{ref}, x \mid <\{\begin{array}{|c|}\hline x? \\ \hline \text{Named}(x,\text{'}Mary\text{'}) \\ \hline \end{array}_{pr.na.}\}, \begin{array}{|c|}\hline \\ \hline t \prec n \ e \subseteq t \\ e\text{: sleep'}(x) \\ \hline \end{array} >>$

(4.67)

$$<t, e''_{ref}, y \mid < \{\mid \begin{array}{|c|}\hline y? \\ \hline \text{Named}(y,\text{‘}John\text{’}) \\ \hline \end{array}\ _{pr.na.}\}, \begin{array}{|c|}\hline \quad \\ \hline t' \prec n\ dur(e'') = t' \\ \neg \begin{array}{|c|}\hline e' \\ \hline e'\text{: sleep'}(y) \\ e' \subseteq e'' \\ \hline \end{array} \\ \hline \end{array} >>$$

As we saw in Section 3.10.2 (the section on conditionals), the representation of the higher TP of the main clause is as in (4.68).

(4.68)

$$\begin{array}{|c|}\hline \quad \\ \hline \begin{array}{|c|}\hline <t, e_{ref}, x \mid < \{\mid \begin{array}{|c|}\hline x? \\ \hline \text{Named}(x,\text{‘}Mary\text{’}) \\ \hline \end{array}\ _{pr.na.}\}, \begin{array}{|c|}\hline \quad \\ \hline t \prec n\ e \subseteq t \\ e\text{: sleep'}(x) \\ \hline \end{array} >> \\ \hline \end{array} \Rightarrow \\ \begin{array}{|c|}\hline <t, e_{ref}, y \mid < \{\mid \begin{array}{|c|}\hline y? \\ \hline \text{Named}(y,John) \\ \hline \end{array}\ _{pr.na.}\}, \begin{array}{|c|}\hline \quad \\ \hline t' \prec n\ dur(e'') = t' \\ \neg \begin{array}{|c|}\hline e' \\ \hline e'\text{: sleep'}(y) \\ e' \subseteq e'' \\ \hline \end{array} \\ \hline \end{array} >> \\ \hline \end{array} \\ \hline \end{array}$$

The store clearing dictated by the Comp node of the main clause now leads to transfer of drefs from stores to suitable Universes of non-presuppositional

DRSs. Here we make the default assumption that we have followed so far, that a dref in store should be transferred to the Universe of the non-presuppositonal DRS to its right. (But, as noted above, this does not affect the question-marked drefs, which are kept in their stores until their presuppositions have been dealt with.) Applying this default procedure leads to the preliminary DRS in (4.69).

(4.69)
$$
\begin{array}{|c|}
\hline
\\
\hline
< x \mid <\{ \begin{array}{|c|}\hline \\ \hline \text{Named}(x,\text{‘}Mary\text{’}) \\ \hline \end{array}_{\ pr.na.}\}, \begin{array}{|c|}\hline t \quad e \\ \hline t \prec n \quad e \subseteq t \\ e\text{: sleep’}(x) \\ \hline \end{array} >> \quad \Rightarrow \\
\\
< y \mid <\{ \begin{array}{|c|}\hline y? \\ \hline \text{Named}(y,\text{‘}John\text{’}) \\ \hline \end{array}_{\ pr.na.}\}, \begin{array}{|c|}\hline t' \quad e'' \\ \hline t' \prec n \quad dur(e'') = t' \\ \neg \begin{array}{|c|}\hline e' \\ \hline e'\text{: sleep’}(y) \\ e' \subseteq e'' \\ \hline \end{array} \\ \hline \end{array} >> \\
\hline
\end{array}
$$

Resolution of the two proper name presuppositions now leads to the transfer of their drefs and conditions into the Universe and Condition Set of the main DRS, as shown in (4.70).

(4.70)

$$\begin{array}{|c|} \hline x \quad y \\ \hline \text{Named}(x,\text{‘}Mary\text{’}) \quad \text{Named}(y,\text{‘}John\text{’}) \\ \begin{array}{|c|}\hline t \quad e \\ \hline t \prec n \quad e \subseteq t \\ e\text{: sleep'}(x) \\ \hline \end{array} \Rightarrow \begin{array}{|c|} \hline t' \quad e'' \\ \hline t' \prec n \quad dur(e'') = t' \\ \neg \begin{array}{|c|} \hline e' \\ \hline e'\text{: sleep'}(y) \\ e' \subseteq e'' \\ \hline \end{array} \\ \hline \end{array} \\ \hline \end{array}$$

4.2.1.1 A note on the claim that proper names are predicates

The treatment of proper names presented above might be perceived by some as a version of the 'Descriptive Theory of Names'. 'The Descriptive Theory of Names' has become the label for a range of different accounts of proper names according to which a name makes its contributions to the semantic content of the sentences in which it occurs via some descriptive content – some predicate that is supposed to be true of the name's referent – that is associated with the name. In early proposals of this general type, hints of which can be found in the work of Frege and Russell, the descriptions in question where supposed to be those in terms of which the user of a name thinks about the referent and which she may use to identify the referent (i.e. distinguish it from other entities). Of a radically different type are theories in which the descriptive content is something that is ultimately uninformative and also has a flavor of circularity, such as the description 'is named N'. Some years ago a version of this second type was proposed as a kind of revival of 'The descriptive Theory of Names' in (Geurts 1997). In Geurts' account the predicate 'is named N' plays a central part, and one which in the author's view justifies classifying his account as a version of The descriptive Theory of Names. This is not the place for a detailed discussion of Geurts' proposal. (But see the footnote attached to the next paragraph) I mention it as an example of an account of how names work in which the predicate 'is named N' appears to be doing most of the semantic work.

In this last respect Geurts' account resembles the one proposed in these Notes. But although in our account the predicate 'is named N' plays a pivotal role as well, it doesn't' seem right to describe it as an instance of 'The Descriptive Theory of Names'. Why it wouldn't be right is not so easy to see from the presentation in this section. but it emerges clearly from the story for which the presentation in this section is a kind of stand-in. As that story goes, 'is named N' predications are crucial to how a speaker can choose a name N to refer to a given referent r, and also to how an interpreter of her utterance can recover the reference from her use of N. But these predications only play an ancillary part in the referential relations between agents and entities that is part of what it is for an agent a to think about an entity r, an it is these relations between agents and objects of their thoughts that the story takes to be the essence of reference; and the 'is named N' predications are also no more than derivative and instrumental in what the story has to say about the 'naming relation', which holds between a name and the entity that it is a name of. I realize that all of this is rather vague and presupposing things that I do not spell out. But I'll leave things as they are and refer the rear to Geurts' paper and the reference that can be found there as well as, once more, to (Kamp 2015) and the as yet unpublished (Kamp 2001,2011). which cam be found on my webpage.[18]

The thesis that proper names have descriptive content can also take another form, embodied in the slogan that 'names are predicates' ((Graff-Fara 2015)). A feature of English and many (perhaps all) human languages that may encourage this view is that names can be more or less unrestrictedly used as

[18] Geurts connects the case he is making for the role that the Named relation plays in the semantics of proper names with the observation that the use of a name can be 'locally justified', with the Naming relation being central to the name's locally justified interpretation. An example that he gives of this possibility is the following sentence.

(4.71) If some parents decide to christen their daughter 'Bambi', then the Disney Company will sue Bambi's parents.

Here, Geurts argues, the occurrence of the name 'Bambi' in the main clause can be understood as anaphoric to the indefinite DP *their daughter* in the *if*-clause, and this is possible because the *if*-clause provides the information that what the indefinite DP denotes is a girl named 'Bambi'. That proper names can be used in this way is an important observation. (Assuming that such uses are acceptable; some speakers find (4.71) only marginally felicitous.) But the possibility of using a name in the way (4.71) illustrates is an issue that must be distinguished from the role that is played by the Naming relation in the use and interpretation of names. According to the story that we cannot tell here because the right framework is lacking, the role of the predication 'Named(x,'$Bambi$') in the interpretation of the second occurrence of 'Bambi' is not different from the role that Naming predications play in the more common and familiar uses of proper names, in which the name refers non-anaphorically – to some individual that can be recovered from the context in which the name is uttered.

common nouns, as in sentences like those in (4.72).

(4.72)a. There are several Marys in this room.

b. I don't believe you do not know any Mary.

In each of these sentences the semantic contribution of the word *Mary* is that of a common noun and the semantic contribution made by it can be adequately paraphrased as 'person called 'Mary''.

Moreover, in many languages names can be accompanied by the definite article also when they are used in the standard referential way, which in English forbids the use of an article. (In some languages inclusion of the definite article is obligatory in the standard referential use of names; in others it is optional.)

Transformations in the opposite direction are possible too. Common nouns are common choices when we are looking for a new name for something. Familiar examples are *Faith*, *Victor*, *Apple*, *Dartmouth*, *Rio Colorado* or *Valley Forge*. In fact, when it comes to expressions that are already part of the language, nouns seem to be preferred over expressions belonging to other grammatical categories (such as adjectives or verbs). For instance, while *Honesty* is an authenticated name, *Honest* would be strange as a name, for a woman or for anything else.[19]

In order to use a noun as name some kind of baptismal act has to be performed first. Some person or group has to decide to make a noun N into the name of a referent r to get the practice of using N as name for r under way. Using names as nouns is not subject to such a constraint. You can use name N as a predicate whose extension consists of all and only those entities that are named N without any preambles. This mechanism appears to be unrestrictively productive.

Especially the unrestricted availability of using names as common nouns may be seen as supporting the view that names are a kind of predicates to begin with, whether they are ever used as nouns or not. In spite of this suggestive evidence, however, I believe that classifying names as a species of predicates

[19]Admittedly there are exceptions to this. For instance, in Roman antiquity, adjectives like *Primus*, *Tertius*, *Quintus* and so on were commonly used as names to be given to the first, third, fifth,.. son. In principle there is no limit to what strings of letters or phonemes can be used as names. Especially when it comes to naming horses or houses name-giving owners have been known to demonstrate astounding feats of imagination and poor taste.

is a move that should be resisted. The syntactic analysis of proper name DPs we have adopted and the things we have hinted at in our description of the role which the Named relation plays in the use and interpretation of names provide a useful backdrop for a discussion of this matter. First, the effect of the Determiner feature $\emptyset_{pr.na.}$ is to turn the semantic representation of the NP (of which the 'Named(x,'N')' Condition is part) into a presupposition. The part that the Condition 'Named(x,'N')' then subsequently plays in the resolution of this presupposition is markedly different from the role that DRS Conditions play in non-presuppositional representations. Secondly, the logical structure of predications of the form 'Named(–,'N')' is crucially different from the semantics of a common noun with the outward shape of 'N'. For instance, consider the use of *Mary* as common noun. Suppose that, faithful to the convention that we have been following though out these Notes, we represent this noun in our DRS language as *Mary'*. It is true that predications involving *Mary'* – DRS Conditions of the form '$Mary'(x)$' – have the meaning that x is an individual that is named 'Mary'. But the role that the common noun *Mary* plays in these conditions is clearly very different from the one it plays in predications of the form 'Named(x,'$Mary$')', in which '*Mary*' is not a 1-place predicate, but an argument phrase that fills one of the slots of the 2-place predicate '$Named$'. (The role that '*Mary*' plays as argument to the predicate *Named* is special in that it refers 'autonymously': it refers to itself, as the very phonologically or orthographically identified linguistic expression that is part of 'Named(x,'$Mary$') as a well-formed string of our DRS language. But the fact that '*Mary*' plays a different role in 'Named(x,'$Mary$')' from the one that '*Mary'*' plays in '$Mary'(x)$' remains.)

The perspective according to which moving from name to noun involves the transformation of one word into another, rather than being a mere testimony to the possibility of using one and the same word in grammatically different ways, is a natural one when we view name-to-noun transitions against the background of other systematic transformations of words into what seem to be different words. Mechanisms for turning words of one grammatical category into words of another category are ubiquitous in natural language. A range of such mechanisms is found in English, ut other languages have them too; and perhaps the availability of such mechanisms is a kind of linguistic universal. To mention just a couple of examples: (i) Agentive verbs in English can be turned into nouns denoting their agents, as in *walk* $\rightsquigarrow$ *walker*, *interpret* $\rightsquigarrow$ *interpreter*, *control* $\rightsquigarrow$ *controller* etc. (ii) Short (mostly monosyllabic) English activity verbs can be turned into nouns that denote the events that the verbs are used too describe: *walk* (the verb *to walk*) $\rightsquigarrow$ *walk* (noun, as in *a walk*), *laugh* (verb) $\rightsquigarrow$ *laugh* (noun), *buy* (verb) $\rightsquigarrow$ *buy* (noun); (iii) English nouns can be turned into verbs: *house* (noun) $\rightsquigarrow$ *house* (verb), *skate* (noun) $\rightsquigarrow$ *skate* (verb), *garden* (noun) $\rightsquigarrow$ *garden* (verb); (iv) *-like* is a suffix that turns nouns into adjectives which are true of things that resemble things in the extension of the transformed noun in some relevant way, as in *child-like*, *bird-like*, *prison-like*. Of these four transformations (i) and (iv) are morphologically overt – the phonological and orthographical shape of the output word is different from that of the input word. For (ii) and (iii) this is not so. These are 'zero morphology' transformations. In this they resemble, you could say, the transition from names to nouns. That resemblance need not prove anything beyond what it is. For instance, there is no compelling reason why you couldn't hold that the verb *walk* and the noun *walk* are different words, but that when a proper name is used as a common noun, it is still the same word, but with a different syntactic distribution and a correspondingly adjusted meaning (much in the way that plural forms of nouns are subject to different morpho-syntactic constraints from their singular forms and have a semantics that differs accordingly). It is unclear to me what should govern such choices – different words or same word, but subject to different syntactic and semantic constraints. But when you look at the common noun uses of proper names from the more general perspective of derivational morphology, two points become clear: (i) it may not be all that clear what the right choice is in their case,and (ii) The dramatic differences between the grammatical properties of proper names and common nouns may well tip the balance in favor of the 'two words' choice. (And as indicated above, (ii) is my choice when you force me to make one.)

In our discussion of proper names that are made into common nouns we have assumed that the noun's semantics is given by the the Named relation. Notoriously, this isn't the only way in which the noun can be semantically related to the name. Other uses of proper names as common nouns are those which target one or more of the salient properties of some particular bearer of the name a (typically a well-known one). A classic example is the use of *Napoleon* in a sentence like (4.73.a).

(4.73)a. Fred behaves like a real Napoleon these days. (What's got into him?)

b. Hilary Clinton is the Angela Merkel of America.

c. My lawyer is a real shark/a true butcher.

d. My uncle Freddie wouldn't know how to hurt a fly. He is a true koala.

e. His Napoleonesque behavior is not easy to put up with.

f. Two members of the new government have as pronounced a tendency to merkel as the Chancellor herself.

g. She is a Meg person.

h. She is a Meg sort of person.

i. Three people in this room Meg.

The use of *Napoleon* in (4.73.a) and that of *Angela Merkel* in (4.73.b) should be seen as examples of a certain kind of metaphorical language use.[20] The kind of metaphor represented by these two examples is quite close to certain metaphorical uses of common nouns that are illustrated in (4.73.c,d).

A second feature that sets these common noun uses of proper names apart from those in which the noun *N* is understood as meaning 'named *N*' is that such metaphorical uses of names aren't restricted to common nouns; they can also take the form of an adjective, like *Napoleonesque* in (4.73.e) , or a verb, as *merkel* in (4.73.f). This too is something that such uses of proper names have in common with common nouns, cf. the adjective*butcher-like* and the verb *to butcher*. On the other hand, attempts to use a proper name as an adjective or verb with an 'is named' meaning, as in (4.73.g), (4.73.h) or (4.73.i) just won't work. To the extent that (4.73.g) is grammatical, it doesn't mean that the subject is a person called Meg, but rather that he or she is a person with Meg-like qualities (whatever those might be). Much the same applies to (4.73.h). This seems an acceptable sentence. But it too only has an interpretation according to which the subject has Meg-like qualities, and to presuppose that one or more people called 'Meg' are or are taken to be distinguished by certain salient properties, or that this is generally the case for people called 'Meg'. Finally, for me (4.73.i) is unsalvageable altogether: There is no way to interpret this string as a well-formed sentence, no matter what meaning is assigned to it.

In short, the mechanisms that are involved in these uses of proper names are clearly different from the rule that turns a proper name *N* into a noun *N* with the meaning 'is named '*N*'' – syntactically in that they allow for the formation of words other than nouns and semantically in that the changes

[20](4.73.b) is of fairly recent date. It differs from (4.73.a) in trying to do two things at once; insult the person named and slander the person whose name is used as noun by insinuating that she and the country that elected her as head of government had negative properties that are fabrications or figments of the speaker's imagination. This and other proofs of the speaker's mastery in the art of flatulent innuendo have added a new dimension to the metaphorical use of proper names.

in meaning are much more varied and depend on context and on the users' imagination in ways that are typical of metaphorical language use more generally and that play no part in the interpretation of N as referring to things that have the name N.

We conclude this section on proper names with a point that hasn't got much to do with proper names as such, but that fills a gap in our earlier discussion about the negation test in Section 4.2.2. There we observed that one reason why negation preserves presuppositions is that it leaves the presuppositions of its input representations untouched, operating only on the non-presuppositional part of the input. This was the case illustrated by the example we were looking at. But there is also a second type of case, that where the presupposition is generated outside the scope of a given negation and the representation of the part containing the presupposition is combined with the part that contains the negation at some point higher up in the tree. Sentence (4.64) offers two illustrations of this, involving the identification presupposition triggered by *Mary* and the identification presupposition triggered by *John*. The representation of the identification presupposition associated with *John* is introduced in the process of determining the representation of the subject DP and thus is in place when the subject representation is combined with the T' representation. The presupposition set of the resulting representation is the union of the presupposition sets of the two daughter representations. For the example under consideration this set just consists of the identification presupposition triggered by *John*; for that is the only member of the DP representation, while the presupposition set of the T' representation is empty.

The case of the presupposition triggered by *Mary* in (4.64) is much the same, and different in only one respect. This presupposition is put in place in the course of building the representation for the *if*-clause. When the *if*-clause representation is then combined with the lower TP representation of the main clause in the way we have assumed when constructing the representation for (4.64), then there isn't even a possibility of interaction between negation and presupposition: the presupposition becomes a member of the presupposition set of the representation of the antecedent of the resulting conditional and the negation remains confined to the non-presuppositional part of the representation of the consequent.

The difference between these two situations in which presuppositions are unaffected by negation – the one where the negation 'hops over' the presuppositions that are part of the representation that it receives as input and the one where it has already done its work by the time the result of its application is combined with the representation of another sentence constituent in which the presuppositions have been triggered – can also be observed when we compare the case where a presupposition is triggered within the subject DPs of a negated clause with that in which the presupposition is triggered within the object DP. If the presupposition is generated within the object DP, the negation hops over it; if it is generated within the subject DP then the presupposition and the negation remain separated throughout the construction of the preliminary representation of the sentence. An example is provided by the following pair of sentences, in each of which an *again*-presupposition is generated within a relative clause, which belongs to subject DP in the first sentence and to the direct object DP in the second..

(4.74)a. Five members who defaulted on their membership dues again didn't go to the annual reception.

b. They didn't invite to the annual reception five members who defaulted on their membership dues again.

4.2.2 Third Person Pronouns

The original Top Down version of DRT which we reviewed in Section 2 was motivated in part by the behavior of 3d person pronouns, and more particularly by 'donkey pronouns', both the donkey pronouns which find their antecedents within the sentence in which they themselves occur and those whose antecedents belong to some other sentence in the antecedent discourse. The handle that the Top Down construction method gave us was given up when we switched from Top Down to Bottom Up. And the promise I made at that point was that what was lost through this transition would be regained eventually, but that this would require the integration into the Bottom Up approach of a formally precise treatment of presupposition.

The time has come to make good on this promise. But to do so, we need to be very precise about what the identification presuppositions are like that pronouns trigger and especially about the ways in which these presuppositions are resolved. In fact, as with proper names it is the resolution constraints that carry nearly the entire load of the account. The representational form of the identification presuppositions triggered by 3rd person pronouns conveys almost nothing about what these constraints are – in tho regard they are like the identification presuppositions for proper names. So once again wee need something else to record what type of expression was the trigger of these presuppositions. We adopt the same kind of device this as we did for proper names: a subscript, $_{3d.p.pr}$, is added to the representations of 3d person pronoun presuppositions. Once more, this subscript is to be seen as a shorthand for the complex statement of the principles that govern the resolution of the subscripted presuppositions.

A complete statement of these principles is complicated by the fact that pronouns can not only be used anaphorically, but also deictically. (As when I say to you, pointing surreptitiously at a woman standing in the opposite corned of the room,'She did her Ph. D. with Tarski.') In these Notes deictic uses are set aside, so here we only have to worry about the anaphoric uses. But even just stating the principles involved in anaphoric resolutions of 3rd person singular pronoun presuppositions is a notoriously complex task. Fortunately, our DRT-based framework makes this task somewhat easier than it might have been had we been using a different general framework. This is still true now that we have made the transition to Bottom Up DRS construction. Most of the general architecture of DRT is not affected by this switch. In particular it is no less true now than when DRSs were constructed top down that DRSs serve as global and local contexts; and the preliminary sentence representations we now construct still have the logical structure that determines which parts of a DRS can serve as local discourse contexts for which other parts. Discourse referent accessibility – or, in the terminology introduced in Section 4.2, what constitutes the locally available context – remains essentially unaltered. (We have already seen a number of examples of the role that the locally available context plays in presupposition resolution, in Sections 4.1.5, 4.2.3 and 4.3.1.)

Before we are in a position to say anything about how pronoun presuppositions can be resolved, we have to make explicit what their representations are like and how these are constructed. And that leads us once more to the question how pronouns are syntactically represented. Since the normal occurrences of personal pronouns take the form of argument phrases, and thus of DPs, this question boils down to: What is the syntactic structure of a DP that overtly consists just of a pronoun?

When we addressed the analogous question for the case of proper names in the last section, I mentioned Longobardi's argument that Italian proper name DPs must have a determiner position, which is filled by the name in case no article is used. I took this argument as motivating the assumption that English proper name DPs, have such an internal structure too, with a Det constituent and an NP constituent.

As far as I know, there is no comparable argument to support the assumption that pronoun DPs have a similar internal structure. Nevertheless I will assume that these DPs too consist of a Det position and an NP complement. Here is the general idea behind this assumption. Pronouns, in English and many other languages, form a small logical space of three dimensions. The dimensions are the three so-called ϕ features, person, number and gender. Each of these has more than one value, how many varies from language to language. For English, the feature person has three values, 1st, 2nd and 3rd, number has two values, singular and plural, and gender has three values, feminine, masculine and neuter. (The morphological manifestations of gender are rudimentary; but they are manifest in third person pronouns and that is what matters for us here.) Even if not all ϕ feature value distinctions are explicitly marked on all pronouns – for instance, there is no differentiating gender marking on first and second pronouns – it is nevertheless natural to think of pronouns as each characterized by a triple of values for each of ϕ features, and indeed, this is a widely adopted way of looking at them. For each pronoun the three feature values that characterize it fully determine the semantic contributions that it can make to the different sentences in which it can occur.

The next point to be considered is what functional roles the different feature values play in the meaning contributions that the different pronouns make. Since we are only considering singular pronouns here, as part of our general policy to leave plurals out of the picture because we cannot do everything, the interpretational differences between singular and plural pronouns can be ignored as well. It suffices to stick to the principle that has been implicit in all we have said about singular pronouns so far, viz. that a pronoun always represents a single entity. The two remaining features, person and gender, play rather different roles. For English pronouns gender information is descriptive information about the referent[21] – the kind of information that so

[21] In languages in which common nouns have 'grammatical' gender (i.e. in which gender will manifest itself overtly in the morphology of some or all complex noun phrases with a common noun as lexical head), such as e.g. German or the Romance languages, where the gender marking of a pronoun imposes a grammatical constraint on its anaphoric antecedent: a pronoun can be construed as anaphoric to a complex noun phrase in the sentence or discourse in which it occurs only if its gender agrees with that of the complex noun phrase. Only when the presumed antecedent NP is not complex in this sense – i.e. when it is a proper name, another pronoun occurrence or a simple demonstrative – does the pronoun's gender normally impose the same semantic constraint that it generally imposes in English. The semantic gender constraint on anaphora that is found in English is apparently not all that common cross-linguistically, not even within the typological family (Indo-European) to which English belongs.

far we always took to be part of the NP constituent of a DP. On the other hand, the information that is conveyed by the person feature has to with how the pronoun's denotation is to be identified. First and second person pronouns get their referents from the utterance context: a first person pronoun refers to the speaker or author of the utterance of which it is part, a second person pronoun refers to the addressee. Third person pronouns are different. They cannot be interpreted by linking them to speaker or addressee as the utterance context makes these available, but have to find their referents in some other way. As noted before, they can be interpreted in two importantly different ways, anaphorically or deictically. Of these two possibilities the second one is out of bounds, because our framework is unsuited for the treatment of deixis, just as it isn't suited for an appropriate treatment of proper names.

The person feature of a personal pronoun thus has to do with what kind of identification presupposition the pronoun triggers. In this regard the values of this feature seem to play a similar role as the Determiner feature $\emptyset_{pr.na.}$ that we adopted as the element at LF which determines what is to be done with the semantic representation of the NP constituent. Recall that the semantics of $\emptyset_{pr.na.}$ takes the form of an operator which transforms the non-presuppositional content of its NP complement into a presupposition. We now assume that this is also the task of the person feature value '3d person'. And as part of this decision we also assume that this feature value is specified as part of the Det constituent of a 3d person pronoun DP, in the form of the silent constituent $\emptyset_{3rd.p.sing.pron}$. [22] Furthermore, the descriptive content determined by the gender feature values is assumed to be specified as part of the NP constituent of 3d person pronoun DPs. In this way the operator determined by the person feature value '3d person' can operate on the semantic representation of the gender information in much the same way as the feature $\emptyset_{pr.na.}$ operates on the descriptive content specified in the NP constituent of a proper name DP.

In this way we are led to a syntactic structure for 3d person pronouns that is a direct analogue of the structure we adopted for the DPs that overtly con-

[22] It is also part of this proposal that the person feature values of 1st and 2nd person pronouns also determine what is to be done with the descriptive content of their NP complements (though as a matter of fact it is a reasonable assumption that there is no descriptive content at all to these pronouns). 1st and 2nd person pronouns can also be treated as generating identification presuppositions. But normally the resolution of those presuppositions will be uninterestingly straightforward: they will be resolved in the utterance context, to speaker and addressee, respectively.

sist of a single proper name. The structure for the pronoun *she* is displayed below in (4.75).

There is one remaining point to be discussed before these preliminary considerations can be put into practice. This is the descriptive content of the three gender values. There is a standard way of specifying these contents as: feminine $\rightsquigarrow$ human and female; masculine $\rightsquigarrow$ human and male; neuter $\rightsquigarrow$ non-human. Pretty much everyone who adopts this specification hastens to say that it is obviously no more than a first approximation[23], that anyone can see counterexamples if they stop to think about this for only a second, that, sure, one could do better than this first approximation, but that it is hard to do really well, and that for present purposes (whatever they may be) it doesn't matter that the approximation isn't better than it is.

The very last decision we have to make about the syntactic representation of 3d pronoun Dps concerns the structure of their NP constituents. The semantics just adopted for the gender feature values specify a conjunction of predicate for the feature values 'feminine' and 'masculine'. In particular, 'feminine' has been identified as specifying the conjunction of 'human' and 'female'. The NP part of the DP *she* will have to represent this conjunction and the question is how it should do this. On this question I am going to make an essentially arbitrary choice from small number of alternatives. We will treat the conjunction as th semantics of a complex noun that doesn't exist as such in English. In the syntactic structure of the DP *she* this noun is represented in the form of the feature value $\emptyset_{feminine}$. The artificiality of this design decision is undeniable. But it has the modest advantage that in enables us to describe the relation between syntax and semantics of 3d person pronoun DPs with the same formal means hat we have been using in dealing with the syntax of the DPs we have been dealing with so far.

(4.75) gives the syntactic structure of the DP *she*. (4.76.a) gives the lexical entry for $\emptyset_{feminine}$ and (4.76.b) that for $\emptyset_{3rd.p.sing.pron}$. The semantic representation for the DP *she* that we get by putting these tw pieces together is given in (4.76.c). As in the case of proper name DPs, this last representation could also be taken as the semantic representation for the 3rd person pronoun *she*. which is then passed up without change to the DP *she*, by anyone who assumes that personal pronoun DPs just consist of a pronoun and do

[23]For instance: *it* can be used to refer to children, *she* can be used to refer to ships, *she* and *he* to refer to mares and stallions, and most readers will be able to come up with further counterexamples to the approximation will with little or no effort.

not have internal structure.

(4.75)
$$[_{DP}\ [_{Det}\ 3rd.p.sing.pron.]\ [_{NP}\ [_{N}\ \emptyset_{feminine}]]]$$

(4.76) (lexical entry for the feature value 'feminine')

$\emptyset_{feminine}$ (noun)
x

a. Sel. Restr:

Sem.Repr: $< x \mid \begin{array}{|c|}\hline \\ \hline \text{human}(x) \\ \text{female}(x) \\ \hline \end{array} >$

b. $\emptyset_{3rd.p.sing.pron}$

Sel. Restr: —-

Sem.Repr: $< x_{ref} \mid \begin{array}{|c|}\hline \\ \hline \text{human}(x) \\ \text{female}(x) \\ \hline \end{array} > \ \rightsquigarrow$

$$< x_{ref} \mid < \{ \begin{array}{|c|}\hline x? \\ \hline \text{human}(x) \\ \text{female}(x) \\ \hline \end{array}_{3d.p.pr} \}, \begin{array}{|c|}\hline \\ \hline \\ \hline \end{array} >>$$

Combining these two entries gives us the semantics of the DP *she/her*, which is just the output representation of (4.76.c):

$$< x_{ref} \,|< \{ \begin{array}{|c|} \hline x? \\ \hline \text{human}(x) \\ \text{female}(x) \\ \hline \end{array}_{\,3d.p.pr} \}, \begin{array}{|c|} \hline \ \\ \hline \ \\ \hline \end{array} >>$$

<u>Exercise</u>: Give similar syntactic and semantic representations for the pronouns *he* and *it*.

We illustrate how the present proposal for the treatment of pronouns accounts for the contributions they make to sentences in which they occur by applying it to one of the classical donkey sentences. The sentence we consider is the conditional donkey sentence (1.44.a) that we have already looked at more than once.

(1.44.a) If Pedro owns a donkey, he beats it.

The LF we adopt for (1.44.a) is a revision of the syntactic tree for (1.44.a) that was presented in Section 1.9. The tree we need, however, has to have the provisions for tense that have been part of our LFs since the beginning of Section 3 and it should also have DPs for *he* and *it* that are like the DP for *she* shown in (4.75) (see the above exercise). This LF is shown in (4.77). The semantic representations for these DPs are shown in (4.78).

(4.77)a.

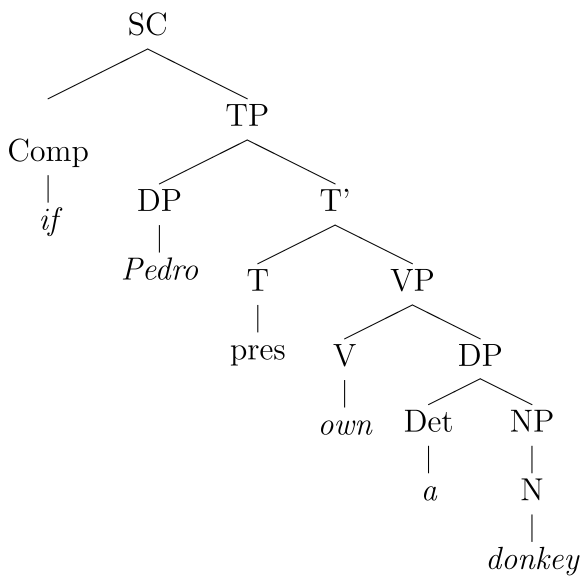

b.

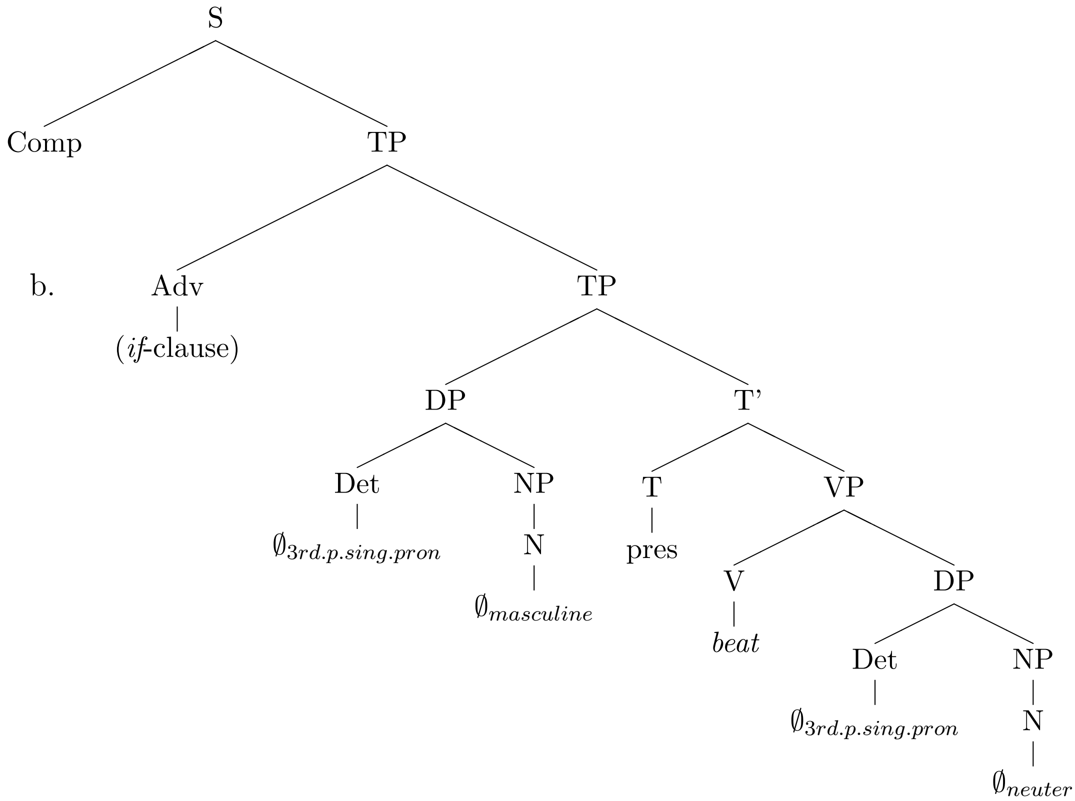

(4.78)a. $< x_{ref} |< \{ \begin{array}{|c|} \hline x? \\ \hline \text{human}(x) \\ \text{male}(x) \\ \hline \end{array}_{3d.p.pr} \}, \begin{array}{|c|} \hline \\ \hline \\ \hline \end{array} >>$

b. $< x_{ref} |< \{ \begin{array}{|c|} \hline x? \\ \hline \text{non-human}(x) \\ \hline \end{array}_{3d.p.pr} \}, \begin{array}{|c|} \hline \\ \hline \\ \hline \end{array} >>$

The semantic representation for the upper TP of the main clause is shown in (**??**) and the preliminary representation for the entire sentence, after default transfer of the non-presuppositional drefs from their stores to the Universes of the non-presuppositional DRSs to their right, in (4.79).

(4.79)

$$\begin{array}{|c|}\hline \\ \hline
\begin{array}{|c|}\hline \\ \hline
<x \mid <\{\begin{array}{|c|}\hline x? \\ \hline \text{Named}(x,\mathit{Pedro}) \\ \hline \end{array}>_{pr.na.}\}, \begin{array}{|c|}\hline t \;\; s \;\; y \\ \hline t \subseteq n \;\; t \subseteq s \\ \text{donkey'}(y) \\ s\text{: own'}(x,y) \\ \hline \end{array} >> \\ \hline
\end{array} \;\; \Rightarrow \\
\begin{array}{|c|}\hline \\ \hline
< u, v \mid <\{\begin{array}{|c|}\hline u? \\ \hline \text{human}(u) \\ \text{male}(u) \\ \hline \end{array}_{3d.p.pr}, \begin{array}{|c|}\hline v? \\ \hline \text{non-h'n}(v) \\ \hline \end{array}_{3d.p.pr}\}, \begin{array}{|c|}\hline t', s' \\ \hline t' \subseteq n \;\; t' \subseteq s' \\ s'\text{: beat'}(u,v) \\ \hline \end{array} >> \\ \hline
\end{array} \\ \hline
\end{array}$$

We now move to the second stage, in which the presuppositions of (4.79) are to be resolved. As we have a seen in our earlier discussion of donkey sentences, the antecedent of the conditional in (4.79) provides the material needed to resolve the pronoun presuppositions. But as things stand this antecedent has an unresolved presupposition of its own. This is a problem that we haven't had to face so far. What are we to do when information we would like to use to locally resolve a presupposition is itself still dependent on some other unresolved presupposition? There is more than one way in which this problem can be dealt with. The one we adopt is a radical form of playing it safe. We categorically prohibit the use of local contexts that haven't yet been cleared of their presuppositional mortgages. As long as not all of its presuppositions have been resolved (with or without accommodation) a local context is not available for the resolution of other presuppositions.

The constraint we will adopt is even more restrictive than the one just described. When a a presupposition occupies an embedded position in a preliminary representation K, then its resolution may be tackled only when there are no unresolved presuppositions on the 'accessibility projection path' that leads from the given presupposition 'upwards' through K to the 'global' context in which the represented sentence is interpreted. (The accessibility projection path is a notion that we have been making use of in more than one place in these Notes. It is connected with the question what drefs are accessible as potential antecedents for anaphoric pronouns, as first discussed in these Notes in Section 1.9, and then again at various points in Section 4, as configurationally defined constraints on presupposition resolution. In essence, the accessibility projection path that starts at a presupposition that occurs in some embedded position of a logically complex DRS K is a chain of positions within K which starts with the position of the presupposition in question and in which connects this position via a number of 'immediate accessibility' links to the position represented by the Universe of the Main DRS of which K is a sub-DRS. For an example., suppose that the given presupposition belongs to the presupposition set of the consequent of a conditional DRS Condition. Then the first link of the accessibility projection path links the position of the presupposition to this antecedent of this conditional DRS Condition. A second link connects this antecedent of conditional DRS Condition to the DRS that contains the DRS Condition as a member of its Condition Set. Suppose further that this last DRS is part of the nuclear scope of a Duplex Condition. Then the next link in the chain will link this nuclear scope to the restrictor of the Duplex CVondition and the link following this one connects the restrictor of the Duplex Condition with the DRS that contains the Duplex Condition in its Condition Set; and os on until the link is reached whose second chain is the Main DRS. A proper definition of the notion 'accessibility projection path starting from position p', where p is some position within a preliminary DRS, requires some more technical machinery, which it would be unhelpful to introduce here. But I hope that this example makes it clear enough how a formal definition would go.)

The formal constraint on pronoun presupposition resolution that I am about to state make use of the notion 'accessibility projection path' and also of a generalized notion of 'preliminary DRS'. This second notion is easily explained. it is connected with the possibility that the presuppositions of a preliminary representation (in the sense of this term in which we have been using it so far) are not necessarily resolved all at once, but that resolution may be a staggered process, in which presuppositions are resolved in some order. When presupposition resolution takes this form, then the normal situation will be that in which one or more of the unresolved presuppositions of a given representation are resolved, leading to a new representation in which these presuppositions no longer occur as presuppositions. If all goes well, then at the end of the process a representation is reached that has on unresolved presuppositions left and that is a DRS in the sense in which DRSs were defined in Section 3. The intermediate stages of this pricess will in general be representations that are like the preliminary DRSs reached at the end of Stage 1in containing representation of unresolved presuppositions. These representations are not preliminary FDRSs in the sense of being the end product of Stage 1 preliminary DRS construction, but they are formally like preliminary DRSs in the old sense in that the set of their unresolved presuppositions is non-empty. All such representations will now be called 'preliminary DRSs' irrespective of where they occur in the chain of successive representations that lead from the output of the first stage to the DRS that is the final result of the second stage. (Strictly speaking the presupposition-free DRSs that result when all presuppositions have been resolved are also preliminary DRSs according to this revised definition; DRSs constitute the limiting case if preliminary DRSs, that where there aren't any unresolved presuppositions. This is a consequence of the new definition of 'preliminary DRS' that entirely harmless, so long as we make sure that the terminology we are using doesn't lead to any misunderstandings.)

The assumption that presupposition resolution can be a process consisting of a number of successive steps, each of which leads to a representation with fewer unresolved presuppositions than the last one, is essential background to the constraint on presupposiotion resolution that we are finally ready to state. If resolution of the presuppositions of a preliminary DRS had to be a single step procedure, then the constraint would prevent many preliminary DRSs from being converted into DRSs, for which this ought to be intuitively possible.

(4.80) A presupposition occurring somewhere in a preliminary DRS may only be resolved when the accessibility projection path starting from it position is entirely presupposition free.

For (4.79) the constraint in (4.80) means that the presupposition triggered by *Pedro* has to be resolved before we can resolve the pronoun presuppositions. As we saw in the last section, an adequate account of the resolution proper name presuppositions cannot be formulated within our current framework, and the best we can do is to approximate it leads to what is shown in (4.81).

(4.81)

$$
\begin{array}{|c|}
\hline
x \\
\hline
\text{Named}(x,\mathit{Pedro}) \\
\begin{array}{|c|}\hline t\ \ s\ \ y \\ \hline t \subseteq n\ t \subseteq s \\ \text{donkey'}(y) \\ s\text{: own'}(x,y) \\ \hline \end{array} \ \Rightarrow \\
\begin{array}{|c|} \hline t'\ s' \\ \hline <u,v \mid <\{ \begin{array}{|c|}\hline u? \\ \hline \text{human}(u)\\ \text{male}(u)\\ \hline\end{array}_{3d.p.pr}, \begin{array}{|c|}\hline v? \\ \hline \text{non-human}(v)\\ \hline\end{array}_{3d.p.pr} \}, \begin{array}{|c|}\hline \\ \hline t' \subseteq n\ t' \subseteq s' \\ s'\text{: beat'}(u,v) \\ \hline \end{array} >> \\ \hline \end{array} \\
\hline
\end{array}
$$

Once this preliminary DRSs has been reached and the antecedent of the conditional DRS Condition is presupposition-free, resolution of the pronoun presuppositions from the presupposition set of the consequent is legitimate.There is not much that we need to add at this point to what has been said about pronoun resolution in earlier parts of the Notes. The old wine remains but needs to be transferred to the newly fashioned bag.

The old wine: Resolving a pronoun must take the form of finding an antecedent among the drefs that can be found in positions that are accessible form the position of the pronoun. The dref α introduced by the pronoun is then set equal to the dref β that is chosen as its antecedent. α and the equation '$\alpha = \beta$' are then added to the local Universe and Condition Set of the pronoun.

After transfer into the new bag: To resolve the referential presupposition introduced by a pronoun a dref must be found in some Universe along the accessibility projection line starting at the position of the presupposition. The chosen dref β must satisfy the DRS Conditions of the presupposition. On the assumption that this has been established, the referential presupposition can be eliminated from the given preliminary DRS while β is substituted for all occurrences of α in the resulting representation.

The requirement that β be established as satisfying the Conditions of the referential presupposition is obvious enough, and it is an aspect of pronoun interpretation that is well familiar to loa of us as listeners and readers. But it is something of which in practice we are unaware more often than we are aware of it, and it is something that is not often drawn explicit attention to in theoretical discussions of anaphora. Classical examples are 'The surgeon was emotionally drained after the difficult but ultimately successful operation. The stress wouldn't have been so intense if the patient hadn't been her son.' It is only after having reached the second sentence that the reader of this two sentence discourse is able to infer that the surgeon was a woman. When I use the word 'infer' here, the use I am making of it is the same in which it was used in Section 4.1.5. Here as there the 'inference' is a side effect of an accommodation that the interpreter sees as more or less forced upon him: He wants to interpret the pronoun *her* as referring to the surgeon. but of course that is possible only if the surgeon is female. So to make things fit he assumes that this must be the case.

In the case of our example the effects of accommodation are on the face of it less dramatic. To resolve the presupposition triggered by *he* in (4.79) in the way we want, viz. by using the dref x in (4.81), we have to mask sure that x satisfies the predicates 'human' and 'male'. Is this something we can infer from (4.81)? 'Yes', someone would answer who takes it for granted that *Pedro* must be the name of a human male. But can tho be taken for granted? Probably not. These days, the sex of someone who bears a certain first name is at best a matter of presumption or default. With *Pedro* the presumption may still be a rather strong one. And given that it is strong, the accommodation that the Pedro of (4.81) will come cheap and the interpreter is likely to be unaware that he is making it. But strictly speaking even interpreting *Pedro* as the anaphoric antecedent of *he* is making a commitment that the referent of *Pedro* is male and even in this commitment there is a bit of accommodation.

The use of x as antecedent for the pronoun *he* also requires establishing that x satisfies the predicate 'human'. Again that x has the name 'Pedro' makes this plausible, in spite of the fact that 'Pedro' could in principle be used as the name of something that is non-human (e.g. somebody's favorite donkey). But the assumption that x satisfied 'human' is also supported by another aspect of the content of the antecedent of the conditional in (4.81). x is said to own a donkey. That is a property that, presumably, only humans can

have: the verb]*own* comes with a selection restriction on its first argument to the effect that it can be appropriately filled only by arguments for which it is given or assumed that they are human. If we take such information about selection restrictions as part of the information that an interpreter can rely on, then no accommodation is need for the predication 'human(x)'.
The resolution of the identification presupposition for *it* in (4.81) is less problematic form the present perspective than the resolution of the *he*-presupposition. In order to use the dref y introduced by the indefinite *a donkey* to resolve this presupposition we need to verify that y satisfies the predicate 'non-human'. But that follows form the fact that according to the antecedent of the conditional in (4.81) y represents a donkey and donkeys are, as a matter of general knowledge, not human. So to use y in the resolution of this presupposition no accommodation is needed.

The result of resolving the pronoun presuppositions in (4.81) is displayed in (4.82). In (4.82) the Conditions that have been accommodated and inferred as part of the resolution of these presuppositions are marked in boldface. (Note well: the boldface bits are not part of the DRS; they have been added only for purposes of display. To obtain the official form of the final representation of (1.44.a) these bits should be removed.

$$(4.82)\quad \begin{array}{|c|} \hline x \\ \hline \text{Named}(x,Pedro)\quad \text{male}(x)\ \textbf{[accomm]}\quad \text{human}(x)\ \textbf{[inferred]} \\ \begin{array}{|c|} \hline t\ \ s\ \ y \\ \hline t \subseteq n\ \ t \subseteq s \\ \text{donkey'}(y) \\ s\text{: own'}(x,y) \\ \text{non-human}(y)\ \textbf{[inferred]} \\ \hline \end{array} \Rightarrow \begin{array}{|c|} \hline t'\ s'\ u\ v \\ \hline \text{human}(u)\ \ \text{male}(u) \\ \text{non-human}(v) \\ u = x\ \ v = y \\ t' \subseteq n\ \ t' \subseteq s' \\ s'\text{: beat'}(u,v) \\ \hline \end{array} \\ \hline \end{array}$$

4.2.2.1 Further Constraints on Pronominal Anaphora

In the discussion above we have mentioned two constraints on the resolution of anaphoric 3rd person singular pronouns: (i) the dref β that is chosen as antecedent has to be accessible from the position of the pronoun presupposition (see (**??**) in the lest subsection) and (ii) it must be possible to attribute to its denotation the properties that the presupposition attributes to its referential argument (i.e. to the question-marked dref α that occurs in the Universe of the presupposition). In PART I we reviewed a number of examples which show the need for modification of this general accessibility constraint: pronominalization was possible in certain cases in spite of the fact that the accessibility constraint was not satisfied as is, and we had to assume the availability of certain operations on the context DRS to render the intended antecedent for the pronoun accessible while preserving the truth conditions of the context DRS. Building these operations into our present formalism with its more complex representations is a bit of technical challenge. But there are no fundamental issues here, so i leave this task to whosoever may feel motivated to carry it out.

But there are further complications, which haven't been mentioned so far. (i) and (ii) are not the only constraints on pronoun resolution. The examples below point to some further constraints. And these constraints differ from the ones considered up to now in that the DRT framework we are using is not equipped to deal with them.

Syntactic Constraints on Pronoun Anaphora

One of these constraints is illustrated by the following sentence pair.

(4.83)a. He chased a woman who loathed a man.

b. A man chased a woman who loathed him.

It is plain that (4.83.a) cannot be understood as expressing the same proposition as (4.83.b): the pronoun *he* in (4.83.a) cannot be interpreted as anaphoric to the phrase *a man* that is embedded within it. To capture such constraints is the aim of Chomsky's Binding Theory, see (Chomsky 1981) and many subsequent publications. The Binding Theory articulates the limits that the syntax of English and other natural languages impose on anaphoric relations, by stating constraints for the different types of noun phrases that allow for anaphoric 'binding' (which for the Binding Theory include all those noun phrases which we have classified as definite). In particular, 3rd person pronouns ('pronominals' in Chomsky's terminology) are subject to a principle, known as 'Principle B', to the effect that a pronoun/pronominal may not 'command' its antecedent. The notion of command referred to is defined in terms of the configurational structure of syntactic trees and there has been some dispute over its exact definition. But all definitions that have been proposed converge on the conclusion that *he* in (4.83.a) cannot be anaphoric to *a man* because the two noun phrases do not stand in the right command relationship: the pronoun commands its putative antecedent.

Another syntactic constraint that the Binding Theory deals with is illustrated by the contrast between (4.84.a) and (4.84,b). The *him* in (4.84.a) cannot be interpreted as anaphoric to the subject DP *John*. To express the proposition that (4.84.a) would express on such an interpretation if that interpretation was possible we have to use the reflexive pronoun *himself*, as in (4.84.b).

(4.84)a. John admired him.

b. John admired himself.

c. Mary compared John to him.

d. Mary compared John to himself.

e. John found a snake near him.

f. John found a snake near himself.

g. John talked to Mary about him.

h. John talked to Mary about himself.

i. Mary talked to John about him.

j. Mary talked to John about himself.

k. John_1 was happy. He_1 admired $^*\text{him}_1/\surd\text{himself}_1$.

We find a similar contrast between (4.84.c) and (4.84.d): the *him* of (4.84.c) cannot be interpreted as coreferential with the direct objet DP *John*; here too we need the reflexive *himself*. These two example pairs suggest that a pronoun cannot be anaphoric to a DP in the same clause. But the facts are complicated. For instance, what is unambiguously expressed by (4.84.f) can also be expressed by (4.84.e), though the sentences differ in that the proposition expressed by (4.84.f) is the only proposition this sentence it can express, whereas (4.84.e) can be understood not only to express this proposition but also allows for interpretations in which *him* is taken to refer to someone other than John. A further compilation is shown by the sentences (4.84.g - j). For (4.84.g) a reading in which *him* is interpreted as anaphoric to the subject *John* is for most speakers impossible or highly marginal; for them (4.84.h) is the only way to say this. But when subject and *to*-object are interchanged, as in the next pair, then it seems possible again to interpret *him* as anaphoric to *John*, see (4.84.i).

The constraints illustrated in (4.84) are syntactic – certain syntactic configurations involving a pronoun and a putative antecedent for it prohibit an anaphoric link between them. The Binding Theory makes use of certain configurational notions to account for the syntactic constraints on pronominal coreference. It is a point of debate among syntacticians whether this effort has ever been fully successful. But that is a debate into which we neither can nor want to enter. We finesse these problems by a subterfuge to which we have resorted on earlier occasions. We declare that since these constraints are syntactic, it is the syntax that should deal with it. And in our terms this means that the syntactic parser that computes the LFs from which our semantic representations are constructed has identified these restrictions. Assuming that to be so, a lll we ask of it to make thus information explicitly available in the LFs it delivers.

Let us assume that this information is given in the form of an *index set* that is associated with each DP in the given LF and that it consists of indices for each of the DPs in the sentence for which the first DP cannot serve as anaphoric antecedent for syntactic reasons. (We use the indices here that coined DPs with their argument slots after lexical insertion for the predicate words that these DPs are syntactic arguments to. In the present setting, where we are only concerned with the semantics of third person pronouns, we may assume that index sets consist exclusively of indies of 3rd person pronoun DPs.)

To implement this idea it is convenient to assume that lexical insertion (i.e. insertion of the semantic representation provided by the lexical entry of the word) is performed for all predicate words occurring in the input LF before any other construction steps are carried out and that the index sets are put in place after these insertions have taken place. Each DP now gets its index set attached to it in the form of a subscript. (Strictly speaking these subscripts take the form of finite lists of DP indices, but the order in which these indices are listed doesn't matter.)

The index sets play their part at both levels of DRS construction, first in the course of constructing the preliminary representation and then during presupposition resolution. (i) When during the construction of the preliminary DRS a DP from the LF is interpreted and as part of that a dref is chosen as its referential argument, then both the DP's index and its index set are transferred to this dref. (ii) When the identification presupposition of a pronoun is resolved, then no dref may be chosen as antecedent for the pronoun whose index set contains the index of the question-marked dref of the pronoun's identification presupposition.

Unfortunately this isn't the whole story about syntactic constraints on pronoun resolution. That it isn't is shown by (4.84.k). An interpretation of this two sentence discourse in which the pronouns are interpreted as indicated informally by the subscripts is impossible. That isn't much of a surprise, since such an interpretation violates the prohibition that a pronoun cannot be interpreted as anaphoric to another DP in the same clause. But as the syntactic constraints on pronoun resolution have been stated, they let this case slip through the net. Nothing we have so far said prevents the two pronouns from being both interpreted as anaphoric to the subject DP *John* of the first sentence.

What (4.84.k) shows is that when two DPs are prohibited from standing in a direct anaphoric relation – neither may be interpreted as anaphoric to the other – then they may not be connected by a chain of such links either. The illustration that (4.84.k) provides of this principle is a very simple one: the two pronouns of the second sentence may not be interpreted as anaphoric to the same antecedent phrase. But it isn't hard to see that the principle is a more general one. For instance, suppose that we replace the first sentence of (4.84.k) by 'John was happy that everybody liked him'. Then one formal option for the interpretation of the pronouns of the second sentence would be for the first to be interpreted as anaphoric to *John* and the second to the occurrence of *him* in the new first sentence. If at the same time the *him* of the first sentence is interpreted as anaphoric to *John*, then we have once more a situation in which the pronouns of the second sentence are linked by an anaphoric chain, yielding an interpretation that is clearly not there.

Evidently a remedy against this can only be found by regarding the preliminary representation which contains the pronoun presuppositions whose solution is at issue as part of a larger structure which also includes the discourse context provided by the DRS for the preceding discourse. (For an example, in the case of (4.84.k), this discourse context will be the DRS of the first sentence.) As more and more identification presuppositions get resolved in the course of the construction of this larger structure, anaphoric coreference networks get established as part of it: Each time a pronoun presupposition is resolved, by choosing an accessible dref β as antecedent, then this establishes an *anaphorically link* between α and β. So at any point in the course of the construction each dref γ determines a 'coreference cluster', consisting of all the drefs δ which are connected with γ by a chain of anaphorically links. We now also assume that presupposition resolution is strictly sequential: At each construction step at most one presupposition is justified. The syntactic constraint on pronoun resolution is now the following: In order that a dref β can serve as antecedent in the resolution of a pronoun presupposition with question-marked dref α the index of α must not occur in the index sets of all drefs δ in the cluster determined by β at that point.

This abstract description of how unwanted intra-clausal pronominal coreference can be avoided may be a little difficult to decipher. So let us see how this works for (4.84.k). Let us suppose that the DRS for the first sentence of (4.84.k) is as in (4.85).

$$(4.85)\ \begin{array}{|c|}\hline t\ \ s\ \ x_1^{\emptyset} \\ \hline t = n\ \ t \subseteq s \\ s : \text{happy'}(x) \\ \hline \end{array}$$

In (4.85) 1 is the index of x. x's index set is empty, since there are no pronouns in the first sentence that are prohibited from using x as antecedent. The preliminary representation for the second sentence may be assumed to have the form shown in (4.86).

(4.86)

$$<t', s', y_2^{\{3\}}, z_3^{\{2\}}\ |<\{ \begin{array}{|c|}\hline y? \\ \hline \text{h'n}(y) \\ \text{male}(y) \\ \hline \end{array}_{\,3d.p.pr}, \begin{array}{|c|}\hline z? \\ \hline \text{h'n}(z) \\ \text{male}(z) \\ \hline \end{array}_{\,3d.p.pr} \}, \begin{array}{|c|}\hline \\ \hline t' \prec n \\ t' \subseteq s' \\ s'\text{: adm'}(y,z) \\ \hline \end{array} >>$$

Suppose now that we first resolve the presupposition for *he*, by using x in (4.85) as antecedent. This leads to the modified preliminary DRS in (4.87.a). It also creates an anaphoric link between y and x and with that the cluster $\{x, y\}$, 'officially' displayed in (4.87.b).

$$(4.87)\text{a.}\quad <t', s', z_3^{\{2\}}\ |<\{ \begin{array}{|c|}\hline z? \\ \hline \text{h'n}(z) \\ \text{male}(z) \\ \hline \end{array}_{\,3d.p.pr} \}, \begin{array}{|c|}\hline y_2^{\{3\}} \\ \hline t' \prec n\ \ \ t' \subseteq s' \\ y = x \\ s'\text{: adm'}(y,z) \\ \hline \end{array} >>$$

b. $\{x, y\}$

At this point it is no longer possible to resolve the presupposition for *him* also to x. The reason is that x now is part of a cluster with a member that contains the index of the question-marked z in its index set. Had we first resolved the *him*-presupposition to x, thence would have run into the same problem when then truing t resolve the *he* to x. Note that in the general description we failed to make room for a formal representation of the anaphoric clusters that are formed as the result of the anaphoric resolutions of referential presuppositions. One way to solve this problem is to list the set of clusters as a separate component, in the way we have done in the example.[24]

Anaphoric possibilities that our current account wrongly forbids

Pronominal anaphora resolution is not only complicated by the syntactic constraints discussed above. There are also complications of a different sort. Most of these point in the opposite direction: there are more options for anaphora resolution than our account allows for in its present formulation. A first glimpse of these ;after complications is provided by the following list of examples. All are variations of the three classical donkey sentence paradigms that were first mentioned in (1.44).

(4.89)a. If Pedro owns a donkey he beats it.

b. If a farmer owns a donkey he beats it.

c. Every farmer who owns a donkey beats it.

d. If he owns a donkey, Pedro beats it.

e. If Pedro owns it, he beats a donkey.

f. If he owns it, Pedro beats a donkey.

[24] The prohibitionism that can be made explicit in the way just shown should not be confused with coreference of two DPs that comes about in some other way than through mere anaphoric linking, such as through an assertion of sameness. An example is the sentence in (4.88).

(4.88) That man just put on John's coat. So he must be John.

This sentence pair cannot be true, unless *he* and *John* refer to the same individual. And there is nothing wrong with interpreting it in such a way that these are its truth conditions. The formal difference between the case presented by the interpretation construction of this sentence pair and the constructions discussed above should be easy to see. The natural interception of *he* in (4.88) is by using the dref for *that man* as antecedent. That is unproblematic because the dregs representing *that man* and *John* do not belong to the same cluster.

g. He beats it, if Pedro owns a donkey.

h. He beats a donkey, if Pedro owns it.

i. Pedro beats it, if he owns a donkey.

j. Pedro beats a donkey, if he owns it.

k. If a farmer owns it, he beats a donkey.

l. If he owns a donkey, a farmer beats it.

m. If he owns it, a farmer beats a donkey.

n. He beats it, if a farmer owns a donkey.

o. He beats a donkey, if a farmer owns it.

p. A farmer beats it, if he owns a donkey.

q. A farmer beats a donkey, if he owns it.

r. Every farmer who owns it beats a donkey.

s. Every farmer beats it, if he owns a donkey.

t. Every farmer beats a donkey, if he owns it.

(4.89.a,b,c) are the old (1.44.a,b,c). When these sentences were first discussed in PART I, it was noted that they are generally judged to be grammatical. (4.89.d,e,f,g,h,i,j) are variants of (4.89.a). The first of these seems perfectly acceptable, which is in accordance with our present account: The *Pedro*-presupposition can be resolved first, yielding a dref at the top level, and this dref can then be used for the resolution of the presupposition triggered by *he*. On the other hand, according to our present account (4.89.e) and (4.89.f) should not be acceptable, unless it can be argued that the indefinite DP *a donkey* can be given a specific reading and its dref placed in the Universe of the main DRS rather than in that of the consequent DRS of the $\Rightarrow$-Condition. But even if such a specific reading of the indefinite is possible that would attribute truth conditions to these sentences which are not the ones that speakers get (assuming that they find these sentences acceptable). On their prominent interpretation the truth conditions of all the sentences (4.89.d,e,f,g,h,i,j) are the same as those of (4.89.a).

As a matter of fact, the sentences from (4.89.d,e,f,g,h,i,j) for which this accessibility problem arises, viz. (4.89.d,e,f,h,j) do not seem any worse than the remaining two, (4.89.g) and(4.89.i). (4.89.g), in particular, sounds quite awkward. So we are facing a double problem here: (a) What makes (4.89.g) and(4.89.i) as awkward as they seem to be? (b) How can it be that (4.89.d,e,f,h,j) are any good at all? We will set problem (a) aside. There may be processing factors at work here, or some combination of syntactic and processing considerations, that explain why these sentences are suboptimal even though our account predicts that they should be interpretable and in such a way that their truth conditions come out the same as those of (4.89.a). But such factors are not what our account is supposed to be able to address in the first place.[25] Problem (b) ion the other hand is one that we cannot afford to ignore. Our account shouldn't rule out interpretations for these sentences that makes them truth=conditionally equivalent to (4.89.a). How can we adapt it so that it doesn't do this?

The next eight conditionals from the list in (4.89), (4.89.k-q), raise much the same issues as those we just looked at. The only difference is that the subject is now the indefinite *a farmer* instead of the proper name *Pedro*. (So these sentences are variants of (4.89.b) in the same way that (4.89.d-j) are variants of (4.89.a).) Once more, acceptability appears to be a matter of degree and as i went out of my way to stress above, questions of graded acceptability are among those that our framework is not designed to deal with. However, the second problem also shows up here, and with a vengeance, because we are now dealing with two indefinites instead of one. In fact, all but one of the sentences in this bunch are affected by this problem: for none of them is our approach able to come up with a semantic representation that assigns them their intuitive truth conditions. And the one sentence for which this is not true, (4.89.n), is out (or marginal at best) for a different reason, which we observed earlier: a sentence-initial pronoun that is the subject of a the main clause cannot be anaphoric to any DP in the sentence.[26]

[25](4.89.g,h,i,j) differ from the preceding sentences in that the *if*-clause follows the main clause rather than preceding it. These sentences are generally felt to be somewhat degraded when compared with their counterparts in (4.89.a), (4.89.d), (4.89.e) and (4.89.f), in which the *if*-clause comes before the main clause; and it may be that they are acceptable only to the extent that it is possible and legitimate for the interpreter to transform them into those counterparts as part of what he has to do to interpret them (putting the *if*-clause 'back to where it belongs', so to speak). A more thorough canvassing would be needed to assess in more detail the degrees to which these sentences are acceptable and in particular which are more acceptable than which. But whatever the outcome of such an investigation, it is clear that within the framework we are using it will be impossible to say much if anything to explain the results. Grades of acceptability are not what our account is deigned for. The predictions that our framework is able to make are all black-and-white; a given interpretation of a sentence is either possible or it is not; grisium non datur.

[26]The impossibility to resolve anaphoric pronouns that are sentence-initial subjects of main clauses sentence-internally looks like a configurational constraint, which should be explained by syntax. It isn't quite clear, however, how a syntactic approach like the Binding Theory would be able to account for this. It would seem that a configurational

Of the remaining three sentences, (4.89.r,s,t), the last two present a new complication,which has to do with the scope relation between *every* and *if*. According to the assumptions web have been making about adverb adjunction and quantifier raising both the *if*-clause and the quantifying DP *every farmer* end up in an adjunction position to the main clause TP node. The scope relation between DP and *if*-clause are not fixed by these assumptions.It might be argued that since for the *if*-clause its position was TP adjunct is its base position and the TP adjunction position of *every farmer* is the result of movement, the latter position should be above the former, so that the quantifier has wide scope relative to the *if*-clause. As a matter of fact, these syntactic speculations are of little import in relation to (4.89.s,t). For in each of these sentences the *if*-clause contains the pronoun *he*, which in the interpretations we are after must be resolved to *every farmer* as its anaphoric antecedent and that is possible only if the DP *every farmer* has scope over the *if*-clause. Given that this is the only one of the two scope relations that could serve our purpose, the central problem with (4.89.s,t) is once more an instance of problem (b): (4.89.s) is a sentence for which our account allows us to construct the interpretation we are after, but for (4.89.t) this is not the case. And yet (4.89.t) seems a much more natural way of conveying the content that (4.89.s) and (4.89.t) share than (4.89.s).

(4.89.r) is also among the sentences for which our current account does not deliver the interpretation we seek. That may not seem such bad result in this case, as (4.89.r) doesn't come across as a particularly goods sentence. But when we reflect on why that might be, another feature of this sentence pops into prominence. (4.89.r) seems to suggest that the same donkey could be owned by a number of different farmers, and since seems an impossibility the sentence strikes us as odd. Joint ownership of a single donkey by a lot of different farmers, in the way in which, say, a combine harvester could be jointly owned by the members of a farmers' cooperative, is implausible to start with, and, even more importantly, joint ownership cannot be expressed by the verb *owns* as it is used in (4.89.r). To see the relevance of this

explanation would have to rely on the assumption that the sentence-final *if*-clause cannot be adjoined to TP (or at an even higher position). For if such an attachment were possible, then it seems impossible to explain why the resolution is possible when the *if*-clause is sentence-initial but not when it is sentence-final. It may be that some kind of processing constraint explains the difference between these two cases. But if that is so, I do not know that anyone has a clear idea how to formulate such a constraint or how this constraint is related to, or interacts with, purely configurational principles like those of the Binding Theory.

consideration, compare (4.89.r) with the sentences in (4.90).

(4.90)a. Every art-collecting tycoon who finds out that it is for sale will bid for a canvas by Picasso.

b. Every art-collecting tycoon who finds out that a canvas by Picasso is for sale will bid for it.

(4.90.a), the direct analogue of (4.89.r), is somewhat less natural than (4.90.b) . But it still seems quite acceptable, and definitely mores than (4.89.r). The natural interpretation of both sentences in (4.90) seems to be something like this: Whenever a painting by Picasso is for sale, then every art-collecting tycoon will bid for that painting. We will see below what this aspect of the sentences (4.89.r) and (4.90.a,b) has to do with problem (b).

At last we turn to an analysis of problem (b). We have seen that this is a problem for most of the sentences in (4.89): If our present account of pronoun anaphora were right, then none of these sentences should be able to get the interpretations in question. And yet some of these sentences seem quite acceptable, and better than some of the sentences for which our account does license the wanted interoperation. But the account of pronominal anaphora proposed in the last section predicts that none of them can have the readings that speakers seem to get for them. Is there anything we can do to remedy this situation?

From one angle the chances may look dim. One central assumption we have been making all along is that anaphoric pronouns must find their antecedents among the accessible discourse referents. But how do we get the referential argument of an indefinite in the main clause of a conditional in a position that is accessible to a pronoun occurring in its *if*-clause? Another guiding principle of DRT since the beginning is that antecedents of conditionals are local contexts for their consequents but not the other way round. That is supposed to follow from the very notion of a conditional, as an information structure in which the consequent says more about the situation or situations posited by the antecedent. This principle – the antecedent is accessible from the consequent, but not the other way round – is non-negotiable.

The solution to this problem therefore has to be found elsewhere. In a nutshell this is what I think the solution must be: All the sentences for which problem (b) arises involve a hidden generic quantifier. This quantifier allows the indefinite to be bound 'piggyback', with the same scope as the dref directly bound by the generic quantifier, in the same away that the indefinite in (4.89.c) is bound ‘piggyback’ by the quantifier *every farmer.*

With the help of a generic quantifier we can semantically represent the first of the sentences in (4.89) that give rise to problem (b), viz. (4.89.e), in the following form.

(4.91)

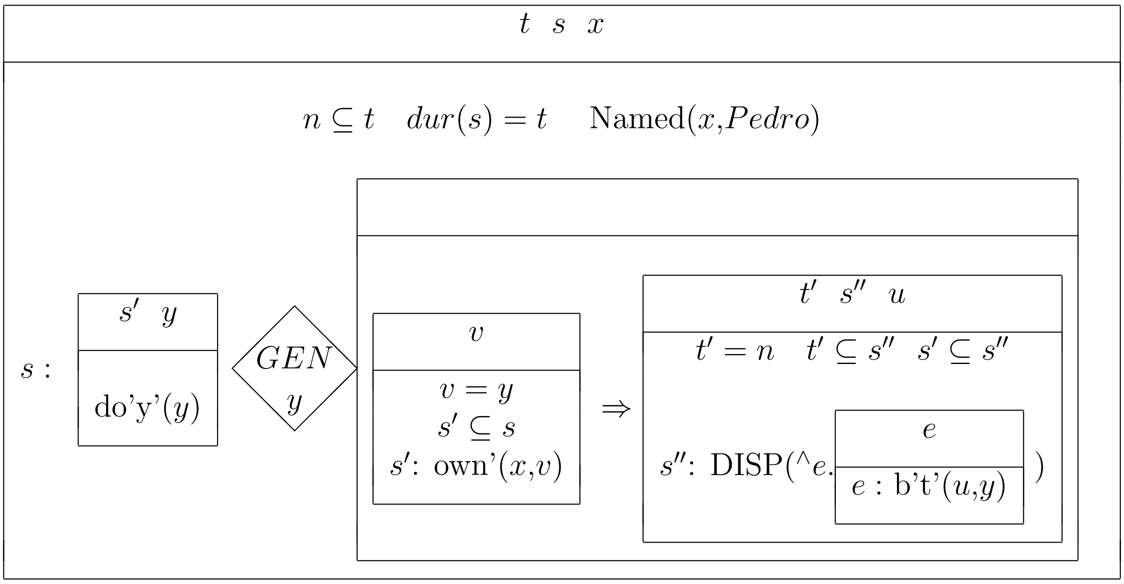

N.B. There are a number of aspects of this representation that are not covered by what has been said about the construction algorithm. First, there is the question what triggers the generic quantification in (4.89.e). I’ll return to this briefly below, but for now this much should suffice: What triggers a generic quantifier is a combination of the simple present tense and the indefinite *a donkey.* And the effect of this is a quantification state s whose content is given by a Duplex Condition with the generic quantifier GEN which binds the referential argument of the indefinite. (More needs to be aid about the principle that allows for the indefinite to be ‘extracted’ from the position it occupies in the LF from which the semantic representation is computed,which must bne part, in one way or another, of the operations that introduces the generic Duplex Condition.) A crucial feature of this operation is that in the resulting Duplex Condition the nuclear scope consists of everything that is left behind after the indefinite has been extracted. This consists of the semantic representation of the (remnant of the) main clause and the *if*-clause; and this representation takes, as always, the form of a conditional DRS Condition. In the preliminary representation the antecedent of this conditional will have the identification presupposition for the pronoun *it* left-adjoined to it. Resolution of this presupposition can now make use of the dref y, leading to the equation '$v = y$' in (4.91). Note also that the Simple Present Tense use of the verb *beat* has been interpreted as involving coercion into a dispositional interpretation.

Other sentences from (4.89) that are affected by problem (b) can be dealt with in similar ways. For instance, (4.89.k) can be given a semantic representation in which the indefinite *a donkey* has been lifted to the restrictor of a wide scope generic quantifier, just as in (4.91), whereas the indefinite *a farmer* can be interpreted locally to the *if*-clause. To get the right logical form for (4.89.q) both *a donkey* and *a farmer* have to be made part of the generic quantification, so that both pronouns in the *if*-clause can be interpreted in the way we want. And the semantic representations for (4.89.r) and (4.90.a) also require generic quantification over the referential argument of *a donkey* or *a canvas by Picasso*; but here the nuclear scope is another Duplex Condition, triggered by the quantified subject DP of the main clause (*every farmer* or *every art-collecting tycoon*). These DRSs enable us to see why such sentences are naturally understood as carrying an implication that there can be several art-collecting tycoons finding out that a given Picasso is for sale or (implausibly) that there can be several framers each owning the same donkey.

While I think that this is in outline the right response to problem (b), it should be clear that several pieces are missing from this outline. First, it is unclear what the different factors are that can trigger generic interpretations. I mentioned the use of the Simple Present tense and the occurrence of a DP beginning with the indefinite article *a*. But certain past and future tense sentences can also get generic interpretations. And the presence of an *a*-indefinite isn't a necessary condition either. Bare plurals can play this part as well, as can – perhaps – be singular and plural definite descriptions (as for instance in 'The duck-billed platypus is an animal of unusual shape and habits living in Tasmania'). I expect there will be other factors too, but do not know what they are. For much detailed information about genericity in natural language see (Carlson & Pelletier 1995*b*).

A further potential problem that we seem to be taking on board is a new source of ambiguity. The point is perhaps most easily explained with reference to (4.89.k). What I said about this sentence above was that in the construction of a semantic representation for this sentence that assigns it the desired truth conditions, *a donkey* has to be made part of the generic quantification, but the indefinite *a farmer* can be given the usual treatment which leads to its referential argument ending up in the Universe of the antecedent DRS of the conditional DRS Condition. But a viable representation can also be obtained when *a donkey* and *a farmer* are both treated as part of the generic quantification. By the same token, nothing seems to be stopping us from giving analyses involving generic quantification for the sentences in (4.89) that are not affected by problem (b). Does this mean that all these sentences are structurally ambiguous? And if not, then why not?[27]

Another matter that needs to be addressed by any theory of genericity is the semantics of the generic quantifier (i.e. of the verification conditions for Duplex Conditions that have GEN as operator in their central diamond). This is a topic that has been extensively discussed in the literature and for which a range of explicit proposals have been made, but without an acknowledged winner. Once again consult (Carlson & Pelletier 1995*b*).

If much in this section has come over as rather impressionistic and ad hoc, then that impression is not mistaken; for much of it has been impressionistic and ad hoc. I am leaving this section in its current state nonetheless, at least for now. Its main point has been to give some taste of the complexity of the data that a theory of pronominal anaphor should be able to account for, including data to be found in the immediate vicinity of the classical donkey pronoun sentences that played such a crucial part in the original conception of DRT and that have been responsible for the central architectural features which DRT has retained from its beginnings to the present day. So it seemed only right to provide and inkling of how much is hidden behind the facade of DRT's original treatment of the classical donkey cases.

One thing at least should have become clear: Among the problems that a full account of pronominal anaphora should be able to handle there are two types which are clearly and importantly different: (i) there are the configurational

[27] I do not think that even if quite a few sentences would prove to be ambiguous in this way that would necessarily count against the account that would treat these as ambiguous. That structurally distinct analyses of the same sentence may lead to the same truth conditions is something that good many accounts of complex sets of syntactic and/or semantic data must allow for. It is nonetheless important to see more sharply how much ambiguity is imported by an account in which indefinites can get generic interpretations as well as those we acknowledged earlier in these Notes. But that will be possible only when an account of genericity and of the roles that indefinite play in it has been worked out ingrate formal detail than we can do here.

constraints that I argued should be left to syntax and (ii) there are the problems that have to do with indefinites as 'generic quantifiers' . These, it appears, can be solved only by assuming that indefinite DPs allow for generic interpretations and that these interpretations allow them to be lifted from their syntactic positions to become quantifiers with a wider scope. It would have been wrong, nearly 40 years after DRT was first conceived, to discuss a phenomenon that was central to ist first formulations without pointing to some of the factors that cloud the illusory clarity of the picture which that account may once have seemed to give us.

4.2.3 Definite Descriptions

For the project of giving a systematic account of the reference conditions for the different types of definite DPs definite descriptions present a special challenge. This is because of their versatility: definite descriptions outstrip all other definite DP types in the range of their different uses.

Some of these uses are deictic. Like other deictically used DPs deictic uses of definite descriptions refer to entities from the non-linguistic context, as for instance when I say to you *The man over there is someone I have met before*, referring to the man we can both see. As noted before, deictic uses cannot be properly treated within our framework and we therefore set such uses of definite descriptions aside in this survey, just as we did for pronouns.

But even when we restrict attention to non-deictic uses of definite descriptions, in which they find their referents without the help of the non-linguistic context, there remain several different use types of definite descriptions that we need to distinguish. One of these is the use that the classical accounts of definite descriptions, starting with Frege, have focused on – that in which a definite description has a proper denotation if and only if there is a unique satisfier of its descriptive content. We had a few things to say about these accounts in the introduction to Section 4, where we noted that this use of definite descriptions was the topic of the early history of Presupposition Theory starting with Frege.

As noted in those introductory remarks, the classical view of definite descriptions, common to Frege and Russell (and many others besides), is that when a definite description refers properly, it refers to the unique satisfier of its descriptive content. What Russell and Frege disagreed about is what follows when there is no unique satisfier. One aim of this subsection is to see what our approach to the reference of definite DPs has to say to each of the two classical accounts, the Frege-Strawson account – according to which failure of unique satisfaction produces failure of denotation, which in its turn entails failure of the sentence containing th description to express a proposition and with that failure to determine a truth value – and the Russellian account according to which failure of unique satisfaction is to be treated as a conjunct of the presupposition expressed, which in the typical cases (when the description is interpreted as having maximal scope) entails that the expressed proposition is false.

It might be thought that our approach is just one more version of the Fregean view that definite descriptions generate presuppositions, and thus that it is obviously more like the Frege-Strawson than the Russell account. We will see, however, that its relations to those earlier accounts are more complicated. One reason for this has to do with an aspect of presupposition that played a part in our introduction to presupposition theory in Sections 4.1.1 - 4.1.4 – especially in our discussion of the 'Hey, wait a minute' test – but which then disappeared from view when we switched in 4.1.5 to the dynamic perspective which has been guiding us from then on. The dynamic perspective led us to distinguish between verification of presuppositions in the given context as it is given to the interpreter and cases where there is no verification by the 'context as given' and where accommodation has to come to the rescue. What we have not yet taken into consideration in our formal accounts of presupposition representation and resolution are cases where accommodation is impossible because the interpreter si aware that the needed accommodation would contradict what he independently knows to be the case. For propositional presuppositions (as distinct from referential ones) such cases have a straightforward description: They are the cases where the interpreter knows the presupposed proposition to be false. Accommodating information that entails a proposeition that one knows to be false is to knowingly adopt a contradiction. That wouldn't just be unsound; it would be what is arguably a psychological impossibility.[28]

[28] This isn't completely right. As already noted by Stalnaker in the seventies **[ref.]**, interpreters will sometimes be prepared to make assumptions that they know to be false in order to justify what they know to be unsatisfiable presuppositions of utterances that the speaker makes, either keeping in the back of their minds that the speaker must be wrong in the presupposition she is making or going along with her in what they understand as a piece of counterfactual reasoning or speculation. Here we set such cases of 'speaker-hearer collusion' aside, assuming that whenever an interpreter discovers that justification of a presupposition would lead to a contradiction, interpretation grinds to a halt and no logical form is assigned to the current sentence.

The use of definite descriptions provides many examples of the impossibility of presupposition accommodation. This is so not only on the assumption that a definite ascription presupposes the unique satisfaction of its descriptive content, but also on the weaker and less controversial assumption that definite descriptions merely presuppose the satisfaction of their descriptive content, unique or non-unique. Consider once more the classical example (4.2), repeated below.

(4.2) The King of France is bald.

The typical recipient of an utterance of this sentence made in the 21-st century will know that the presupposition of the definite description *the King of France* contradicts what he knows to be for a fact: that there is no current King of France. What can the recipient who knows that France doesn't have a king do in reaction to this utterance other than to say something like "Hey, wait a minute"?[29] Such a reaction is a sign from the interpreter to the speaker that his interpretation of her utterance has reached an impasse. It may be the start of a correction process, perhaps with the speaker withdrawing her utterance in the end. Our framework does not provide us with the means to say anything about how such correction processes might go. All it permits us to state is that the interpretation has failed and thus that no logical form has been assigned to the utterance.

This verdict is in line with the Frege-Strawson perspective: The impossibility to construct a logical form for the utterance entails the impossibility to take it as expressing a proposition, and so the question whether the expressed proposition is true or false cannot even be raised. But the matter is further complicated by the role of negation in natural language (or at least natural languages like English). Among the ways in which negation is used in English there is a family of uses described in considerable detail in the work of Horn ((**?**)). As Horn has noticed, negation does not always target the truth value of the utterance (and reverse it); it can target any aspect of it that renders the utterance suboptimal: the way in which a word in it is pronounced or the way in which a word is spelled, the apparent lack of 'force' of an expression that renders the utterance inadequate to the state of affairs it describes or the message it wants to get across, the social 'register' of the weirds to grammatical constructions used, the politeness conventions connected with

[29]The only alternative would be to 'let this pass' and to interpret this and following utterances by the speaker as premised on this non-actual background assumption, see the last footnote 28.

the social relations between speaker and addressee (like the use of *tu* and *Vous* in French), and so on. These interpretational options exist in particular for negation expressed through the use of the English particle *not*, see (4.92).

(4.92)

a. Pope Gregory VII didn't excorporate Henry the Fourth. He excommunicated him.
b. Things didn't go àwry. They went awrỳ.
c. The concert didn't go well. It was a roaring success.
d. (Spoken by Fritz Reiner:) The president didn't come to Fritz's concert. For you it is "Mr. Reiner".

The first sentences of each of these 2-sentence utterances can be understood as a statement that what is in the scope of the negation is not true for the reason mentioned in the second sentence. But note well that the second sentence is needed to make clear what is wrong with what is in the scope of the negation of the first. Without this additional information it is quasi-impossible to interpret the first sentence in such a way.

False presuppositions are also among the improprieties that a negation can target. Thus the first sentence of (4.93) can be understood as saying that what its negation applies to is not true because, as the second sentence makes clear, failure of the presupposition carried by the subject phrase *The current King of France*.

(4.93)The King of France isn't bald. There is no current King of France.

Given what was said in Section 4.1 about projection, the first sentence of (4.93) would seem to be a case of a presupposition that does not project. After all, there is no requirement that the presupposition – of there being a unique current King of France – is satisfied in the context in which (4.93) is uttered. But note well that the reasons for non-projection are very different here from those considered earlier. Thus far, a presupposition could fail to project because it was locally verified. In cases like (4.93) the reason why the presupposition doesn't project is not its local satisfaction, but rather the open admission that it isn't satisfied at any level, local or global. We will refer to such cases, in which an utterance is interpreted as stating the non-satisfaction of one of its own presuppositions, as cases of *presupposition cancellation*.

Sentences in which negations are interpreted as denials that a presupposition is true go some way in the direction of a Russellian treatment of presupposition failure. Such sentences are true for the reason that on Russell's Theory of Descriptions a sentence is true in which an improper description is construed as occurring within the scope of negation, while this negation does not occur in the scope of any other logical operators. (According to the Theory of Descriptions the sentence is true because it contains the failing presupposition as a false conjunct within the scope of the negation. The scope of the negation will be false since it has a false conjunct; so the negation itself, and therewith the sentence as a whole, will be true). But note well that this isn't going in Russell's direction very far. For as we have seen, both from the examples in (4.92) and the one in (4.93) Horn-type interpretations are not easily available. It is only the explicit presence of the information provided by the second sentences of the examples in (4.92) and (4.93) that makes these interpretations accessible. And for unnegated sentences with improper descriptions, in which there is nothing that can trigger cancellation, the tension between speakers' intuitions and Russell's Theory of Descriptions remains. Such sentences, one cannot help feeling, aren't simply false; there is something else that is wrong with them. For such unnegated utterances, as well as for negated ones in which the negation cannot be understood as stating presupposition failure, we will stick with the proposal made above that when the context excludes the truth of a presupposition, the interpretation aborts. In such cases no coherent interpretation will be forthcoming. (In such cases, we noted, only a conditional interpretation might be possible for the interpreter, in which he goes along with the speaker or author in spite of his knowledge that the presupposition is false.)

In one sense this is a confirmation of the Frege-Strawson view that when a presupposition fails without being cancelled, then no proposition is expressed and no truth value determined. But in our set-up this is no 'threat to the logic'. In this setting logical notions such as logical entailment are directly defined for logical forms (i.e. for the DRSs that are the 'formulas' of our Logical Form Formalism). The usual definition of logical entailment (or 'logical consequence') between DRSs, as preservation of truth in models, imposes on the LFF the logic that, to the extent one can tell, Frege was after. We have seen how the semantics of the LFF can be extended to fragments of natural languages: Methods can be developed for assigning formulas of the LFF as logical forms to sentences of given natural language fragments, with the stipulation that the semantics of the sentences of the fragment is to be that of the logical forms assigned to them. (For instance, a sentence S1 from the fragment can be defined to logically entail another sentence S2 iff the logical form of S1 logically entails the logical form of S2.) In this way we obtain a classical logic for sentences from the fragment. Sentences from the fragment with improper descriptions – descriptions that cannot be verified with our without accommodation and that aren't cancelled either – won't be assigned a logical form and so cannot stand in logical relations to other sentences. So they cannot affect the logic of the fragment, and in particular cannot do any 'damage' to it. The only 'negative' consequences for the logic of the fragment will be that it is a partial logic, in the sense that it doesn't extend to all its grammatically well-formed sentences. But for those sentences that are included the logic is just that of their logical forms.[30,31]

[30]The argument above contains one obvious oversimplification. Whether the presupposition triggered by a definite description can be justified often depends on the utterance context. This, we have seen, is true of Presupposition Theory's 'Urexample', the Russell sentence 'The king of France is bald': Utterances of this sentence between the seventh and the eighteenth century were or would have been unproblematic. But this is not so for utterances made at the time when Russell came up with this example or at any time thereafter. If the context dependence of presupposition justification is taken into account, then we have to acknowledge that which sentences of a given fragment get a logical form assigned to them and thereby enter into logical relations may vary from context to context. But even if that is true, it remains true as well that for any given utterance context C the set of those sentences of the fragment that are assigned a logical form in C will have the logic imposed upon them by their logical forms, and that will always be the same logic, even if the set of sentences to which it applies varies as a function of C.

[31]There is an aspect of Frege's worry that these considerations cannot set aside. When natural language is used in doing mathematics, definite descriptions will often be used to denote mathematical objects (e.g. natural numbers). Whether such descriptions properly denote – i.e. whether its descriptive content has a unique satisfier – is sometimes very hard to determine. In fact, we know form the work of Gdel, Turing and Church that there cannot be a general algorithm for answering such satisfaction questions; for an algorithm there will be some such questions that are too hard for it to solve. In practice this means that when we use natural language in doing mathematics we will often not be in a situation to determine whether a given sentence we consider using is one with a well-defined logical former not. A good way to avoid this kind of problem is to make the presuppositions of definite descriptions explicit as separate propositions. When the presupposition is not satisfied, then we have a false but well-defined statement and so the conjunction of it and the sentence with the given definite description occurrence is also false irrespective of what one might want to say about that sentence. And when the presupposition is satisfied, then the statement is true, the sentence with the definite description is well-defined and the truth value of the conjunction of the two will be that of the sentence.
This comes close to Russell's Theory of Descriptions. The only difference is one of information organization. A central feature of DRT, and of the approach in PART II of these Notes as a special case of that. is the incremental interpretation of multi sentence texts and discourses. This makes it possible to split the content of a sentence S with a definite description into a first sentence S1 that states the satisfaction conditions of the description and a second sentence S2 that states the content of S on the assumption that S1 is true. This way of thinking about incremental logical forms as capturing the content of certain

In an approach like ours, in which preliminary semantic representations are computed from syntactic structures, it is not hard to incorporate certain forms of Horn negation, which target various non-truth-functional aspects of the input representations on which those negations act. This is so in particular for Horn negations that target presuppositions. The main problem for such an implementation, however, would be to formulate criteria which decide when a negation is used in this way. We have seen that such interpretations do not come for free – the context has to provide clear clues that such an interpretation is intended. But to determine exactly what forms this triggering information can take is something else. This is one of the tasks that are left for future work.

Unique Satisfaction?

In the discussion above we have been assuming that resolution of the identification presupposition of a definite description requires verifying that its descriptive content has a unique satisfier. But when that requirement is taken literally, it is plainly wrong. In fact, in this simple form the requirement is so absurd that nobody could be supposed to have seriously entertained it. Counterexamples to the unique satisfaction requirement in this simple form, starting with Strawson's examples like 'The book is on the table', are everywhere and it is hard to believe that anyone could have been unaware of them.

Nevertheless, despite the obviousness and the ubiquitousness of such examples there is something to the unique satisfaction account that seems right, and right also for them. When descriptions like *the book* and *the table* do their job, picking out some particular book and some particular table, that is because their overt descriptive contents – the predicate 'book' and the predicate 'table' – succeed in selecting unique satisfiers from some independently restricted search domain, a domain that contains just one book and just one table. Such restricted search domains are somehow determined by the contexts in which the descriptions are used. According to this diagnosis definite descriptions can succeed in referring because their descriptive content is uniquely satisfied in the search domain determined by the context in

---

sequences of natural language sentences is alien to the general views about the relation between language and logic that were prevalent at the time when Russell conceived his Theory of Descriptions and against the background of which the Theory of Descriptions should be understood. A logical form (DRS) one part of which states the unique satisfaction conditions of the descriptive content of the definite description of S while a second part states the truth conditions of S on the assumption that these unique satisfaction conditions are true is precisely the logical form for the Russellian analysis of S.

which they are used.

The claim that this is the right diagnosis of how definite descriptions succeed in selecting their referents has not been without contestants. Not everybody seems to agree with tho explanation of how descriptions like *the book* or *the table* can do their job. The objections deserve careful consideration, but that is a task we reserve for the next section. Our conclusion t here will be that not all definite descriptions should be construed as imposing a unique satisfaction requirement relative to some independently determined search domain. For now, however, we will focus our attention on descriptions that do impose such a requirement.

4.2.3.1 Constructing identification presuppositions for definite descriptions

Just as we did for proper names and third person pronouns, we assume that it is the determiner constituent of the DP that is responsible for introducing its identification presupposition. Thus the identification presupposition of a definite description is triggered by the determiner *the*. The central part of the formal specification of our account of definite descriptions is therefore the formulation of the lexical entry for this determiner.

The lexical entry for *the* is like those for proper name and pronoun DPs in that it turns the non-presuppositional representation of the NP constituent into the identification presupposition of the DP. One difference with the earlier two entries is that the identification presupposition of a definite description contains a subsidiary presupposition for the contextual domain restriction predicate C, within the extension of which the descriptive content must find its unique satisfier. There is also a (potential) second difference between the presuppositions triggered by *the* and the presuppositions triggered by names and pronouns. This second difference has to do with the principles that govern presupposition resolution. For proper names and pronouns we have assumed that their presuppositions are referential presuppositions whose resolution requires finding a dref that can serve as the antecedent for the question-marked dref of the presupposition. In particular, the anaphoric resolution of a pronoun presupposition requires that such a dref can be found in the locally available context. The identification presuppositions of definite descriptions differ in that they require unique satisfaction of their descriptive content, consisting of their explicitly given descriptive content in conjunction with the value for the context predicate dref C that has been determined through resolution of the subsidiary presupposition. (The identification presupposition for C must be resolved before resolution of the identification presupposition for the definite description can be considered.) When stated in this way, resolution of the main presupposition of a definite description sounds like verification of a proposition (the 'unique satisfaction presupposition') rather than the identification of an antecedent, and of course that is in keeping with the way in which the presuppositions of definite descriptions have been thought of for most of the modern history of presupposition theory, starting with Frege.

The implementation of this informal description of satisfaction requirements for definite descriptions presents us with a mild quandary. Should we follow the tradition and treat the presuppositions of definite descriptions as propositional presuppositions or go for uniformity in our treatment of definite noun phrases and treat the presuppositions of definite descriptions like those of the definite DPs that have been discussed in the last few sections as referential presuppositions? From a purely formal point of view either option is possible. I opt for uniformity. So the identification presuppositions of definite descriptions will be formally treated as referential presuppositions, like those for other types of definite DPs. But that of course doesn't alter the fact that resolving these presuppositions amounts to verifying the unique satisfaction condition. That may be possible even though the context DRS as given doesn't have an accessible dref that represents the intended referent. So, if we are to treat the identification presuppositions of definite descriptions as referential presuppositions, which are resolved by identifying their question-marked dref with one from the discourse context, then it must be possible to add this antecedent dref once unique satisfaction of the descriptive content has been established.

Allowing such dref introductions is of course a way of admitting that the presuppositions of definite descriptions really *are* propositional, whereas those for third person pronouns are not. The decisive point is that resolution of the identification presupposition of a definite description can make use not only of accessible drefs that belong to the locally available context by virtue of how this context representation has been constructed from the input LF or LFs but also drefs that can be added to this representation on the basis of inferences drawn from this context. This is part of the story about Partee's ball example (2.52) which we discussed in Part I: The constraints on pronominal anaphora that are revealed by the pair of discourses in (2.52) disappear when the pronouns are replaced by suitable definite descriptions. For more discussion see Section 4.3.5.

The entry for *the* in (4.94) below is exclusively for singular *the*, i.e. for the definite article *the* that occurs as determiner of singular noun phrases. There is no good justification for this restriction. It is motivated solely by the general decision in Part II of the Notes to restrict attention to singular noun phrases. As a matter of fact there are strong reasons for preferring a joint treatment of singular and plural *the*. For these are not just different uses of what looks like the same word; the most natural analysis of singular and plural *the* is as instances of what is a single word semantically as well as morphologically. The semantic contribution of *the* can then be stated as creating a DP that refers to the total set of satisfiers of the descriptive content.

The form this takes is that of a presupposition triggered by *the* to the effect that there is at least one satisfier. The difference between singular and plural definite descriptions is then that the number feature 'sing' contributes the information that this set is a singleton (or an atomic individual, if one adopts a mereological ontology), whereas the information contributed by the feature 'plur' is that the set consists of two or more elements (or is a non-atomic individual).

The lexical entry for singular *the* is given below in (4.94). Once again, the entry takes the general form we have been using for the lexical entries of operators, and like the entries for the determiner features of proper name and pronoun DPs the operator turns the non-presuppositional representation of the NP complement into the identification presupposition of the output representation for the DP. (Note well, however: The NP part of a definite description can be of unbounded complexity because the NP may contain relative clauses of arbitrary nesting depth. This means that in general the NP representation may have presuppositions of its own. These become subordinate presuppositions of the identification presupposition that is the output of applications the *the* operator, jointly with the new C-presupposition.)

(4.94) (lexical entry for singular *the*)

the (Det)

Sel. Restr:

Sem.Repr: $< x_{ref}, \beta_2, ., \beta_n \mid < \{PR_1, ., PR_m\}, K >> \rightsquigarrow$

$$\left\langle x_{ref} \mid \left\langle \left\{ \left\langle C, \beta_2, ., \beta_n \mid \left\langle \left\{ \begin{array}{|c|} \hline C? \\ \hline C(x) \\ \hline \end{array}, PR_1, .., PR_m \right\}, K' \right\rangle \right\rangle \right\}, \begin{array}{|c|} \hline \\ \hline \\ \hline \end{array} \right\rangle \right\rangle$$

Here K' is the DRS that represents the unique satisfaction presupposition. Its Universe consists of the question-marked dref x, together with the drefs $y_1, ..y_r$ in case there are any such drefs in this Universe of K. Its Conditions say that x' is the unique satisfier of the contextually reinforced predicate provided by the DRS K. That is, K' has the following form (4.95). (RestrDescrCont is the union of the Condition Set of K and the Condition '$C(x)$'.)

(4.95)
$$\begin{array}{|c|} \hline x? \quad y_1 \ \ldots \ y_r \\ \hline \text{RestrDescrCont} \\ \begin{array}{|c|} \hline x' \\ \hline \text{RestDescCont}[x'/x] \\ \hline \end{array} \left\langle \begin{array}{c} \forall \\ x' \end{array} \right\rangle \begin{array}{|c|} \hline \\ \hline x' = x \\ \hline \end{array} \\ \hline \end{array},$$

Evidently there are uses of definite descriptions where the descriptive content is meant to select a unique satisfier independently of context. Such descriptions are common when natural language is used in stating and discussing mathematical or scientific propositions, but they also occur in more mundane uses of language. Prominent examples of such 'self-sufficient' definite descriptions that are found in ordinary speech are descriptions involving superlatives, like *the tallest building in the US*, *the last emperor of China* and so on. For these no retrieval of a contextual predicate C is needed or wanted. We can subsume such cases under the general definition of identification presuppositions for definite descriptions that was made explicit in (4.94) by assuming that they are the cases in which the resolution of the C-presupposition takes the 'vacuous' form of resolving C to a tautological predicate such as $\lambda x.x = x$.

To show how the treatment of definite descriptions works out for 'incomplete' descriptions like *the table* or *the laptop* we go through the representation construction for sentence (4.96). Suppose that this sentence is uttered by a speaker A who is addressing an addressee B whom she is about to show Room 323, in which there is, at the time when she makes her utterance, just one table and just one desktop.

(4.96) The desktop is sitting on the table.

Suppose it is known to both A and B that there is just one table in Room 323 and just one desktop, and that this knowledge is part of the Common Ground between them. We may assume that this bit of Common Ground can be represented in the form shown in (4.97), in which *Room 323* is treated as a proper name of the room in question. (For simplicity the non-verbal predications involving the nouns *table* and *desktop* and the location predicate *in'* are represented as time-independent and *sit on* is treated as a transitive verb.)

(4.97)

$$
\begin{array}{|c|}
\hline
r \quad y \quad z \\
\hline
\text{Named}(r, \textit{Room 323}) \quad \text{desktop'}(y) \quad \text{table'}(z) \\
\text{in'}(y,r) \quad \text{in'}(z,r) \\
\begin{array}{|c|}\hline y' \\ \hline \text{desktop'}(y') \\ \text{in'}(y',r) \\ \hline\end{array} \left\langle \begin{array}{c}\forall \\ y'\end{array} \right\rangle \begin{array}{|c|}\hline \ \\ \hline y' = y \\ \hline\end{array} \\
\begin{array}{|c|}\hline z' \\ \hline \text{table'}(z') \\ \text{in'}(z',r) \\ \hline\end{array} \left\langle \begin{array}{c}\forall \\ z'\end{array} \right\rangle \begin{array}{|c|}\hline \ \\ \hline z' = z \\ \hline\end{array} \\
\hline
\end{array}
$$

As input LF for the preliminary DRS construction for the mentioned utterance of (4.97) we assume (4.98).

(4.98)

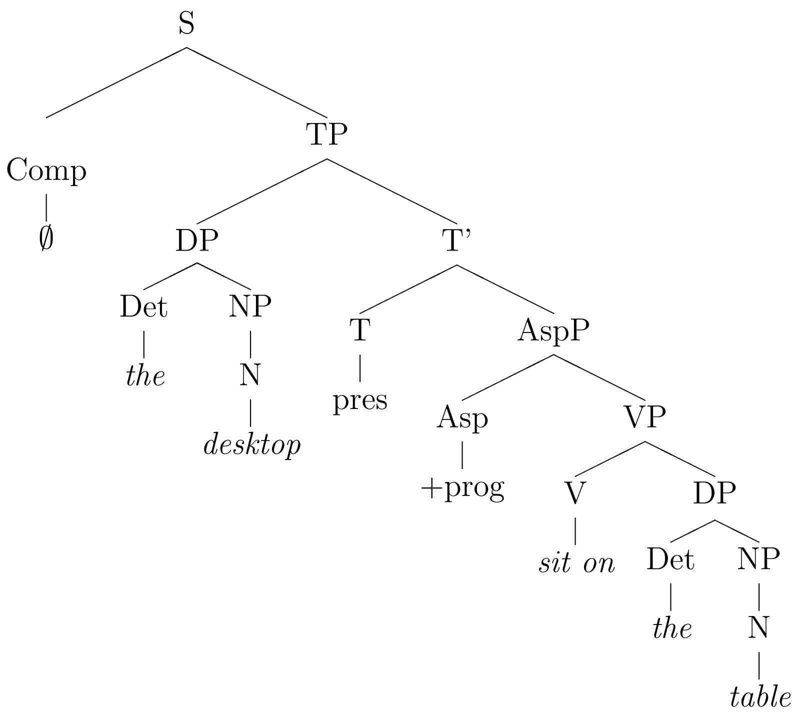

We assume the by now familiar type of lexical entry for the common nouns *desktop* and *table*. The semantic representation for the DP *the table* – resulting from applying the *the*-operator to the representation of the NP *table* (which coincides with the lexical entry for the noun *table*) – is given in (4.99) and the VP-representation that results from combining this representation with that of the 'verb' *lie on* in (4.100).

(4.99)

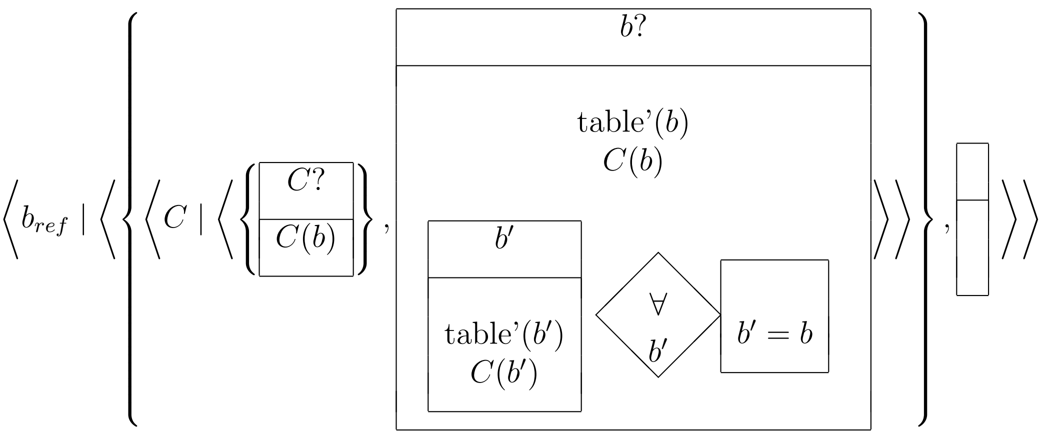

(4.100)

$$\left\langle e_{ref}, b \mid \left\langle \left\{ \left\langle C \mid \left\langle \left\{ \begin{array}{|c|}\hline C? \\ \hline C(b) \\ \hline \end{array} \right\}, \begin{array}{|c|}\hline b? \\ \hline \text{table'}(b) \\ C(b) \\ \begin{array}{|c|}\hline b' \\ \hline \text{t'le'}(b') \\ C(b') \\ \hline \end{array} \left\langle \begin{array}{c} \forall \\ b' \end{array} \right\rangle \begin{array}{|c|}\hline b' = b \\ \hline \end{array} \\ \hline \end{array} \right\rangle \right\rangle \right\}, \begin{array}{|c|}\hline \\ \hline e : \text{'s-n'}(\underline{x}, b) \\ \hline \end{array} \right\rangle \right\rangle$$

The subject phrase *the desktop* is represented in the same way as the 'direct object' DP *the table* (see (4.99)). Furthermore, the construction of the representations of VP, AspP, T' and TP are familiar. The result of these steps is the representation in (4.101).

(4.101)

$<t, s_{ref}, b, a \mid$

$$< \{ \left\langle C \mid \left\langle \left\{ \begin{array}{|c|}\hline C? \\ \hline C(b) \\ \hline \end{array} \right\}, \begin{array}{|c|}\hline b? \\ \hline \text{table'}(b) \\ C(b) \\ \begin{array}{|c|}\hline b' \\ \hline \text{t'le'}(b') \\ C(b') \\ \hline \end{array} \;\; \langle \forall\, b' \rangle \;\; \begin{array}{|c|}\hline b' = b \\ \hline \end{array} \\ \hline \end{array} \right\rangle \right\rangle,$$

$$\left\langle C' \mid \left\langle \left\{ \begin{array}{|c|}\hline C'? \\ \hline C'(a) \\ \hline \end{array} \right\}, \begin{array}{|c|}\hline a? \\ \hline \text{desktop'}(a) \\ C'(a) \\ \begin{array}{|c|}\hline a' \\ \hline \text{d't'}(a') \\ C'(a') \\ \hline \end{array} \;\; \langle \forall\, a' \rangle \;\; \begin{array}{|c|}\hline a' = a \\ \hline \end{array} \\ \hline \end{array} \right\rangle \right\rangle \},$$

$$\begin{array}{|c|}\hline \\ \hline t \subseteq n \;\; t \subseteq s \;\; C(b) \;\; C'(a) \\ s\text{: PROG}(\hat{}e. \begin{array}{|c|}\hline e \\ \hline e\text{: sit-on'}(a,b) \\ \hline \end{array}) \\ \hline \end{array} >>$$

Next, those drefs that occur in the store of (4.101) but not in any of its presuppositions are transferred to the Universe of the non-presupposiitonal DRS which follows the store. This transforms (4.101) into (4.102).

(4.102)

$<b', a' \mid$

$$< \{ \left\langle C \mid \left\langle \left\{ \begin{array}{|c|}\hline C? \\ \hline C(b) \\ \hline \end{array} \right\}, \begin{array}{|c|}\hline b? \\ \hline \text{table'}(b) \\ C(b) \\ \begin{array}{|c|}\hline b' \\ \hline \text{t'le'}(b') \\ C(b') \\ \hline \end{array} \left\langle \begin{array}{c} \forall \\ b' \end{array} \right\rangle \begin{array}{|c|}\hline b' = b \\ \hline \end{array} \\ \hline \end{array} \right\rangle \right\rangle,$$

$$\left\langle C' \mid \left\langle \left\{ \begin{array}{|c|}\hline C'? \\ \hline C'(a) \\ \hline \end{array} \right\}, \begin{array}{|c|}\hline a? \\ \hline \text{desktop'}(a) \\ C'(a) \\ \begin{array}{|c|}\hline a' \\ \hline \text{d't'}(a') \\ C'(a') \\ \hline \end{array} \left\langle \begin{array}{c} \forall \\ a' \end{array} \right\rangle \begin{array}{|c|}\hline a' = a \\ \hline \end{array} \\ \hline \end{array} \right\rangle \right\rangle \},$$

$$\begin{array}{|c|}\hline \text{t} \quad \text{s} \\ \hline t \subseteq n \quad t \subseteq s \\ s\text{: PROG}(\hat{}e. \begin{array}{|c|}\hline e \\ \hline e\text{: sit-on'}(a,b) \\ \hline \end{array}) \\ \hline \end{array} >>$$

The two identification presuppositions in (4.102) can be resolved only after resolution of the C- and C'-presuppositions on which they depend. In the context in which we have been assuming that (4.96) is uttered the obvious resolution for both predicates C and C' is ‘objects in Room 323’. The result of these resolutions of the C- and C'-presuppositions can be represented as shown in (4.103) and (4.104). These representations make a modest use of lambda-abstraction[32]. With this notation we can represent the predicate ‘objects in Room 323’ as $\lambda v'$.in’(v',z). Since we represent the domain restriction presuppositions also as referential presuppositions, we need a dref in the context DRS with which the question-marked drefs C and C' can be identified. So we need a principle that allows us to extend the context DRS with such a dref. The principle we adopt is that any predicate that can be defined on the basis of information entailed by the context DRS can be used as ground for the introduction of such a dref, and that the discourse context representation may be extended with this dref together with Conditions to the effect that it represents this predicate. For the case at hand, and assuming that C'' is the predicate dref introduced into (4.97) on the strength of this principle, the extension of (4.97) (provoked by the resolution of the presuppositions for C and C' and licensed by the principle) is as in (4.103).

(4.103)

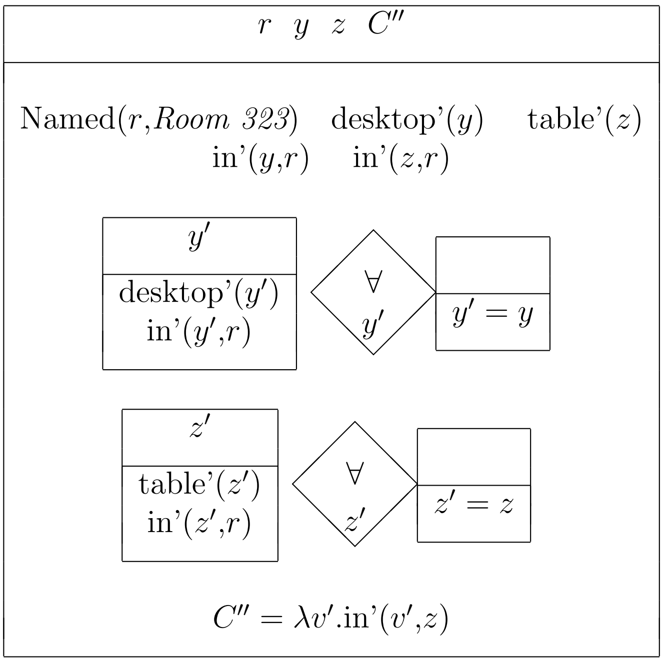

Resolution of the C- and C'-presuppositions can now take the form of setting both C and C' equal to C''. These equations are added to the non-presuppositional DRS of the preliminary representation (4.102). In the wake

[32]Adding this much λ notation to DRS languages is unproblematic. For details see (Kamp et al. 2011)

of these identifications C and C' can now be transferred to the Universes of the non-presuppositional parts of the identification presuppositions for *the desktop* and *the table*, and the C- and C'-presuppositions can, now that they have been resolved, be eliminated. This gives us the (still preliminary) representation in (4.104).

(4.104)$<b, a \mid$

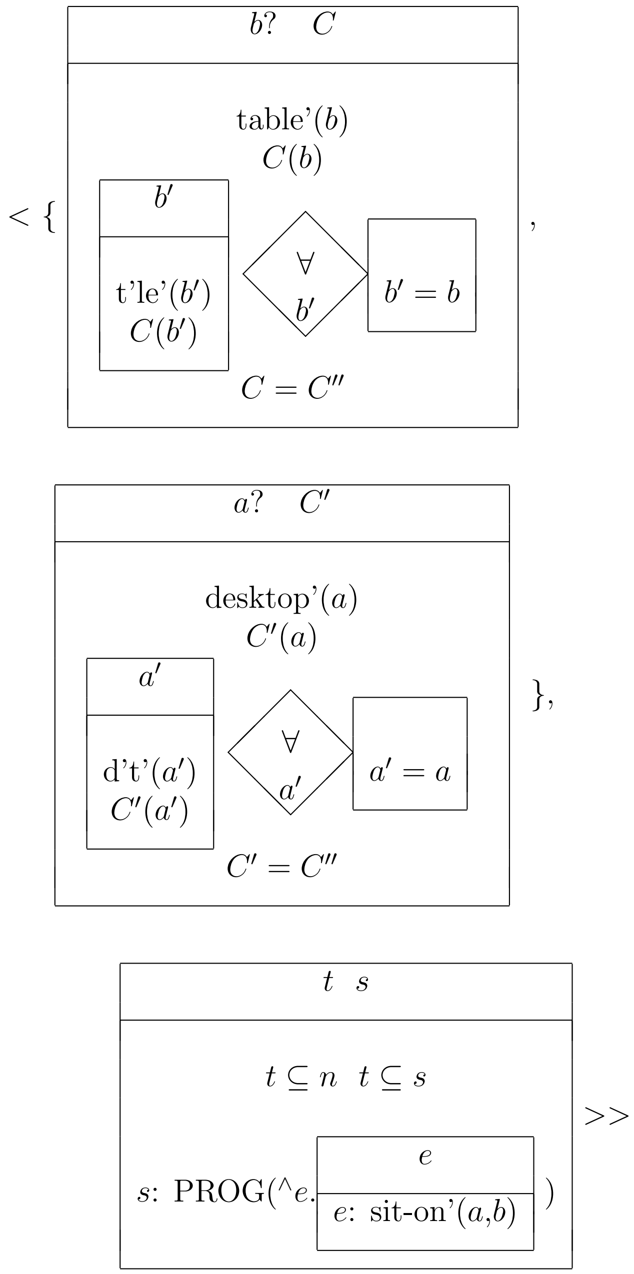

The identification presuppositions for the descriptions *the desktop* and *the table* in (4.104) are still to be resolved. But note that on the basis of the extended discourse context (4.103) and the identification of C with C'' we can infer that there is a unique satisfier of the conjunction of the Conditions 'table'(b)' and 'C(b)' (i.e. that the total content of the C-presupposition of (4.104) is satisfied). So, according to the principle for the resolution of definite descriptions mentioned earlier we are entitled to add a dref a' to the Universe of the context DRS (4.103) and Conditions to its Condition Set which state that a' is the unique satisfier; and. likewise, we may add a dref b' and Conditions to the effect that b' represents the unique table in Room 323. These additions transform (4.103) into (4.105).

(4.105)

$$
\begin{array}{|c|}
\hline
r \quad y \quad z \quad C'' \quad a' \quad b' \\
\hline
\text{Named}(r, \textit{Room 323}) \quad \text{desktop'}(y) \quad \text{table'}(z) \\
\text{in'}(y,r) \quad \text{in'}(z,r) \\
\begin{array}{|c|}\hline y' \\ \hline \text{desktop'}(y') \\ \text{in'}(y',r) \\ \hline\end{array}
\;\langle \forall\; y' \rangle\;
\begin{array}{|c|}\hline \\ \hline y' = y \\ \hline\end{array} \\
\begin{array}{|c|}\hline z' \\ \hline \text{table'}(z') \\ \text{in'}(z',r) \\ \hline\end{array}
\;\langle \forall\; z' \rangle\;
\begin{array}{|c|}\hline \\ \hline z' = z \\ \hline\end{array} \\
C'' = \lambda v'.\text{in'}(v',z) \\
\text{desktop'}(a') \quad C''(a') \\
\begin{array}{|c|}\hline a'' \\ \hline \text{desktop'}(a'') \quad C''(\text{a''}) \\ \hline\end{array}
\;\langle \forall\; a'' \rangle\;
\begin{array}{|c|}\hline \\ \hline a'' = a' \\ \hline\end{array} \\
\text{table'}(b') \quad C''(b') \\
\begin{array}{|c|}\hline b'' \\ \hline \text{table'}(b'') \quad C''(\text{b''}) \\ \hline\end{array}
\;\langle \forall\; b'' \rangle\;
\begin{array}{|c|}\hline \\ \hline b'' = b' \\ \hline\end{array} \\
\hline
\end{array}
$$

The presuppositions for *the desktop* and *the table* can now be resolved by setting a equal to a' and b to b'. Formally this leads to the addition of the equations '$a = a'$' and '$b = b'$' to the Condition Set of the non-presuppositional DRS of (4.104). The drefs a and b may then be transferred from the store to the Universe of this DRS and the identification presuppositions can then be eliminated in their turn. These operations turn (4.104) into (4.106).

(4.106)

$$\begin{array}{|c|}\hline t\ \ s\ \ a\ \ b \\ \hline \\ t \subseteq n\ \ t \subseteq s \\ \\ s\text{: PROG(}^\wedge e.\begin{array}{|c|}\hline e \\ \hline e\text{: sit-on'}(a,b) \\ \hline \end{array}\text{)} \\ \\ a = a'\ \ b = b' \\ \\ \hline \end{array}$$

(4.106) is an improper DRS because of the free occurrences of the drefs a' and b'. This problem is resolved by merging (4.106) with the context DRS (4.105).

4.2.3.2 Definite Descriptions and Anaphora

One of the remarkable facts about the history of the theory of reference since Frege and Russell is the radical difference between the dominant views about pronouns and definite descriptions. Pronouns were seen and treated as variables – as the 'variables of natural language' in the words of Quine – and definite descriptions as 'referential expressions', which select their referents via their descriptive content. This dichotomy seems particularly implausible when we turn to context in which pronouns and definite descriptions appear to be competing for the same tasks. Compare the sentences in (4.107).

(4.107)a. If Pedro owns a donkey he beats it.

b. If Pedro owns a donkey and a mule he beats it.

c. If Pedro owns a donkey and a mule he beats it, but loves it.

d. If Pedro owns a donkey and a mule he beats the donkey.

e. If Pedro owns a donkey and a mule he beats the donkey but loves the mule.

f. If Pedro owns a donkey he beats the donkey.

(4.107.a) is the donkey sentence for which we presented an explicit treatment in Section 4.3.2. In this sentence the use of the pronoun *it* is fully felicitous. That is not so for the *it* of (4.107.b). The reason for this difference between (4.107.a) and (4.107.b) would appear to be that in (4.107.b) *a donkey* and *a mule* seem equally good choices as anaphoric antecedent for the pronoun. There just is no way one can see why *it* should be understood as anaphoric to *a donkey* or as anaphoric to *a mule*. Note well, though, that by itself this is not an explanation of why we do not perceive (4.107.b) as merely ambiguous, but as in some sense ungrammatical. We will come back to this point later, when we compare (4.107.b) with sentences in which a pronoun can be construed in two ways, – as anaphoric to one antecedent or as anaphoric to another – but where the judgment is simply that the sentence is ambiguous, but not ill-formed. (4.107.c) is much like (4.107.b). In principle this sentence should allow for four different readings with two possible resolutions for each of the occurrences of *it*. But in fact the sentence seems just as ill-formed as (4.107.b).

What matters at this point is how (4.107.b) and (4.107.c) compare on the on hand to the well-formed (4.107.a) and on the other to the well-formed (4.107.d) and (4.107.e), in which the occurrences of *it* in (4.107.b) and (4.107.c) have been replaced by definite descriptions. These examples seem to tell us, the definite descriptions *the donkey* and *the mule* and the pronoun *it* cover the anaphoric needs that arise in connection with the sentences in (4.107). Moreover, the way in which they do that is plain enough: Definite descriptions are to be used when the pronoun cannot differentiate between two or more possible antecedents, whereas their descriptive content does make the needed distinction. But when the descriptive content of the definite description isn't needed for this purpose and the pronoun will do as well, then the pronoun tends to be preferred. An example illustrating this last observation is (4.107.f). This sentence isn't exactly ungrammatical. But it is awkward and someone who reads the sentence may feel an itch to replace *the donkey* by *it*, turning (4.107.f) into (4.107.a).[33]

[33] There is an alternative to *it* in (4.107.a) that is pretty much as felicitous as the pronoun and better than the description *the donkey*. This is the demonstrative phrase *this donkey*. The reason why *this donkey* is better than *the donkey* in (4.107.f) has to do with the non-uniqueness implication of complex demonstrative phrases. For instance, the *this* of *this donkey* implies that the referent is one from a set of several donkeys to each of which the speaker might have been referring. To explain exactly how *this* points to the intended antecedent we would need to say more about the functions of *this* and *that* in

When looked at in the way we just did in connection with the sentences in (4.107), pronouns and descriptions may seem very much of a kind: expressions that are designed for the same tasks, but which bring different resources to these tasks, with the effect that sometimes a pronoun will do, whereas in others, where pronouns are not up to the task, there will be need for a suitably chosen description. But is there any conflict between this perspective and the view described at the outset of the section, according to which pronouns are like variables and definite descriptions are expressions that select their referents via their descriptive contents? The explorations of this section will lead us to the conclusion that strictly speaking there isn't a conflict: Anaphoric pronouns and anaphora descriptions find their antecedents through applications of the same general principles. But the way in which these principles are typically applied to pronouns and the way they are mostly applied to descriptions differ considerably, encouraging the impression that pronouns and descriptions are as different as the classical view takes them to be.
One reason why pronouns and descriptions aren't as different as the classical view may seem to suggest has to do with the question what unique satisfaction of descriptive content comes to. We saw in the last section that for most descriptions that we encounter in ordinary writing and speech unique satisfaction is a tenable principle only when it is qualified by the possibility of restricting the search domain: the descriptive content of the description has to be uniquely satisfied within a restricted search domain that is chosen in the light of information provided by the context. In our reconstruction of the interpretations of *the desktop* and *the table* in sentence (4.96) the context, represented by the DRS (4.97), was given by the situation in which the utterance of (4.96) was assumed to take place. It seems a reasonable hypothesis that the interpretations of the anaphoric descriptions *the donkey* and *the mule* in (4.107.d) and (4.107.e) take the same form, with as only difference that it is now the discourse context derived from the 'local context' provided by the *if*-clause that is responsible for the choice of search domain and the verification of the uniqueness conditions within that domain. (Our

non-anaphoric settings – deictic settings, in which there is a direct pointing to the referent itself, as physical object in the environment in which the utterance is made, and pictorial settings, in which there is an effigy of the referent, as when I say to you, pointing at a figure in a photograph, "This is my uncle Tim as a young man". Since the framework we are using in these Notes is unsuited for these non-anaphoric uses of demonstrative DPs, and because a proper account of their anaphoric uses is hardly possible when these non-anaphoric uses aren't considered as well, it seemed better to exclude demonstratives from our discussions altogether. For a discussion of demonstrative DPs within a DRT-based setting, see (Kamp 2001,2011). For a little more about why a proper treatment of demonstratives is impossible within the framework of these Notes see the paragraphs devoted to demonstratives towards the end of this section.

DRT-based framework is well-suited to bring out the similarity of the interpretation procedures for the non-anaphoric descriptions *the desktop* and *the table* on the one hand and the anaphoric descriptions *the donkey* and *the mule* on the other. This is because in either case resolution of the identification presuppositions of the descriptions is a procedure that makes use of the context *representation* (in the form of DRSs or of more complex structures in which the relevant DRSs occur as constituents).)

Applying by now familiar construction principles to the *if*-clause of (4.107.d,e) we obtain the DRS K_C in (4.108). (For some of the details of the construction see the part headed 'The Polymorphism of Conjunctions and Disjunctions' of Section 3.10.2.)

(4.108)
$$\begin{array}{|c|}\hline t \quad s_1 \quad s_2 \quad y \quad z \\ \hline t = n \quad t \subseteq s_1 \quad t \subseteq s_2 \\ \text{donkey'}(y) \quad \text{mule'}(z) \\ s_1\text{: own'}(p, y) \quad s_2\text{: own'}(p, z) \\ \hline \end{array}$$

> (N.B. The dref p that is introduced as referential argument for the proper name *Pedro* will end up in the Unviverse of the DRS that contains the conditional Condition of which (4.108) is the antecedent in its Condition Set.)

A simple assumption we can make for the search domain provided by K_C is that it consists of the individuals represented by the individual-representing drefs p, y and z that have been introduced as representatives of the DPs *Pedro*, *a donkey* and *a mule*. Since verification of the unique satisfaction conditions is to be by inference from K_C it is natural to identify this domain with the set $\{p, y, z\}$ of the representing drefs. The verification of the unique satisfaction of the descriptive contents 'donkey' and 'mule' then takes the form of showing that exactly one member of the search domain $\{p, y, z\}$ represents an individual that is a donkey and exactly one represents an individual that is a mule. These verifications are intuitively straightforward, but it is well to see exactly what the various assumptions are that they depend on. Some of the premises on which the verifications depend are items of 'world knowledge' that are not explicitly represented in K_C and a mild form of accommodation is needed as well.

The accommodation concerns the individual represented by p. Intuitively that individual is a human being. but how do we know this? There are two clues. One is that *Pedro* is typically used as a name for human males. And the other is that p occurs as the subject argument of the predicate 'own', which often, and perhaps typically, is a human. But these are no more than hints. In principle anything can be called 'Pedro', including your dog, your toy tortoise, your favored mule and so on. And owners need not be humans, they can also be companies and other kinds of legal 'persons'. Perhaps there is enough information implicit in K_C from which it can actually be inferred that p represents a human being. But there is no need to get hung up on the question whether the assumption that the 'Pedro' of the sentences in (4.107) is a man has the status of an inference or an accommodation: To be on the safe side let's assume that it is an accommodation.[34]

In addition to this accommodation the verifications also rely on robust bits of world knowledge, to the effect that no donkey is a mule and that no human being is either a mule or a donkey. With these bits of information unique satisfaction of the predicate 'donkey' within $\{p, y, z\}$ can now be established as follows: (i) that y represents a satisfier of the predicate 'donkey' is stated explicitly as part of K_C; (ii) that z does not represent a satisfier follows from the fact, recorded in K_C, that z represents a mule together with the world knowledge that no donkey is a mule; and (iii) that p does not represent a satisfier follows from the accommodation that Pedro is a human male together with the world knowledge that people aren't donkeys.

The unique satisfaction condition for *the mule* is verified analogously.

4.2.3.3 Anaphoric Resolution Strategies for Descriptions and Pronouns

Suppose that what we have so far said about sentences (4.107.d,e) is right. Then we can conclude that there are two things that the discourse context provided by the *if*-clause of these sentences supplies towards the interpretation of the descriptions *the donkey* and *the mule*: (i) the domain within which the unique satisfier has to be found and (ii) information about that domain which guarantees that the descriptive content is uniquely instantiated within it. With regard to (ii) there is nothing more that needs to be said right now.

[34] The information that p represents a male human is also needed to justify the interpretation of the occurrences of *he* in (4.107.d,e) as anaphoric to *Pedro*. So the accommodation does double duty, you might say.

But an observation should be made in relation to (i). A point that wasn't made explicit in our discussion of *the table* and *the desktop* in (4.96), but that becomes important when we want to compare the definite descriptions of (4.96) with those of (4.107.d,e) is this: what in (4.96) was supposed to determine the choice of search domain – the set of objects in Room 323 – is independent from the information which licenses the inference that within this domain there is exactly one desktop and exactly one table. The search domain was suggested by the fact that speaker and addressee are standing in front of Room 323 and that the speaker is about to show the addressee this particular room from the inside. The information about this domain that guarantees unique satisfaction within it of the predicates 'table' and 'desktop' derives from general knowledge of how rooms in the given building (or rooms on the given floor, or those on the given floor that serve a particular purpose, such as that of a faculty office) are furnished and equipped. From our discussion above of the sentences (4.107.b,c,d,e) it is less clear that the choice of search domain and the verification of unique satisfaction make use of distinct information sources. Both rely on the Discourse Context, the DRS K_C in (4.108) that results from applying the DRS construction algorithm to the *if*-clause that is shared by the sentences in (4.107). In particular, we assumed that the search domain is given by the set of all drefs that are introduced in the course of that construction, i.e. the set $\{p, y, z\}$. For the examples (4.107.b,c,d,e) this assumption works well enough. But as we will see, anaphoric resolutions of the identification presuppositions of definite descriptions sometimes require proper subsets of the domain determined by all drefs in the Universe of K_C. In such cases the choice of the search domain will depend on information or considerations that go beyond what is encoded in this representation.

We will return to the question of search domain choice below. For the moment just the following observation, relating to the sentences in (4.107.b,c,d,e): The interpretations of *the donkey* and *the mule* in (4.107) might also have been reconstructed as follows: among the subsets of the set of all drefs made available by the *if*-clause there is in particular the set consisting of the donkey-representing dref y and the mule-representing dref z. (This set arguably has a certain prominence because its drefs represent the conjuncts of the conjunctive DP *a donkey and a mule*. For all we know at this point this syntactic fact may suffice to qualify $\{y, z\}$ as possible search domain.) Obviously the predicates 'donkey' and 'mule' select unique satisfiers within this smaller, two-element set as well. So this set would also have done as search domain for the verification of the unique satisfaction conditions. This situation, that more than one set could have been chosen as search domain but where each of the possible choices leads to the same interpretation, is a familiar feature of domain restriction, for quanifiers as well as for definite descriptions.

I have just spoken, somewhat informally, of the search domain determined by the drefs y and z and referred to that set, somewhat sloppily, as '$\{y,z\}$'. We can also think of this domain in a different way, viz. as represented by a 'plural' dref – W, say – which represents the set of the individuals represented by y and z. More precisely, and adopting the mereological approach to the distinction between singulars and plurals advocated in Section 2, we can introduce W as the mereological sum $y \bigoplus z$ of the drefs y and z.[35]

In Section 2 it was observed that summation is one of a number of operations on DRSs which yield additional drefs that can serve as the antecedents of plural pronouns. This set of dref-creating operations is logically restricted in that it must be incapable of mimicking the operation of set subtraction. Some other operations from this set were mentioned in Section 2, but the question was left open whether the operations that have been mentioned characterize the set exhaustively. (See also the discussions of (Kamp & Reyle 1993), Ch. 4.) In what follows we will proceeed on the assumption that a complete set Op has been given and that Summation is one of the operations in this set. (Op will be assumed to be finite, an assumption that has been made implicitly all along but that never really mattered until now.) Given this set Op of operations it is possible to extend any DRS K to its 'existential closure':

[35]To repeat the essential features of the mereological approach: (i) instead of distinguishing between individuals and sets of individuals, mereology distinguishes between atomic individuals (the individuals of the set-theoretic approach) and non-atomic individuals, corresponding to sets of two or more individuals in the set-theoretic approach. (There are no counterparts to the singleton sets of Set Theory, or, put differently, no distinction between individuals and their singleton sets.) (ii) The denotations of plural definite and indefinite descriptions, as well as of most occurrences of plural pronouns, are non-atomic individuals, whereas atomic individuals are the denotations of singular pronouns and singular definite and indefinite descriptions whose nominal heads are count nouns. (iii) The basic relation of a mereological ontology is the part-whole relation, denoted as $\preceq$. This is a transitive, antisymmetric and reflexive relation between individuals (atomic and non-atomic). Other mereological relations and operations can be defined in terms of $\preceq$. For instance, the Summation operation $\bigoplus$ can be defined as the operation that when applied to individuals a and b yields as result $a \bigoplus b$ the smallest individual (in the sense of $\preceq$) such that $a \preceq a \bigoplus b$ and $b \preceq a \bigoplus b$. (iv) When attention is restricted, as it will be here, to the denotations of count DPs, all non-atomic individuals d will be 'uniquely decomposable into atoms'. Speaking somewhat informally, there will be a unique collection D of atomic individuals such that d is the 'sum' of the individuals in D. Often the collection D is a finite set $\{d_1, ..., d_n\}$, in which case $d = d_1 \bigoplus ... \bigoplus d_n$.

the DRS K' that results from iterated application of members of Op to K and that has the property that further applications of operations from Op do not lead to the introduction of new individuals or sets, i.e. to new drefs for atomic or non-atomic individuals that are not provably co-existential with drefs that have been introduced already. (When the Universe of K is finite, which it will be in all cases of existential closure considered here, then such a K' will be 'reached' in a finite number of operation applications. In such cases K' can be identified with the first DRS which stabilizes in this way that is reached when the operations from Op are iteratively applied to K in some predetermined order. We refer to the Universe of the DRS K' as 'Acc(K)'. 'Acc(K) may be read as:'the drefs made available by K'.) The intuition behind this terminology is that when K is the representation K_C of the discourse context for the interpretation of a definite description, then the representations of the possible search domains for its interpretation will be among the members of Acc(K_C). In fact, we adopt the following assumption:

(4.109)All drefs from Acc(K) are available as representations of possible search domains for the anaphoric interpretations of definite descriptions.

In the light of what has been said so far, it might be thought that only plural drefs from Acc(K) – i.e. those drefs that represent non-atomic individuals – qualify as representations of search domains. However, it will be argued below that the singular drefs from Acc(K) should also be included.[36]

[36]The identification of Acc(K_C) with the set of possible search domain choices for anaphoric interpretations of definite descriptions is applicable straightforwardly in, for instance, those cases where the definite description that is to be interpreted occurs in a single-clause sentence S2 that follows another sentence S1 in a 2-sentence discourse <S1,S2>. In such cases K_C is the DRS constructed for S1. But as we have seen, the contextual information that is available for the interpretation of anaphoric and other presupposition-triggering expressions cannot always be identified in this simple way with a single DRS. Rather, the available contextual information is in general the amalgamation of all the information stored along the 'accessibility projection line' that reaches up from the embedded position of a presupposition-triggering expression within a complex preliminary representation to the top level of that representation, in accordance with the definition of DRS accessibility. (In fact, as we noted earlier, the sentences in (4.107) are mild instances of this more complex situation, in that in the semantic representation for any of these sentences the dref p for Pedro will eventually end up in the Universe of the main DRS, which contains the conditional Condition triggered by the *if*-clause as a member of its Condition Set whereas the drefs introduced for the indefinite DPs *a donkey* and *a mule* belong to the Universe of the antecedent DRS of the conditional Condition. Thus p will not be part of the same DRS Universe as the drefs y and z introduced for the DPs

(4.110) gives a summary of where we have so far got with our analysis of anaphoric resolutions of the identification presuppositions of definite descriptions.

(4.110)Suppose that a definite description δ with descriptive content P is interpreted anaphorically in relation to the existential closure K_C of the available discourse context information.

a. The possible search domains for the resolution of the identification presupposition of δ are all and only those that are represented by drefs in Acc(K_C).

b. For any choice X from Acc(K_C) the identification presupposition for δ is *resolved relative to* X iff it is shown that within the non-atomic individual represented by X there is a unique satisfier of P. (That is: that the Condition 'There is a unique atomic individual that is part of X and satisfies P' is true.)

When anaphoric search domains are identified with the mereological sums represented by drefs from Acc(K_C), this also provides us with a handle on cases like (4.107.b,c) where pronominal anaphora is marginal or impossible. The remarkable thing about the sentences in (4.107.b,c), we noted, is that they are not just ambiguous, but ungrammatical, or nearly so. This appears to be closely related to the fact that the two possible antecedents for the occurrences of the pronoun*it* in these sentences are constituents of a single conjunctive DP *a donkey and a mule*. So far we have been assuming in our discussion of these sentences that the drefs y and z introduced for the DPs *a donkey* and *a mule* can be 'summed into' a dref W which represents the mereological sum of what is represented by y and z individually; in this way a representation that already has the drefs y and z can be extended to one that also has a dref W which is related to these two via the equation '$W = y \bigoplus z$'. But perhaps, this is not the right way to think of how the semantic representation of the first sentence of (4.107.b,c) is to be constructed from that sentence. Perhaps the DP *a donkey and a mule*, a single argument phrase of the sentence, should be treated as a semantic unit in the following sense:

a donkey and *a mule*.) In cases where the discourse context is spread over different levels of a complex DRS, the applications of the operations in *Op* tend to be more complicated. Nevertheless, for these more complicated cases it is also possible to define an existential closure in the form of a DRS K_C, and a set of drefs that are accessible as possible search domain representations can then again be identified with the Universe of that DRS. We won't give the somewhat cumbersome technical definitions of K_C and Acc(K_C) for the more general case. But I will assume that this has been done and that Acc(K_C) is well-defined also for these more complicated kinds of discourse context.

the DP as a whole gives rise to the introduction of a dref W which is then specified further, by the internal structure of the conjunctive DP *a donkey and a mule*, as the sum of two drefs that represent the denotations of its two conjunct DPs *a donkey* and *a mule*. But – this is the crucial consideration – this 'defining specification' of W might well be something that belongs to a lower, more deeply embedded level, which is not visible from the outside and therefore not available for the resolution of identification presuppositions for anaphoric expressions from the next sentence. More specifically, we may think of the definition of W as given by a DRS that contains the defining Condition '$W = y \bigoplus z$' as a Condition in its Condition Set and the drefs y and z in its Universe, but which is a sub-DRS of the DRS whose Universe contains W. Thus by modifying the DRS construction rules for conjunctive DPs along these lines the DRS for the *if*-clause of (4.107.d,e) we obtain is the one shown in (4.108). (Formal details of rule modification and DRS construction are omitted, but the radar ought to have little difficulty with filling in these details.) [37]

$$
(4.111)\quad
\begin{array}{|c|}
\hline
t \quad s \quad W \\
\hline
t = n \quad t \subseteq s \\
\begin{array}{|c|}
\hline
y \quad z \\
\hline
\text{donkey'}(y) \quad \text{mule'}(z) \\
W = y \bigoplus z \\
\hline
\end{array} \\
s\text{: own'}(p, W) \\
\hline
\end{array}
$$

The DRS which specifies the definition for W in (4.111) is on the one hand a genuine sub-DRS of the DRS (4.111) in the sense that the drefs in its Universe are not accessible at the level of the DRS containing it. But as far as truth conditions are concerned this sub-DRS makes the same contributions to the truth conditions of the DRS containing it that it would when merged with

[37]One difference between (4.111) and the earlier DRS (4.108) is that the new DRS has the single Condition 's : own'(p, W)' whereas the earlier one has the pair of Conditions 's_1: own'(p, y)' and 's_2: own'(p, z)'. This difference is inessential and irrelevant to the point we are discussing. The distributing semantics of the verb 'own' licenses the expansion of 's: own'(p, W)' into the two conditions 's_1: own'(p, y)' and 's_2: own'(p, z)' (where $s = s_1 \bigoplus s_2$). The details of this expansion belong to the theory of plural uses of verbs and thus to a chapter of semantics that has been excluded from PART II of the Notes.

that DRS. That is, the truth conditions of (4.111) are the same as those of (4.112).

(4.112)

$t \ s \ W \ y \ z$
$t = n \quad t \subseteq s$ donkey'(y) mule'(z) $W = y \bigoplus z$ s: own'(p, W)

Treating the content of the *if*-clause of the sentences in (4.107.b,c,d,e) is, you might say, merely a convenient formal device for keeping the drefs y and z out of sight. When there is no longer any need for this because all presupposition resolutions have been carried out, DRSs like (4.111) can be simplified through merging, leading to DRSs like (4.112).[38]

As noted earlier, pronominal ambiguity isn't in itself a ground for ungrammaticality. Consider (4.113).

(4.113) John told Bill that he had made a mistake.

The pronoun *he* in (4.113) has two possible interpretations, one on which it refers to John and one in which it refers to Bill. But although this sentence feels genuinely ambiguous, it doesn't come across as in any way ill-formed. When offered out of the blue as here, it conveys the impression that if one only knew more about the context one would know how to resolve the ambiguity. One obvious difference between (4.113) and (4.107.b,c) is that there is no reason why the drefs for *John* and *Bill* end up in positions in which they are inaccessible from the pronoun *he*. Another example that suggests even more strongly that what makes (4.107.b,c) so bad is that *a donkey* and *a mule* are conjuncts of a larger DP is the following pair (4.114.a,b) of variants of a well-known example of Rooth. (Rooth's original example served a very different purpose, having to do with information structure; see (Rooth 1992).)

[38]In order that the ploy exemplified in (4.111) works the way it is meant to it should be the case that once the drefs y and z occur in the subordinate positions they occupy in (4.111), neither they nor any other drefs representing the denotations of *a donkey* or *a mule* will end up in the set Acc(4.111). That this won't happen can be proved formally only on the basis of a fully explicit definition of how Acc(K_C) gets computed and that in its turn requires a fully explicit definition of the operations set *Op*. Since no such definition is given here, a corresponding proof cannot be given either. The reader will have to take it on trust that things do work out the way they should.

(4.114)a. An American farmer and a Canadian farmer met at a cattle market. He sold him a pig.

b. An American farmer met a Canadian farmer at a cattle market. He sold him a pig.

The first sentences of (4.114.a) and (4.114.b) are near-synonyms. But nevertheless there is a striking difference between (4.114.a) and (4.114.b). (4.114.b) seems perfectly acceptable – arguably with a slight preference for the interpretation according to which the American farmer sold a pig to the Canadian farmer; but the converse interpretation isn't impossible either. (4.114.a), on the other hand, seems quite awkward, whether we try to interpret it as saying that the American farmer was the one who sold the pig or that it was the Canadian farmer.

Perhaps (4.114.a) isn't quite as bad as the account just proposed for (4.107.b,c) should make us expect. But we shouldn't be too surprised by this, for there may be additional mechanisms that can endow the conjunct DPs with differentiating profiles. Suppose for instance that – implausibly, but never mind – it is known that American farmers are pig farmers and that Canadian farmers only grow crops, but occasionally buy a pig in order to make sausages for their private enjoyment. Then, in the light of these distinct roles that American and Canadian farmers are known to play, the interpretation according to which *he* is anaphoric to the American farmer and *him* to the Canadian farmer, does no longer seem all that bad.

4.2.3.4 More about Anaphoric Resolution Strategies for Definite Descriptions and Pronouns

We started this section by mentioning a widespread view about definite descriptions and pronouns according to which their interpretations involve very different principles. In the framework developed in the Notes, I repeat once more, the difference comes to this: anaphoric pronouns get their interpretation by being identified with accessible drefs in the discourse context – we will from now on refer to this strategy as 'Strategy 1' – and definite descriptions get interpreted by showing that their descriptive content has a unique satisfier; we call this strategy 'Strategy 2'.

We already saw in Section 3.4.4 that Strategy 2 is a tenable strategy only if it allows for restrictions of the 'search domain' within which the descriptive content is to find its unique satisfier; without this qualification the strategy would be hopelessly at variance with the facts. Once the need for such a proviso has been acknowledged, however, the hard question that then presents itself is what the mechanisms are that make restricted search domains available. (We briefly touched on this question in our discussion of sentences (4.107.d) and (4.107.e). But there we only observed that the wanted interpretation could be achieved irrespectively of whether the dref p for Pedro was or wasn't included in the search domain.) For the case of anaphorically interpreted descriptions we made a first proposal for the constraints on search domain choice: the possible search domains are those that are represented by drefs from the set $\mathrm{Acc}(K_C)$. But it was left open whether the choice of search domains is subject to additional constraints.

Whether or not there are such further constraints and whatever they may be like, one thing is clear. Once search domain choice is included as part of Strategy 2, this strategy becomes a *two-step strategy*, which consists of (a) using the discourse context to determine the search domain, and (b) showing, making use of the discourse context once more, that the descriptive content has a unique satisfier.

(4.115) gives a schematic description of the two strategies and their uses.

(4.115)1. There are two strategies for the anaphoric interpretation of third person pronouns and definite descriptions, Strategy 1 and Strategy 2.

2. Strategy 1 consists in choosing a suitable dref from the set $\mathrm{Acc}(K_C)$ of contextually available discourse referents as antecedent for the anaphoric DP.

3. Strategy 2 consists of two steps: (a) choosing a dref X from $\mathrm{Acc}(K_C)$ as representative of the search domain. (b) showing that the descriptive content of the anaphoric DP is uniquely satisfied within the search domain X.

4. Strategy 1 is used for pronouns, Strategy 2 for definite descriptions.

Two questions will dominate the remainder of this section:

Q1: Is there a perfect alignment between Strategy 1 and Strategy 2 on the one side and pronouns and definite descriptions on the other? (That is, is it true, as claimed under 4., that anaphoric pronouns are always interpreted via Strategy 1 and definite descriptions always via Strategy 2?)

Q2: Are Strategy 1 and Strategy 2 best seen as distinct and independent strategies or should they be regarded as alternative implementations of what is at heart a single more general strategy?

Definite Descriptions that seem to behave like Pronouns

The examples in (4.116) indicate that there may be no perfect line-up of the pronoun-definite description distinction and the Strategy 1-Strategy 2 distinction.

(4.116)a. One of the ten balls is not in the bag. It is under the sofa.
b. Nine of the ten balls are in the bag. It is under the sofa.
c. One of the ten balls is not in the bag. The missing ball is under the sofa.
d. Nine of the ten balls are in the bag. The missing ball is under the sofa.
e. One of the ten balls is not in the bag. The ball is under the sofa.
f. Nine of the ten balls are in the bag. The ball is under the sofa.

The pair (4.116.a,b) was discussed at length in Section 2. Let's recall the point it illustrates. The first sentences of (4.116.a) and (4.116.b) express the same proposition (in the sense that they have the same truth conditions); but nevertheless they provide different discourse contexts. The discourse context provided by the first sentence of (4.116.a) makes a dref available for the missing ball, which the pronoun *it* in the second sentence can use as antecedent. (Or, to put the matter using the terminology just introduced: in this case Acc(K_C) contains a dref that represents the missing ball and that can be used to interpret the pronoun.) The first sentence of (4.116.b) makes no such dref available, so the pronoun *it* in the second sentence cannot be interpreted in the way one would want to. Here Acc(K_C) does not contain a dref for the missing ball. The difference between the pronoun in (4.116.a) that can be interpreted as referring to the missing ball and the pronoun in (4.116.b) that cannot be so interpreted is thus explained assuming that pronouns must be interpreted according to Strategy 1.

This much just repeats what was said about (4.116.a) and (4.116.b) in Section 2. But (4.116.c,d) add a further twist to this. When the pronoun *it* is replaced by the definite description *the missing ball*, the difference between the two sentence pairs disappears. (4.116.c) and (4.116.d) are both good. This too can be explained on the basis of what has been said about definite descriptions and Strategy 2. The definite description *the missing ball* can select its intended referent from the search domain whose representative has been chosen from Acc(K_C). Let us assume that this is the set of ten balls represented by the plural dref Y that is introduced as referential argument for the DP *the ten balls* and that will therefore be a member of Acc(K_C). (Recall the treatment of the ball-sample in Section 2.) According to the information carried by the discourse context there is just one element in that set which satisfies the descriptive content 'missing ball' (provided that predicate is understood as 'ball missing from the bag').[39]

When described along these lines, (4.116a,b,.c,d) would seem to confirm the difference between pronouns as expressions to be interpreted via Strategy 1 and definite descriptions as expressions to be interpreted according to Strategy 2. However, the last two examples of (4.116), (4.116.e) and (4.116.f), show that as our story has been told so far this may not be quite right. In these last two sentence pairs we find what looks like a replication of the contrast between (4.116.a) and (4.116.b). The definite description *the ball* can be interpreted as referring to the missing ball in (4.116.e), but not in (4.116.f).

That *the ball* cannot be used to refer to the missing ball in (4.116.f) seems explicable in terms of the assumptions we have made. The crucial point here is that the search domains made available by the discourse context contributed by the first sentence of (4.116.f) all have the property that if they contain any ball at all, then they are supersets of the set of nine balls of which the first sentence says that they are in the bag. (I omit the details of the proof of this statement.) Given this claim, it is clear that the interpretation of *the ball* in the second sentence will abort, since there is no possible search domain in which the predicate 'ball' has a unique satisfier – let alone that the unique satisfier is the ball that is missing.

But if this is right, how then do we explain that *the ball* can be interpreted in the way it is meant to in (4.116.e)? This is what I take to be the right response to this question: Contrary to what we said earlier about the alignment between pronouns-definite descriptions and Strategy 1-Strategy 2, at least some definite descriptions can also be interpreted according to Strategy 1. If *the ball* in (4.116.e) is one of these, then the felicity of (4.116.e) can be accounted for in just the same way in which we could account for the felicity

[39] As often with search domain choice the search domain isn't fully determined. In the case at hand another choice for the search domain would have the set of all individuals introduced by the discourse context, including the mentioned bag. For the selection of a ball by unique satisfaction of the content 'missing ball' this would have made no difference.

of (4.116.a) with its pronoun *it*.

The description *the ball* in (4.116.e) is an example of a definite description that can be interpreted via Strategy 1 but that cannot be interpreted via Strategy 2 (since as we saw in connection with (4.116.f) there is no search domain with more than one individual within which the descriptive content 'ball' could select a unique satisfier). There are also cases – if admittedly they are fairly rare – where pronouns have to be interpreted by an application of Strategy 2. An example is provided by the mini-story in (4.117)

(4.117) The pub was popular with both men and women. At 8.30, when there already was a lively crowd, a couple came in that drew everybody's attention. He was wearing a bowler hat, she had a contraption on her head that was somewhere halfway between a crown and a cauliflower.

The discourse context for the interpretation of *he* and *she*, which is obtained by constructing a DRS for the first two sentences, has a dref for the couple, but no drefs for its members. (Again, this is a claim that requires formal demonstration, but that is something that cannot be provided here.) However, couples typically consist of a male and a female partner, and the pronouns *he* and *she* can be used to each select the partner from the couple that fits its descriptive content. Also note that (4.117) doesn't seem to differ in any significant way from the discourse that we get when we replace *he* by *the man* and *she* by *the woman.* Both for the pronouns *he* and *she* and for the descriptions *the man* and *the woman* interpretation involves choosing the dref for *the couple* as search domain and then using the descriptive contents to select their unique satisfiers within that domain.[40]

The last few examples have shown that the opposition between Strategy 1 and Strategy 2 does not line up perfectly with the opposition between pronouns and definite descriptions. Inasmuch as there is a line-up at all, this is

[40] There is a somewhat idiomatic flavor to (4.117), and other examples that according to what we have been saying ought to be just as good may seem less felicitous. Consider for instance (4.118).

(4.118) The car drove straight into John and his dog, who were standing on the sidewalk. He died but it survived.

Some people I have asked find (4.118) a little funny. But that may well be because the contrast it implies strikes them as not quite proper: the questions of life or death for an animal is somehow incommensurable with the life-or-death question for human beings. Another factor may be that *it* is awkward in contexts which seem to impose some kind of contrastive prominence on it. In (4.118) *it* seems to aspire to this kind of prominence because of the apparent contrast with *he*.

at best a general tendency; pronouns and descriptions insolubly tied to their preferred strategies. Moreover, now that we have made this observation, it seems natural to raise the following question: Are Strategy 1 and Strategy 2 really as radically different from each others their descriptions in (4.115) suggest?

Arguably they are not. There is a way of looking at the two strategies that makes Strategy 1 into a special case of Strategy 2. Suppose that among the permissible choices of search domains in executions of Strategy 2 there are also those represented by singular drefs from $\mathrm{Acc}(K_C)$ (i.e. drefs that stand for atomic individuals). Such a domain will consist of a single individual (in mereological terms: it will just *be* an atomic individual). Selecting the unique satisfier of some predicate P from such a 'domain' reduces to verifying that P is satisfied by the one individual the domain contains.

It will be clear that when Strategy 2 is applied in this way, it provides the wanted resolution for *the ball* in (4.116.e): choose the 'search domain' given by the dref for the missing ball, viz. the dref that was introduced as referential argument of the DP *one of the ten balls*. Then it will follow trivially from the information that has been included in the discourse context K_C about the individual represented by this dref that this individual satisfies the descriptive content 'ball'.

Moreover, if we understand Strategy 2 as including cases where a search domain is chosen that contains a single individual, then there is no longer any reason why pronouns cannot be interpreted using this strategy as well. In Section 4.3.2 we noted that English third person pronouns carry rudimentary descriptive content: 'female human' for *she*, 'male human' for *he* and 'non-human' for *it*. Assume that *it* contributes this content in (4.116.a). Then (4.116.a) can be interpreted by choosing the dref for the missing ball as search domain and verifying that this dref represents something that is not human. But that verification is nearly trivial, in view of the generally available knowledge that balls aren't human.

In fact, once we admit single individuals as search domains, we can define Strategy 1 as a special case of Strategy 2 – that case in which the chosen search domain consists of a single individual. Let us adopt this simplification. Then there is just one strategy for the resolution of the identification presuppositions of both pronouns and definite descriptions. Applying this strategy can take two distinct forms, depending on whether the chosen search domain is or is not a singleton. In the former case there is no further need to select one element from the chosen set, as there is only one element to begin with; all that the descriptive content of the description or pronoun contributes is the constraint that the one element of the search domain should satisfy it.

It might be objected that putting things this way is a little misleading. The role of the descriptive content should be seen as that of *guiding* the choice of the search domain from the beginning: choose a singleton search domain for which the second part of the procedure – verifying that the chosen element satisfies the descriptive content – will succeed. In this way the applications of our general strategy which involve the choice of a singular dref can be seen as a one-step procedure, such as Strategy 1 was described in (4.115).

When on the other hand the chosen search domain consists more than one element, then more work remains for the descriptive content. It must now select one element from the chosen set, and it can do that only if there is just one element in this set that satisfies the descriptive content. So the resolution of the identification presupposition is complete only when unique satisfaction has been established within the search domain. For such cases the description in (4.115) of Strategy 2 as a genuine two-step procedure seems right.

The upshot of this discussion can be summarized as follows. There is a single anaphoric resolution strategy for definite descriptions and pronouns. It consists of (a) choosing a search domain and (b) selecting a single element within that domain via unique satisfaction of descriptive content. But what this comes to heavily depends on whether the chosen domain consist of one or of more elements. When a one-element domain is chosen, then the strategy reduces to what we have been calling Strategy 1, and there is no reason why we shouldn't continue to call these applications of the general strategy by that name. The other applications, in which the domain consists of several elements, are in spirit pretty much like what we had in mind when we first described Strategy 2. So it seems reasonable to go on using 'Strategy 2' for these other applications.

At this point it may look as if the unification of Strategy 1 and Strategy 2 is little more than a purely formal exercise, which won't do much to help us better understand what anaphoric pronouns and descriptions have in common. I believe however that there is some real insight that can be culled form this way of seeing the two strategies as specializations of a single one, and I hope that what follows below will also convince the reader.

4.2.3.5 Accommodation

According to our last specification of Strategy 1 establishing satisfaction of the descriptive content by the element of the singleton domain is a necessary part of applying it. When we go back over the instances of pronoun interpretation that have been presented in these Notes, it is not hard to ascertain that they all conform to this requirement. Or better: they all conform to the requirement in one sense; but there is another sense in which some do not. One example where the requirement is fulfilled in the one sense but not in the other is one that we have encountered repeatedly. This is the use of the pronoun *he* to refer back to a preceding occurrence of the proper name *Pedro*. The interpretation of *he* as anaphoric to *Pedro* is justified only when it can be shown that *Pedro* refers to a human male. This assumption may be plausible enough and in the examples involving *Pedro* that have been discussed it would probably have been made it even if there was no occurrence of *he* or *him* to require it. But still it is an accommodation that needs to be made to justify the anaphoric link between name and pronoun.

Many accommodations that are strictly speaking required to make the resolution of an anaphoric DP perfect are of this innocent sort. They are assumptions that interpreters would be likely to make in any case and they almost feel like inferences, even if it is often not all that clear on what basis they are drawn. But there are also cases where the interpretation of the anaphoric expression can add about the denotation that it ends up sharing with its antecedent and do that because at an accommodation that it more or less forms upon the interpreter.

(4.119)a. One of the videos showed a conversation between a doctor and a patient. She was trying to tell him that he would only have another six months to live.

b. One of the videos showed a conversation between a doctor and a patient. He was trying to tell him that he would only have another six months to live.

c. One of the videos showed a conversation between a doctor and a patient. She was trying to tell him that she would only have another six months to live.

In both (4.119.a) and (4.119.b) the interpretation of the pronouns requires accommodation. In both there is a strong presumption, based on world knowledge about what typically goes on between doctors and patients, that the first pronoun of the second sentence must refer to the doctor and the second and third pronoun to the patient. To get these pronoun interpretations in (4.119.a) we have to accommodate the information that the doctor is female and the patient is male. Likewise, to get such interpretations in (4.119.b) the accommodation has to be that both the doctor and the patient are male. In each of these cases the formal reconstruction is something like this. Suppose that in the discourse context K_C constructed from the first sentence the drefs representing the doctor and the patient are x and y. In the case of (4.119.a) accommodation will add to K_C the Conditions 'woman(x)' and 'man(y)'. These Conditions can then be used to verify the interpretation of *she* as anaphoric to x and the interpretation of *him* and *he* as anaphoric to y. Likewise for (4.119.b).

(4.119.c) has been thrown in partly for the fun of it. But it also serves as a further illustration of the role that accommodation can play in such cases and of the constraints that the descriptive contents of pronouns can impose on what interpretations are possible. (4.119.c) has two possible interpretations, neither of which conforms to our expectations about exchanges between doctors and patients: (i) the doctor is trying to tell the patient that she, the doctor, has only six months to live (perhaps as part of suggesting to the patient that he look for another doctor); (ii) the patient is trying to tell the doctor that she (the patient) only has another six months to live. (Perhaps the patient knows this because of what some other doctor or doctors have told her and she wants this doctor to do something for her during the six months she has left.) The fact that these and only these are possible interpretations of (4.119.c)[41] is an indication of the fact that for both the doctor and the patient there are only two gender options each – male or female (sorry, the example is a little old-fashioned) – and so that there are only four possible gender accommodations that we can make for the pair of these two human beings.

Accommodation can be equally important in connection with the interpretations of definite descriptions. Consider the examples in (4.120).

(4.120) I had five red balls, four green ones and a blue one. Nine were in the bag.

a. * The ball was under the sofa.

[41] if we exclude the possibility that the pronouns could refer to individuals other than the doctor and patient that are mentioned in the sentence

b. $\surd$ The missing ball was under the sofa.

c. * The green ball was under the sofa.

d. $\surd$ The blue ball was under the sofa.

e. $\surd$ The largest ball was under the sofa.

f. $\surd$ The rubber ball was under the sofa.

(4.120) is a series of variants of the ball example paradigm in which the discourse context for the final sentences – these are the sentences explicitly marked (a) - (f) – is provided by a couple of preceding sentences, the first of which provides some additional information about the set of ten balls of which the second sentence tells us that nine are in the bag. The first continuation of these two sentences, (4.120.a), is another instance of the failure of *the ball* that we noted in relation to (4.116.f). That (4.120.b) is a good continuation, we also saw earlier (cf. (4.116.d)). The failure in (4.120.c) is of the same nature as the failure of (4.120.a): The first sentence tells us that there are four green balls in the search domain. So unique satisfaction is not only not confirmed, but is explicitly contradicted, just as in (4.120.a).[42] (4.120.d) on the other hand is fine, since according to the first sentence there is just one blue ball in the search domain.

But the variants that matter most for the point at issue are (4.120.e) and (4.120.f). Both of these seem intuitively acceptable, but evidently they involve accommodation. In (4.120.e) the required accommodation is that the ten balls aren't all the same size and that one of them is larger than all the others. In (4.120.f) it is that the balls weren't all made of the same material, that one was made of rubber and the others from some other material or materials. It is striking with how much ease these accommodations are made. But of course, when the interpreter has contradicting information,

[42](4.120.c) raises a question that also arises in connection with other examples, but this is a good point to draw attention to it. According to the picture that is emerging, *the green ball* is bad only when there is no legitimate choice of a search domain in which there is just one green ball (and in which it is that ball that is missing from the bag). How can we argue that no such search domain is available in the context provided by the first two sentences? In the light of the assumptions we have made, the way to prove this would have to be that of showing that no such domain is represented by a dref in Acc(K_C). This is something that can be proved only when a full specification of the operations in *Op* has been given, something we haven't done here. A formal proof can be given if we adopt the proposal for *Op* that is given in (Kamp & Reyle 1993). But even when the premises for such proofs are available in the precise form in which they are needed, 'impossibility' proofs of this sort tend to be laborious and can often they be quite hard.

then accommodation is no longer easy and in fact it becomes impossible: Accommodations that are incompatible with what the interpreter considers to be the case are out.[43]

(4.121) provides another illustration of how accommodation can help with the interpretation of anaphoric DPs, and also of when it cannot help.

(4.121)a. Pedro owns two animals that he uses to pull his cart. He loves it.

b. Pedro owns two animals that he uses to pull his cart. He loves the donkey.

The pronoun *it* in the second sentence of (4.121.a) is uninterpretable as referring to one of the two animals mentioned in the first sentence.[44] This is another instance of the phenomenon of which we have seen a number of examples. The description *the donkey* in (4.121.b) is different. (4.121.b) may sound a little surprising or abrupt. But one feels that the sentence pair can be made sense of by accommodating the information that one of the animals mentioned in the first sentence is a donkey and that the other animal is not. In fact, that is the assumption that someone who reads (4.121.b) will make almost unthinkingly, as a form of charity towards the speaker and because it is the only form that charity can take in this case.

The difference between (4.121.a) and (4.121.b) should be clear enough at this point, but it deserves to be explicitly stated once more. The difference between the unsuccessful *it* in (4.121.a) and the successful *the donkey* in (4.121.b) is not simply the difference between pronouns and definite descriptions. It is, rather, the difference between an anaphoric expression whose content can in principle select a unique satisfier within the chosen search domain – in this case the content is the predicate 'donkey' – and an anaphoric expression whose content is incapable of doing that because the content can be inferred to be true of several members of the search domain. (In the case of (4.121.a) the descriptive content of *it* is true of all members.) If it can be demonstrated that the descriptive content is multiply instantiated, then no accommodation about the search domain will be able to put matters right.

[43]But recall the qualification in footnote 28 in this connection.

[44]The pronoun *it* can of course be understood as referring to the state of affairs that consists in Pedro having two animals to pull his cart. But that is the only interpretation which (4.121.a) seems to allow for. Such uses of pronouns and other anaphoric DPs as referring to abstract objects like propositions, states of affairs an the like are set aside here. (See (Asher 1993).)

And if there no other search domain available for which this problem does not arise, then the anaphoric expression simply cannot get the interpretation it wants.

So much for the role of accommodation in the anaphoric resolution of the identification presuppositions of definite descriptions and pronouns. This brings us to the end of the examples of this section and the observations they were brought in to illustrate. In the remainder we will first reflect on the implications of what we have just gone through and then make a few more general methodologically tinted remarks on the unnatural limits that have been imposed on the material discussed in this section, sand more generally on the over-all constraints that have been imposed on the choice of topics in Section 4.3.

4.2.4 Wrapping up the central concern of this section

We started out with the assumption that there are two distinct anaphoric interpretation strategies for definite descriptions and third person pronouns. We referred to these as Strategy 1 and Strategy 2 and conjectured that the first is just for pronouns and the second just for definite descriptions. But then we found that we could unify these two strategies by defining Strategy 1 as a special case of Strategy 2. So we ended up with just one strategy which applies to pronouns and definite descriptions alike, but one with two sub-strategies, which for practical purposes are just the Strategy 1 and a complementary strategy, which can be identified with the original Strategy 2.

The unification that makes Strategy 1 and Strategy 2 complementary versions of a single strategy helps us to see that there is less of a difference between anaphoric pronouns and definite descriptions than is suggested by the classical view I mentioned at the outset of this section. That is arguably a good thing. But we should be careful not to let ourselves be seduced into underestimating the differences that can be observed in the actual use of pronouns and definite descriptions. One difference has to do with where intuitively the real work is done in resolving the presuppositions of these anaphoric expressions. When the 'search domain' is a singleton – as it is in applications of Strategy 1, and that is by far the most common strategy for the resolution of pronouns – then choosing that 'search domain' is in essence all there is to the resolution. The descriptive content of the anaphoric expression isn't without significance, but what work it does is best seen as taking the form of guiding the choice of this 'search domain'. When on the other hand the chosen search domain is not a singleton set, then its choice is more naturally described as the first of two steps, which needs completion by verifying that the descriptive content has a unique satisfier within the chosen domain and identifying that satisfier as the referent.

There is also a second intuitive difference between typical applications of Strategy 1 and Strategy 2. Verbal discourse has a great deal of structure. The simplest structure, but an important one no less, is the linear order of speech and texts: the expressions that are the substrings of spoken or written language utterances make up a linear sequence in which there is a certain distance between any two of them, defined by the number of expressions (some number ≥ 0) by which they are separated. But of course there is much more structure than that. Grammatically well-formed sentences have their internal syntactic structure; and there is also structure of a more pragmatic or rhetorical sort, such as topic-comment structure or the structure imposed by rhetorical relations. All these forms of structure have their impact on the use of pronouns and descriptions. Or, putting things more in the terms in which the discussion has been conducted in this section: they all play their part in determining which DP a speaker or author should choose for a given anaphoric task: should she use a pronoun or a definite description, and when a definite description, which description?

An often made observation is that pronominal anaphora is heavily dependent on distance. By and large anaphoric pronouns find their antecedents in close vicinity – elsewhere in the same sentence, or in the immediately preceding sentence, or perhaps within the sentence immediately before that one. Only rarely is a pronoun's antecedent to be found farther back than that. The evidential support for this observation is overwhelming, but it has proved difficult to turn it into a proper theory, in which the distance between a pronoun and some other DP makes a quantifiable contribution to deciding whether this second DP is the antecedent of the pronoun. The role that topic structure plays in the choice between pronouns and definite descriptions has been studied, with significant, but nonetheless only partial success, by Centering Theory ((Grosz, Joshi & Weinstein 1995), (Walker, Joshi & Prince 1998)). By and large there is a preference for using pronouns to refer back to antecedents with topic status, while non-topics are anaphorically resumed by definite descriptions or demonstratives. But the effects of topic-comment structure on the choice between pronouns and definite descriptions is compacted by the dynamics of such structure: As a discourse or text unfolds, topics 'shift': the role of topic shifts form one protagonist to another and such shifts can be linguistically realized in different ways, with different consequences for anaphoric options. Centering Theory has been right in paying close attention to this dynamic aspect of topic-comment structure; but, as said, its successes have thus far been partial Another important factor is focal contrast, which as we have seen disfavors the pronoun *it*, but also affects the choice between pronouns and descriptions in other ways.[45]

The importance of discourse structure for anaphora is visible in the form of certain definite descriptions. The descriptive contents of the descriptions occurring in the examples of this and the last section were all 'referent-related': they all expressed properties of the individuals that the anaphoric expression and its antecedent dref stand for. But not all descriptive contents are like this. Typical descriptions whose content is wholly or partly 'discourse-related' (rather than 'referent-related') are *the former*, *the latter*, *the last-mentioned ..*, as well as (often) *the first*, *the second* and so on. An anaphoric resolution of a DP like *the last-mentioned man* will have to zero in on the last DP in the given discourse or text that can be recognized as describing or referring to a man. A curious feature of such descriptions is that while their head nouns typically suggest that their denotations are 'ordinary individuals' – the DP *the last-mentioned man* clearly intends to refer to a man – the adjectival modifier is meant to be evaluated in relation to the structure of the discourse. The correct description of the resolution conditions for definite descriptions with discourse-related modifiers is an intriguing puzzle; but it is

[45]In the literature on reference and anaphora a free, frequent and deplorable use is made of the term 'salient': what qualifies one expression as anaphoric antecedent for some other expression, it is often claimed, is its 'salience', or its 'high degree of salience'. Within the setting in which the problems of anaphora have been discussed here salience would be relevant – assuming it is relevant anywhere – in connection with the choice of 'search domains' (whether singleton or non-singleton domains): only choices of 'salient' domains are possible domain choices. Here, as in other contexts in which the term 'salient' has been used, the best that can be said about its use is that it can be helpful as pointer in the direction of a problem. Recognizing the exact nature of the problem is another matter. And solving the problem once it has been properly recognized is another matter altogether.

Splitting anaphora resolution into the different tasks mentioned in our characterization of Strategy 1 and Strategy 2 may be seen as a first step in the direction of recognizing the nature of the problem: of decomposing it into a number of sub-problems that are defined in operational terms. Specifying the choice of search domains as limited to the set $\mathrm{Acc}(K_C)$, describing the nature of unique satisfaction and the various kinds of accommodation that can be involved in its verification may be seen as first steps in the direction of a solution . But they are only first steps. It should be plain from what has been said in the body of the text that more details about applications of Strategy 1 and Strategy 2 will have to be chartered before we can feel confident that we have a reasonable grasp of what 'salient' could mean in the context of DP anaphora.

a project in its own right and one that is to be taken up elsewhere.
The last few paragraphs have pointed at further factors that a good theory of the interpretation of anaphoric definite DPs will have to take into account, but that is all they have done. Alas, doing better isn't easy. For one thing, the framework of these Notes would need substantial modification. The DRT-based representations that we have been using do not encode the mentioned aspects of discourse structure, because these aspects were not seen as directly relevant to the problems that DRT was originally designed to deal with. Since the beginnings of DRT extensions have been developed that record some aspects of discourse structure that were missing from DRT in its original form and also from the framework we have been using here, most notably the SDRT of Asher and Lascarides and the Layered DRT of Geurts and Maier (Geurts & Maier 2010). But even if these extensions capture much of the additional structure that is needed, putting all this into a single representational format remains a dating task, which it would be quite impossible to undertake here.

4.2.5 Broadening Perspectives and Coverage

A drawback of the exploration of DP anaphora that this section has engaged in is that it has been confined within boundaries that are not natural ones. There are in particular three types of data that we have not considered and that ideally should have been included: plurals, demonstratives and bridging descriptions. In this last part of Section 4.3.5 I will say a little about each of these, mostly about why they have been excluded in spite of the reasons against this. In the case of bridging more will be said, since this is an aspect of the use of definite descriptions that should be part of an account of definite descriptions in any case, quite apart from the question whether it can throw further light on the forms of anaphora that have been considered in this section up to this point.

4.2.5.1 A. Plurals

The most blatant violation of the principle that the range of phenomena targeted in an investigation should have natural boundaries has been the omission from this section of plural anaphoric DPs. In fact, this is an objection that can be raised against Section 4.3 in general. Some of the benefits of a joint treatment of singular and plural noun phrases were pointed out in Section 2 of PART I of the Notes; and there have been occasional references to those benefits in PART II as well, most recently in connection with the lexical entry (4.94) for *the* in Section 4.3.4.

Had plural anaphoric descriptions and pronouns been included among the expressions investigated in the present section, the general principles that govern the interpretation of anaphoric pronouns and descriptions would have looked rather different from the way they emerged from our look at just the singular DPs. I am not suggesting that the description that has been given of the interpretation of singular anaphoric definite descriptions and pronouns is wrong. But I am fairly convinced that including their plural counterparts would have led to more comprehensive generalizations, more elegant formulations and perhaps also to some deeper insights into the nature of DP anaphora. (For one thing, unique satisfaction, which has played such a central role in the argumentations of this section, would be replaced by the combination of a couple of distinct notions: (a) the formation, triggered by a 'definiteness feature' of the maximal set of satisfiers (within the chosen search domain) and (b) the distinction associated with the features 'singular' and 'plural' between singleton and non-singleton sets (see e.g. (Kamp & Reyle 1993), Ch. 4, (Brasoveanu 2007).)

One reason why the exclusion of plural anaphoric DPs from the investigation in this section may have struck the reader as incongruous is that we have been making a liberal use of plural discourse referents – drefs representing non-atomic individuals in the mereological ontology we adopted – as constituents of discourse context representations. That decision was more or less forced upon us, since without plural drefs we wouldn't have been able to say anything of substance about the workings of Strategy 2 and the ways in which it differs from Strategy 1. But admitting plural drefs as elements of semantic representations is one thing; spelling out the processing rules for the plural DPs that introduce such drefs is a quite different commitment; and charting the processing options for anaphoric plural DPs is yet another. And a proper treatment of plural DPs carries with it a range of further problems, for instance those that have to do with collective interpretations of plural uses of verbs and distributivity. These problems are notoriously hard and that is why the decision was made early on the conception of these Notes to keep plurals out of PART II. Our use of plural drefs in this section contravenes that general decision; but as I said without them it would have been impossible to say much of substance even about the anaphoric behavior of singular DPs.

In the longer run this can of course be no excuse for shunning the treatment of plural DPs, and of plurality more generally. In fact, an integrated treatment of singular and plural noun phrases (and other parts of speech) is one of the most obvious and most urgently needed extensions of the syntax-semantics interface of English that we are developing in PART II of the Notes.

4.2.5.2 B. Demonstratives

Another problematic omission from this section are demonstrative DPs. Pronouns, definite descriptions and demonstratives are the three DP types that have anaphoric uses and that in itself is a reason for studying them together. I do not think that the inclusion of demonstrative DPs in the explorations of this section would have led to significant differences in the formulation of the processing principles for the two types of anaphoric DPs we have been looking at, or even in the organization of their presentation. But since pronouns, definite descriptions and demonstratives are the three DP types that can be used anaphorically, the question how they cover the needs for nominal anaphora between them is of interest in its own right. Arguably the principal value of such 'division-of-linguistic-labor' studies is cross-linguistic: Study the way in which the possible linguistic devices that can be used for a certain type of job in language L1 divide that job between them, study the way this is done by the available devices in language L2 and compare the ways in which the two languages cut this pie. How much replication is there in where the cuts ware made? How much evidence can we find that there are natural places to make them? Although these Notes are devoted solely to English and cross-linguistic comparisons are not within its mandate, a comprehensive study of all the devices that English has for nominal anaphora would be a valuable first step towards such comparative studies of anaphora, in which English is one of the languages compared.

This would have been a strong reason for including demonstrative DPs among the ones considered in this section. Why they haven't been included even so was already touched upon early on in this section (see footnote 33): The primary uses of demonstrative DPs are deictic; without a proper account of those uses an account of their other uses would be either uninformative or skewed. But deictic reference is a topic for which our framework is unsuited. A proper treatment of the deictic uses of demonstratives would require a wholesale overhaul of that framework, of a magnitude that would be quite inconceivable at this point. (The DPs that have been considered in this section, pronouns and definite descriptions, have deictic uses as well, but for them the deictic uses do not seem fundamental to our understanding of their non-deictic uses in the way they seem to be in the case of demonstratives.

So for them focusing on the non-deicitc uses while ignoring the deictic ones isn't as much of a problem.)

4.2.5.3 Bridging

The exclusions described under A and B were exclusions of types of expressions. The one to be addressed in this last part of Section 4.3.5 is the exclusion of a certain way of using definite descriptions, the central topic of this and the preceding section. In the literature on definite descriptions this use is known as 'bridging'.

A typical example of bridging, to which we will return at length below, is found in (4.122).

(4.122)I got a book from Amazon yesterday but I am going to return it, because the front cover is soiled.

Intuitively the definite description *the front cover* in the *because*-clause of (4.122) is interpreted via the DP *a book* in the first clause, viz. as the front cover of the book that this DP speaks of. It is tempting to think of *a book* as the 'anaphoric antecedent' of *the front cover* in a sense comparable to that in which, for instance, *a donkey* in (4.107.d) is understood as the anaphoric antecedent of the definite description *the donkey*. In this sense bridging appears to be similar to the cases of anaphora we have been looking at in this section (as well as at a number of earlier points in these Notes). The main purpose of the following remarks is to become clearer about how anaphora and bridging are related: Is bridging just another form of anaphora? Or is anaphora a special kind of bridging? Or are bridging and anaphora phenomena that resemble each other only in that both involve the identification of 'antecedents'?

Since bridging is one of the ways in which definite descriptions, the DP type to which this and the last section have been devoted, can be used, a serious discussion of bridging has to be on our list, irrespective of whether or how much it affects the analysis of DP uses that have been considered in these sections. Unfortunately I do not feel ready for a serious analysis of bridging, for one thing because I lack the necessary knowledge of the bridging literature. So what follows in these last pages of Section 4.3.5 is to be seen as a kind of promissory note for something better, which I hope to be able to deliver before long.

Exactly what bridging is, or what is the best way to describe it, has been a source of protracted controversies and I am under no illusion that the characterization I will give of bridging here will do much to resolve the controversies. But the following description in (4.123) should do as a starting point for the remarks that will follow:

(4.123)A definite description δ is used as a *bridging description* when its interpretation takes the following form:

a. Identification of a 2-place relation R, by a procedure that is still to be specified.

b. Identification of an 'antecedent' α in the discourse context. (For us this comes to choosing a singular dref from Acc(K_C).)

c. The denotation of δ is then the unique entity that stands in the relation R to the denotation of α. (That there is a unique such entity is part of the presuppositions of bridging descriptions.)

We will have more to say about the choice of R as we go along. For now there is just one question about the possible choices for R that I want to mention: Is the Identity relation = a possible choice for R? This question is special importance in connection with the relationship between bridging and the anaphoric phenomena we have been looking at up to this point. Because we will come back to this question later, it will be useful to give labels to the two possible answers to this question:

A1: The Identity relation = is among the possible choices for R.

A2: = is not among the possible choices for R.

The forms of DP anaphora we have been discussing so far have all been cases where the anaphoric DP refers to the same individual – or, more generally, has the same semantic value or values – as its anaphoric antecedent: The anaphoric DP 'inherits' its referent (or its semantic value or values) from its antecedent, you might say. At least for the core cases of bridging this is not so: the bridging description δ and the DP α that is identified as its antecedent refer to distinct entities. To avoid terminological confusion I will in what follows use the term 'coreference anaphora' for what we have been referring to thus far as 'anaphora'. Unless indicated otherwise, we will be using the term 'bridging' for cases of non-coreference anaphora. And the unmodified term 'anaphora' will be used in the way in which it is used in some of the bridging literature: as a term that encompasses both coreference anaphora and bridging in the sense of non-coreference anaphora. (Common to all cases of anaphora in this broad sense of the term – we already said as much above – is that they all involve the identification of antecedents. In our terms: They all involve the choice of an antecedent in the form of a dref from Acc(K_C).)

Are coreference anaphora and bridging distinct phenomena or are they different manifestations of what in its essence is one and the same phenomenon? This question is directly connected with the alternative answers A1 and A2 to our question above. Coreference anaphora will be the result when (4.123) is applied to a bridging description δ and $=$ is chosen as the relation R. So if the correct answer is A1, then we can think of coreference anaphora as the special case where the choice for R is $=$. Or, more carefully formulated, if A1 is true, then one form that coreference anaphora can take is that of bridging with $=$ as choice for R. (But of course this does not exclude the possibility that the effects of coreference anaphora may also be achieved via a different procedure. We will return to this point.) If on the other hand the correct answer to our question is A2, then no conclusion about the relation between bridging and coreference anaphora follows. For it is still possible that other choices than $=$ for the bridging relation R still allow for 'local reflexivity', i.e. it might be that xRx, where x is the denotation of α.

Looking more closely at out first example

The time has come to look at some bridging examples in detail. We start with the example already given:

(4.122) I got a book from Amazon yesterday but I am going to return it, because the front cover is soiled.

As noted, the phrase *the front cover* in the *because*-clause of this sentence is naturally understood as referring to the front cover of the book mentioned in the first conjunct. How do recipients of the sentence arrive at this interpretation? According to (4.123) the bridging interpretation of *the front cover* involves the identification of (i) a bridging relation R and (ii) an antecedent α. α in this case is the DP *a book*. And R is the relation that holds between books and their front covers – the relation such that if x is a book and xRy[46], then y is the front cover of x. Crucial to the choice of R is evidently

[46]I am assuming that the antecedent provides the first term of R and the bridging description the second term. Thus if R is the relation 'front cover of', then 'xRy' is to be

the head noun *front cover* of the bridging description. In the most typical cases of bridging this head noun is a relational noun, or it is one that can be easily interpreted as one. *cover* (and with it *front cover*) is a good example of this. The relation R triggered by the head noun constrains the choice of the antecedent α: α's denotation must be a member of the Domain of R. In the present case: The referent (or semantic value or values) must be the sort of thing for which the predicate 'has a front cover' is well-defined.

In order that the interpretation of *the front cover* as the front cover of the semantic value or values of the chosen antecedent *a book* be acceptable the relation 'is front cover of' has to be what I will call *locally functional*: for the value x of the antecedent (or for each of its values, if the evaluation of the sentence involves more than one value) there must be exactly one y such that the relation holds between x and y. For the given relation, viz. the relation 'front cover of' local functionality is plausible enough. For first, pretty much every book has a front cover and, second, the thought that a book would have more than one front cover, if perhaps not completely impossible, is certainly far-fetched. So even if it isn't an unquestionable truth that every book has one front cover, the assumption that the book or books spoken of in (4.122) does or do have a unique front cover comes at very little cost; it is an accommodation that an interpreter of (4.122) will normally be happy to make – an act of charity towards the speaker that comes as more or less free of cost.

Some more examples

Two issues have emerged from this discussion of (4.122) that are critical for the interpretation of bridging descriptions in general: (i) How does the content of the bridging description determine or guide the choice of R? (ii) How is local functionality determined for the chosen R? The next example helps to see a little better what may be involved in each of these questions.

(4.124)It was pleasantly cool in the room. Someone had left the window open.

The bridging description in this example is *the window*. Intuitively it seems clear that this phrase must be understood as referring to a window of the room mentioned in the first sentence. But how do we get to this interpretation of (4.124)? First, there is the choice of the relation R on the basis of the bridging description's head noun *window*. *window* isn't in any strict sense a relational noun. True, a window is typically a window in some wall, for the

read 'y is front cover of x' (or as 'x has y for its front cover').

purpose of light coking through the wall or for being able to see the other side, then the window also counts as a window of that room. Furthermore, modern windows can be opened in order to let the air pass between inside and outside. But not every window is the window of some wall. Prefabricated windows often are not the windows of any wall or of any room, so long as they haven't been been assigned to some particular wall or mounted in it.[47] Even so, interpreting *window* as a relational noun, and more specifically, as standing for a relation between windows and rooms, tends to be unproblematic when such an interpretation is wanted on independent grounds.

A difference between this example and the last one is that in the present case local functionality is not quite as straightforward. It is part of standard world knowledge that some rooms have one window, some rooms have more than one window and there are also rooms that haven't got any windows at all. (A sadly common feature of many teaching rooms in contemporary institutions of secondary and higher education.) So minimally the interpreter of (4.124) must assume that the room referred to in the first sentence is one with at least one window. When the interpreter didn't have this information already, then he will have to accommodate it. But mostly this accommodation will be unproblematic too

This is the one half of what has to be assumed towards local functionality of the room-window relation. The other half is that if x is the denotation of *the room* in (4.124), then there is only one y to which x stands in this relation. A simple way in which this second requirement can be satisfied is when the room has just one window. (In that case (4.124) is true just in case that one window is open.) But this doesn't seem to be the only scenario in which (4.124) can be used felicitously. The sentence also seems acceptable in situations where the room has more than one window, but where the window in question – the one that someone had left open – is stood out in some significant way. Perhaps it was the only window in the room that could be opened at all, or the window that people would open if they wanted to open any window at all. Perhaps it is even enough if this window was the one that as a matter of actual fact had been opened by someone on the given occasion, with nothing else to set it part form the other windows in the room.

As far as I can tell, the interpreter's charity can go very far in these cases, almost as if it is enough if the speaker of (4.124) had one window in mind, wanting to say of that window that someone had left it open. Perhaps there

[47]A similar point also applies to book covers. But there is seems a little more contrived.

are some limits to how far charity can reach. For instance, if the room has two identical windows and both have been left open, plain accommodation of local functionality may not be licensed. I am not altogether sure of even this, however. For all I know it might still be possible to entertain the assumption that the speaker had one of those two windows in mind when she uttered (4.124) and accept the utterance on that account, accommodating local uniqueness almost against one's better conscience. In fact, I doubt that there are sharply defined limits to be found here and I will make no further effort to discover if there might be.

Many of these same aspects of bridging interpretations are also illustrated by one of the classical examples from the bridging literature is (4.125).

(4.125)Joyce entered the room. The chandelier was shining brightly.

Here the bridging description is *the chandelier*. It is naturally understood as referring to a chandelier in the room of the first sentence. Again the first step in the bridging interpretation of *the chandelier* is to choose a relation R on the basis of the head noun *chandelier*. We might say that once again this requires interpreting *chandelier* as a relational noun – as meaning something like 'chandelier of x', where x can be a room. But a difference between this case and the last one is that it is not a typical property of rooms that they have chandeliers. Rooms *can* have chandeliers, and some of them do. But for a room to have a chandelier, rather than some other kind of lighting fixture, is more like an exception than the rule.

That, however, is apparently no real obstacle to interpreting *the chandelier* in (4.125) the way people do. And the reason why this way of interpreting the description comes so easily to us probably has to do with the fact that it *is* a common feature of rooms that they have some kind of lighting fixture (or even more than one), even if the fixture or fixtures are not often chandeliers. Thus in this case the existence part of local functionality can be best understood as facilitated by the general knowledge that rooms for the most part have lighting fixtures, together with the knowledge that chandeliers are a kind of lighting fixture. That in the case at hand the lighting fixture at issue is a chandelier can then be seen as some 'non-at-issue' part of the information that is contributed by the second sentence of (4.125).

As in our previous example there is also the uniqueness part of local functionality. Superficially that part may look a little easier in this case: even when the room in question had several lighting fixtures, it may well have been the case that only one of those was a chandelier. When we probe a little further, however, we can see that this doesn't constitute any real advantage vis-à-vis the window example. There are also rooms with several chandeliers, and when the room referred to in (4.125) is one of those, then the acceptability of the statement depends on the same kinds of factors that the acceptability of (4.124) depends on when it is about a room with several windows.

Bridging and Coreference Anaphora

So much for now about the details of applying (4.123) to individual examples. What we have been able to learn from these examples will help us to see a little more clearly how bridging, as characterized in (4.123), is related to coreference anaphora involving Strategy 1 or Strategy 2. The first observation in this connection is a correction of something said above (or at least implied by something said above). In discussing the cases of rooms with more than one window, or with more than one chandelier, we spoke of accommodations towards the existence and the uniqueness parts of the local functionality of the relation R. What was said about existence accommodations is unproblematic. For instance, in the case of (4.124) the accommodation that the room in question has at least one window, or at least one window with some additional properties, can be understood as an assumption about the room-window relation that has been chosen via the relation interpretation of the noun *window*. But in connection with the uniqueness accommodations there is a problem. Consider again the case of a room with several windows. Even if the interpreter knows this to be the case, he may arrive at an acceptable interpretation of (4.124), it was said, so long as he can accommodate that among those windows there was one that stood out in some way and that it was to that window that the description *the window* is referring. But what does such an accommodation say about the chosen relation R?

One answer that might cross the mind is that although the room-window relation that has been chosen as bridging relation R is not locally functional in this case – in that for the given room x there is more than one y to which x stood in this relation, the accommodated information selects a unique window from among those windows. But note that this is actually not consistent with the way in which local functionality of R has been stated. Local functionality of R requires that for the given room x there is one and only one y such that xRy. By assumption that is not the case for this R and it is not something that accommodation can alter. The only consistent way of explicating what accommodation can accomplish in this kind of situation is to provide some information that selects a unique satisfier of that information from the non-singleton set of windows of room x. But note well that the descriptive content 'window' of the bridging description itself can make no contribution to this uniquely selecting information. For by assumption this content is satisfied by all members of the non-singleton set: all members of the set are windows.

An alternative explication of what 'accommodations' towards the local uniqueness of R can do is that they *signal revision*. If the interpreter knows or suspects that the room has more than one window, that is a reason for choosing as R not the simple room-window relation spoken of so far, but a sub-relation of this relation, which only holds between a room and a window of it which has some additional features.

On this second explication what has just been described as accommodation isn't accommodation in the sense in which the term is usually understood. It is more like coercion, such as when a stative VP like *be obnoxious* is reinterpreted as an activity VP when *be* is used in the progressive; see Section 3.5.2. Certain conflicts between constraints on interpretation are resolved by revising some part of the interpretation that had already been adopted.

More importantly for our main concern, there is, according to this second explication, no place anywhere in the procedure outlined in (4.123) for the choice of a non-singleton set from which some descriptive content can select a unique satisfier. So on the present view there is no way of recasting Strategy 2 as a form of bridging. The only successful choices of R are those for which the set $\{y\colon xRy\}$ is a singleton and there are no other sets that the bridging interpretation procedure allows to enter the scene.

On the first explication the situation is different. Here it seems at least formally possible to see the choice of a non-singleton search domain as the special case of choosing a bridging relation R that fails local uniqueness. The choice of a non-singleton domain that is part of Strategy 2 would then have to understood as the choice of a 'degenerate' relation R – a relation that is degenerate in the sense that its first argument has no real role to play. (Formally: a relation R that is degenerate in the sense that there is a fixed set Y such that for any argument x the set of y such that xRy is Y; so R is just the set Y dressed up as a relation.) That such degenerate relations should be among the possible choices for R may seem rather suspect as it is. But there is also another point that makes this reduction of Strategy 2 to bridging implausible. In applications of Strategy 2 the content used to select a unique satisfier from the search domain is given (in whole or part) by the anaphoric description itself. (The descriptive content of the description may have to be reinforced by additional context-supplied predicates, but even then it can be thought of as the core of the predicate that is doing the selecting.) If on the other hand Strategy 2 is defined as a special case of bridging along the lines just suggested, then the descriptive content can play no part in selecting a unique member from the set $\{y\colon xRy\}$. When R is a bridging relation (degenerate or not), then all y such that xRy satisfy the descriptive content of the bridging relation. For that is how the bridging is chosen in the first place.

The upshot of these considerations is that reducing Strategy 2 to bridging is a formal possibility in the best of cases. But even when possible at all, the reduction is a purely formal one, which has little to say for itself otherwise.

What about a reduction to bridging of coherence anaphora via Strategy 1? As far as I can see, the only form such a reduction could take is the following: An application of bridging is case of applying Strategy 1 iff the bridging relation is *locally reflexive*, in the sense that if x is the denotation of the chosen antecedent and xRy, then $y = x$.If this is the case and r is also locally functional, then such applications do assign to the anaphoric description the denotation of its antecedent. But are locally reflexive relations possible choices for R? In connection with thiis question we start by returning to the question we raised early on in tho discussion of bridging, with its two possible answers A1 and A2. If A1 is the correct answer, then the Identity relation R is among the possible choices for R; in that case Strategy 1 coreference anaphora can be seen as a form of bridging, that where $=$ is the bridging relation chosen. If the correct answer is A2, then nothing yet follows, for there might be other choices for R that are locally reflexive. To see a little more clearly what this possibility would come to, let us have another look at an example already discussed, viz. (4.116.e).

(4.116.e) One of the ten balls is not in the bag. The ball is under the sofa.

So far our story about the interpretation of *the ball* in (4.116.e) was that the dref introduced by *one of the ten balls* is chosen as antecedent, and the descriptive content 'ball' of *the ball* is then used to confirm that this is a legitimate choice. (The verification of this is trivial in that any possible referent or semantic value of the phrase *one of the ten balls* must of course be a ball and thus satisfy the predicate 'ball'.)

When Strategy 1 anaphora can be reduced to bridging in the way indicated, and *the ball* in (4.116.e) is interpreted as a case of bridging in this sense, then it must be possible to choose a bridging relation in this case that yields the intended coreference interpretation, which makes *the ball* coreferential with *one of the ten balls*. Can there be such a bridging relation? I doubt it, but have no conclusive argument to offer. As we already saw when discussing the reducibility of Strategy 2, one of the properties that the chosen relation must have is local reflexivity; for otherwise the resulting interpretation will not be coreferential. So one way to show that reduction is not possible would be to argue that bridging relations never are locally reflexive. But that, it seems would be trying to show too much. For an example consider the utterance (4.126), made by A to B in (4.126) in the following setting: B has been complaining to A about everything and everyone always being against him, how unfair the world has always been and so on and has been going on and on about this. At one point A can't stand it any longer and says:

(4.126)Speaking of worst enemies. In your case it is quite clear who the worst enemy is. If you would just listen to yourself for a minute, then you would also see: it is none other than you yourself.

The phrase at issue in (4.126) is the description *the worst enemy*. I take it this description is naturally interpreted as a bridging description. Its head noun is relational and points to the 'enemy'-relation – the relation that holds between x and y iff y is an enemy of x. And the antecedent of the bridging interpretation of this description is B, the denotation of the pronoun *your* in (4.126).

This is a case in which the bridging relation 'worst enemy of' is locally reflexive, since it holds between B and B. So on the assumption that this is a good example of a bridging description, it shows that some bridging relations can be locally reflexive relations. People can be their own worst enemies (and not infrequently they are and when A is right, then B is one of them).

But is this the right kind of bridging relation for a reduction to bridging of Strategy 1? I do not think so. One feature of the 'worst enemy'-relation is that it is *potentially irreflexive*: the relation could have been irreflexive; or put more formally: there are possible worlds w in which the extension of the relation is irreflexive in the sense that for no x in the domain of the extension it is the case that x stands in the relation to x. But can such relations, which are potentially irreflexive in this sense and for which local reflexivity can only be contingent, be the right relations for the reduction of Strategy 1. Once more, I don't think so. In applications of Strategy 1 the coreference relation between the anaphoric DP and its anaphoric antecedent is *not* contingent:

identity between the referent or semantic value of the anaphoric description and its anaphoric antecedent is built into the Strategy 1 interpretation procedure, in our implementation through the addition of the Condition $y = x$ which stipulates y to be equal to x. A bridging relation that is locally reflexive only as a matter of contingently would require an extra interpretation step, which would consist in verifying that the relation is locally reflexive and make use of information that would play no part in an application os Strategy 1.

So bridging relations that are only contingently locally reflexive won't do. What we want are bridging relations that are locally reflexive by necessity, and for which local reflexivity can always be inferred from the way they are specified. the Identity relation = is such a relation and so would restricted Identity relations, such as the identity relation which would only be defined for balls or the identity relation that is only defined for the ten balls mentioned in the foist sentence of relation (4.116.e). There will be other such relations as well, although I find it hard to think of plausible examples.

How much plausibility is there to the assumption that there are necessarily locally reflexive relations among the possible bridging relations. Not much, I cannot help feeling. There may be no way to settle this question other than be stipulation, excluding such relations from the repertoire of bridging relations that are available in applications of bridging in the sense of (4.123). But here is a further consideration that weighs against seeing such relations as possible bridging relations. All the bridging relations that have been encountered in the bridging examples we have considered can be thought of as 'possession' relations – mostly in a very abstract and attenuated sense of possession for sure, but still possession relations in that they can be paraphrased using the verb *have*. A book can be said to 'have' a front and a back cover, and a room can be said to 'have' one or more windows, or 'have' a chandelier. But a donkey (or anything else for that matter) cannot be said to 'have' itself in this sense of *have*; Identity, absolute or restricted, is not a possessive relation in this sense. So if bridging relations must always be possessive relations, then Identity relations cannot be bridging relations. And the same would intuitively seem to apply to any relations that are necessarily locally reflexive. If that is true and if all bridging relations are possession relations in the weak sense hinted at, then Strategy 1 coreference anaphora cannot be reduced to bridging.

To conclude this discussion of the relation between bridging and coherence anaphora: Al though we haven't found any conclusive evidence that nei-

ther Strategy 1 nor Strategy 2 coreference anaphora can be seen as forms of bridging, the prospects for conceptually plausible reductions appear to be very slim. Of course this does not mean that bridging never leads to coreference. (4.126) is a case in point. But intuitively such cases seem quite different from what is normally understood by anaphora and that we have been referring to as coreference anaphora in this section for the sake of terminological perspicuity.
I conclude this discussion of bridging with an observation that I have not found elsewhere (which may not mean much given my limited knowledge of the bridging literature), which do not have an explanation for, but that I find intriguing enough to mention nonetheless. The observation concerns a restriction on the use of bridging descriptions in English which is not found in certain other languages. (It doesn't hold for German, at least not in the same form in which it holds for English.)

The discussion of bridging interpretations as applications of the procedure in (4.123) would seem to imply that if the relation R determined by the descriptive content of the bridging description is a function, then the application would be a particularly straightforward one – none of the special assumptions that we found often have to be made to secure local functionality would be needed – and the sentence containing the bridging description should be expected to be maximally felicitous. But curiously the sentences in (4.127) suggest that that is not so.

(4.127)a. The accident left her crippled for the rest of her life. She lost the right foot.

b. The accident left her crippled for the rest of her life. She lost her right foot.

c. Joan had grown up like an orphan. The father had never bothered to find out how she was doing.

d. Joan had grown up like an orphan. Her father had never bothered to find out how she was doing.

e. Joan had grown up like an orphan. Joan's father had never bothered to find out how she was doing.

f. Joan had never met the father.

g. Joan had never met her father.

(4.127.a) is awkward and carries a suggestion that the speaker is non-native. Perhaps there are special contexts in which the use of the definite article in this sentence is acceptable, but the natural way to express what (4.127.a) is trying to say is the sentence in (4.127.b), with the possessive *her* in lieu of the definite article *the*. A similar contrast can be observed between (4.127.c) and (4.127.d): (4.127.c) is awkward if *the father* is to be interpreted as Joan's father; the better way to say this is (4.127.d), with the possessive *her*, or (4.127.e) with the saxon genitive *Joan's*. Finally, the contrast between (4.127.f) and (4.127.g) is even more telling. (4.127.f) just cannot be interpreted in the way that we can interpret (4.127.g) by taking *her* as anaphoric to Joan. In (4.127.f) *Joan* cannot be the bridging antecedent of *the father*.

These examples seem to indicate that bridging interpretations are blocked in cases where the head noun of the bridging description seems to present us with a relation R that is obviously a function (and thus one for which local functionality comes for free). I confess that I have no clue what it is that makes such putative bridging examples bad. As noted, this particular constraint on bridging does not hold for German. (In a way that makes the constraint only more puzzling, since in all other aspects of anaphora and bridging discussed in this and the last section German definite description seem to behave like English ones.) The German translations of the definite descriptions in (4.127.a,c,f) all allow for bridging interpretations in which the subject DP is the antecedent. The most striking difference can be observed in connection with (4.127.f). The literal German transition of (4.127.f), 'Joan hatte den Vater nie getroffen', can be interpreted as meaning that Joan had never met her own father, although the interpretation according to which she met somebody else's father is possible as well and in out-of-context presentations may be somewhat preferred.

I leave the facts illustrated by (4.127) for what they are. (Any hints of what is behind this constraint will be grateful received.)

Whatever the reason may be for the constraint revealed by these examples, it is a constraint that would appear to have one important consequence for the methodological question that has dominated much of the discussion of the last few pages: How (if at all) are bridging and coreference anaphora formally or conceptually related? One of the issues that came up in our discussion of this matter was whether Identity relations can serve as bridging relations. In the light of what is revealed by the sentences in (4.127) this now seems even less likely than it already did. For these relations are obviously functional. If functional relations are generally prohibited from acting as bridging relations then that would rule out identity relations as well.

Here ends what little these Notes have to offer on the topic of bridging. As I said, these remarks are to be understood as a kind of promissory note for something more substantive. But I hope that the above remarks have thrown some light at least on the possible relations between bridging, in the form we have assumed here, and coreference anaphora as discussed earlier in this section.

This is also the end of Section 4.3.5.[48]

---

[48](The following footnote has been in existence from the time the first lines of Section 4.3.5 were written. As the section gradually took shape, the footnote has been in search of its right attachment point. At long last I decided to attach it right here at the very end of the section. In its present position the footnote can be read in either or both ways: as a comment on an alternative approach to the semantics and pragmatics of pronouns and definite descriptions or as a qualification of what I said at the outset of the section, when I made the observation that according to the view which has dominated much of the history of theorizing about the logic and semantics of reference and anaphora pronouns and descriptions are birds of very different plumage. When read in this second way, what the footnote has to say could have come much, much earlier in the section – attached, say, to the early paragraph in which the traditional view about pronouns as variables and definite descriptions as expressions that determine their referents through unique satisfaction was first dscribed. But the footnote can also be seen as a natural coda to the section as a whole, which draws its interest from much that has been said in the course of the section, including the discussion of bridging and coreference in the last fourteen pages.)

The opening remarks of Section 4.3.5 about the very different conceptions of pronouns and definite descriptions that can be found in the philosophical literature on reference are in need of at least one substantive qualification. This qualification is related directly to the 'donkey sentence phenomena' that were first mentioned in these Notes as one of the central motivations for DRT (see Section 2). There is a way of thinking about 'donkey pronouns' and their anaphoric antecedents that is very different from the conception that led to DRT as its formal implementation. We can recognize the point of departure for this alternative view in some of the things that were said about donkey pronouns by Geach (Geach 1962), who brought the donkey pronoun phenomena within the spotlights of contemporary logic and linguistics. Geach introduced the term 'pronoun of laziness', as a way of conveying the view that donkey pronouns are used as abbreviations of more complex erxpressions, which can be reconstructed in some way or other from the contexts in which those 'abbreviations' occur. This suggestion has been taken up by several philosophers and linguists, most notably by Evans (Evans 1980), (Evans 1977), Cooper (Cooper 1979), Neale (Neale 1990), Heim (Heim 1990) and Elbourne (Elbourne 2005). (Some of these approaches are known as 'E-type' and others as 'D-type'. There are important differences between these, but for the concerns of this footnote those do not matter.) Common to all these proposals is the idea that a proper treatment of donkey pronouns can be obtained by finding definite descriptions that correctly capture their contribution. Such approaches run into two problems. The first is to find a principled way of translating pronouns into definite descriptions. Generally valid translation recipes based solely on the linguistic form of the contexts in which the pronouns appear (consisting of the sentences containing the pronouns as constituents, and perhaps additional material when the pronoun's 'antecedent' occurs in another sentence) are hard to come by, something of which Evans, the first proponent of a formal version of this general approach, was already clearly aware. But the more serious problem is this. The approach presupposes that a good account of definite descriptions is available already and can be taken off the shelf. Given such an account, it would be possible to transform sentences with donkey pronouns into sentences in which the pronouns have been replaced by the corresponding definite descriptions and then apply the account of definite descriptions to those latter sentences.

But what could that account of definite descriptions be? Mostly one isn't told, as if the matter were too obvious to deserve an explicit statement. At the time when the first proposals along these lines were stated the only account that gave something like a formal way of stating what the truth conditions of sentences with definite descriptions are was Russell's Theory of Descriptions. We discussed some of the problems with this theory in Section 4.3.4. Now it is in principle possible that these problems would not arise for the special kinds of definite descriptions into which pronouns are being translated according to the translation algorithms that those approaches also do or should provide. But most of the existing proposals simply aren't explicit enough to enable us to determine whether that might be so. This I believe to be true also of the most recent proposal along these lines, which can be found in the work of Elbourne ((Elbourne 2005) and subsequent work). Elbourne's model-theoretic semantics is a version of situation semantics and his account of the truth-conditional contributions made by definite descriptions is also stated in his situation-theoretic terms. This means in particular that the unique satisfaction conditions that are part of Elbourne's account of definite descriptions have to be evaluated in situation-theoretic terms; that is, unique satisfiers of descriptive contents have to be found within certain situations; so in this regard situations function in his theory in a similar way as the restricted search domains that are part of the account presented in this section. The difficulty I have with Elbourne's account is the seemingly unconstrained use it makes of quantification over situations. This makes it rather hard – for me, at least – to reconstruct the details of the syntax-semantics interface that must be part of the general account from the sample analyses Elbourne provides – both of the formal definition of that interface and of the motivations behind it. Also, the model theory of the situation-theoretic meta-language in which Elbourne's syntax-semantics interface is formulated is difficult to reconstruct (from the publications I have seen).

In the light of what has been argued in this and the last section of the Notes, there is a further question that must be raised for those approaches (except that of Elbourne; see above). We have seen that many descriptions do not select a unique satisfier in any absolute sense. If they select a unique satisfier at all, then only within a restricted (and often heavily restricted) search domain. If the background account of definite descriptions on which an E-type or D-type approach must rely makes no provision for restricted search domains, then the descriptions that do require domain restrictions will have to be treated as 'descriptions of laziness' and a translation algorithm that transforms them into descriptions that are not subject to such restrictions will have to be specified as well. Perhaps the additional burden this places on accounts of the kinds in question isn't all that large: assuming that a suitable translation algorithm can be given for pronouns, extending that algorithm to definite descriptions might well be a lesser headache. Still, it is curious that this point hasn't been made more of than it has been in those E-type or D-type approaches that I am familiar with.

To summarize this footnote, E-type or D-type approaches to pronominal anaphora share with the view that has emerged in this section that both see pronouns and definite de-

This is the end of what these Notes have to say about the presuppositional approach to the reference of definite noun phrases, and therewith the end of Section 4.3. As noted a number of times before, Section 4.3 couldn't be described as a survey of DP reference, as there is too much that is missing. The most important omission, once more, is the absence of the deictic dimension of reference, and that is a defect which cannot be removed without some radical changes in the general approach that is being developed in these Notes. Somewhat paradoxically perhaps, the direct relationship between deictically referring expressions and their referents in the real world in which the languages are used to which these expressions belong makes it necessary to also be much more explicit than we have been in these Notes about the psychological dimension of language use: A viable account of deixis is possible only within a framework that also has a good deal to say about the relationship between natural language expressions – expressions belonging to languages that have an existence and identity that are largely independent of their individual speakers – and the mental representations that lead to and result from their production and interpretation. For the beginnings of such a more radical departure from the traditions, rooted in Montague Grammar, that are upheld as standards by a considerable part of the formal semantics community today see (Kamp 2016) as well as the forthcoming (Kamp 2019).

scriptions as reference devices of essentially the same kind. But the difference between the two (E-type or D-type approaches on the one side and the account developed here on the other) is nonetheless a profound one. The E-type or D-type approaches (arguably with the exception of Elbourne) try to reduce pronouns to descriptions while relying on the problematic assumption that a viable account of definite descriptions is already in place. What that theory of descriptions is supposed to be is often left for the reader to guess. At least for the earliest proposals, such as Evans', the presupposed account of definite descriptions is presumably Russell's Theory of Descriptions. But when this theory is used as part of an account of the behavior of pronouns in natural languages, the problems with it cannot be simply ignored.

The approach pursued in Section 4.3 sees reference and presupposition as constituting a general problem for all definite DPs. The problem unpacks differently for the different types of definite DPs, and also for different types of use of the same DP-types. And definite descriptions, because of the remarkably wide range of different ways in which they can be used, present the hardest unpacking problem of all. Their analysis should come ay the end of an analysis of how definite DP reference works, not at the beginning.

4.3 Tense and Aspect II: Tense and Aspect with Presuppositions

We ended our discussion of Tense and Aspect in Section 3 with the sobering conclusion that in spite of the many pages we had devoted to this topic none of the phenomena we mentioned in the opening pages of PART II had come within range. The reason,we noted, was that all those phenomena have to do with a presuppositional dimension to the semantics of the tenses and to other linguistic devices that contribute temporal and/or aspectual information.

We already noted in Section 3 that what appears to be the oldest explicit sign of awareness of this dimension is what Reichenbach has to say about the Past Perfect. Here is one of the quotations with which the chapter on tense of (Reichenbach 1947) begins:

(4.128)But Philip ceased to think of her a moment after he **had** settled down in his carriage. He thought only of the future. He **had** written to Mrs. Otter, the massière to whom Hayward **had** given him an introduction, and had in his pocket an invitation to tea on the following day.
(From: W. Somerset Maugham, *Of Human Bondage*)

Recall also the pair of examples in (3.1) that we used to illustrate the core of his observation.

(3.1) a John proved the theorem in twenty lines. Mary proved it in ten lines.
b John proved the theorem in twenty lines. Mary had proved it in ten lines.

Reichenbach observed that interpretations of Past Perfects involve a 'Reference Time' that is located in the past of the Speech Time while the described eventuality is in the past of the 'Reference Time'. Or, to put the matter in the more procedural terms of interpretation via discourse representation construction, interpreting a Past Perfect requires identifying a Reference Time tp in the past of the utterance time n and then locating the described eventuality in the past of tp. When a Past Perfect occurs in context, then the Reference Time must be retrieved or reconstructed from the context, and if things work out as they should, a time or event that is salient in the context will offer itself as the natural choice. (3.1.a), we saw, is an example of this; here it is the event described in the first sentence that provides the reference

Time for the second.

Reichenbach's Reference Times have been taken to not only play their part in the semantics of 'two-dimensional' tenses like the Past Perfect, but also in mediating the kind of temporal discourse structure that is typical of narrative prose, where a series of event sentences in the Simple Past is typically understood as describing a corresponding sequence of events using 'temporal iconicity' in the sense that each next sentence in the series describes an event that followed the one described by the preceding sentence. To account for this it has been proposed that each event sentence introduces not only a dref for the event it describes but also a corresponding Reference Time following it, which can then be used as location time for the next event, when the next sentence is processed, which introduces that event ((Partee 1984), (Hinrichs 1986)). It is not completely clear from Reichenbach's own quite brief account of his theory of tense in (Reichenbach 1947) whether he intended or even envisaged this role for his Reference Times. But if Reference Times are to play this part as well, then the notion of a single reference time per tensed clause becomes a problem. Recall in this connection the 'extended flashback' (3.4) that was discussed briefly at the outset of Section 3 and that is repeated here:

(3.4) Bill arrived at noon. He had got up at six thirty, had cooked himself a full breakfast, and had washed up after finishing it. He had left the house in time to catch the 7.54 train at the central station.

Recall our diagnosis of the roles that Reference Times play in this discourse: (a) for each of the Past Perfect clauses, the RT should be noon (the time when Bill arrived); (b) for all but the first Past Perfect clause C the event described by the immediately preceding Past Perfect clause provides an RT that the event described by C is interpreted as following more or less immediately.

For these reasons (Kamp & Rohrer 1983*a*), (Kamp & Reyle 1993) propose to distinguish between two types of times that play a Reference Time-like part in the interpretation of these and other sentences, the *Temporal Perspective Point*, or *TPpt*, and the *Reference Point* or *Rpt*. The TPpt embodies Reichenbach's idea about the semantics of the Past Perfect: it identifies a past vantage point from which the described eventuality is seen as being situated in the past (and thus as even farther into the past than the TPpt itself). The Rpts on the other hand play their part in the interpretation of narrative sentence successions (which can be in the Past Perfect, as in extended flashbacks like that in (3.4), but may also take the form of successions of Simple

Past tense sentences, or even of sequences in which the sentences have yet some other tense, e.g. the Simple Future). We adopt this distinction. For our present considerations only TPpts are directly relevant.

One difference between Reichenbach's analysis of the Past Perfect and ours is that Reichenbach treated the Past Perfect as an unanalyzed form, on a par with the Simple Past, the Simple Present and other non-prefect tense forms. For us the Past Perfect is a combination of the Perfect and the Simple Past and the function of the past TPpt is to provide a location time for the result state that the perfect delivers as output. As far as the location of the described eventuality is concerned the net effect of the two accounts will be the same: the event described by the sentence is located in the past of the past TPpt. But our analysis leads to a further question, which does not arise when the Past Perfect is treated as an unanalyzed tense form. If we want to follow Reichenbach in his insistence on TPpts as an essential ingredient in the analysis of past perfect sentences, then we must allow that at least some occurrences of the Simple Past (those in which it is applied to a perfect) trigger the identification of a past TPpt.

But which occurrences? A priori there are two possible answers to this question. The first is that all occurrences of the Simple Past involve the identification of a past TPpt and the second that some Simple Pasts involve this but others don't. The first answer is the simpler one, so arguably we should adopt it if there are no clear pointers towards an alternative. But in fact there are good reasons for not wanting to adopt the simple answer. To explain these, however, we have to look beyond the Past Perfect and the Simple Past, to the general analysis of tense that we are working towards.

Reichenbach offers a general 'two-dimensional' account of tense, in which every tense form is characterized by a pair of temporal relations, one between Speech Time and Reference Time and one between Reference Time and 'Event Time'. (Thus the Past Perfect is characterized by $<\prec,\prec>$ (R $\prec$ S, E $\prec$ R).) Our aim is to follow him in that we too want an analysis in which each of distinct tenses we recognize determines a pair of relations, between TPpt and the utterance time n, and between the described eventuality ev and the TPpt. The first relation imposes a constraint on where the TPpt is to be located in relation to n, and the second relates ev temporally to the TPpt.

In an account of this kind – whether Reichenbach's or the one we are in the process of developing – there would seem to be only one plausible option for the characterization of the Simple Present, viz. as the form that is to be used to talk about what is the case at Speech Time: it should be characterized by the pair $<\equiv, \equiv>$, where '$\equiv$' stands for temporal coincidence. In other words, the Reference Time (or, for us, the TPpt) coincides with the Speech Time n and the described eventualityy coincides with the Reference Time/TPpt (and thus also coincides with the Speech Time). This indeed is the characterization that Reichenbach offers, and we see no alternative but to adopt it as well.

In our discussion of the present tense in Section 3.7.2 we noted that what we call its 'standard' use of the present tense is felicitous only when its input is a state description (or, in terms less specifically geared towards our DRT-based implementation, when the input has imperfective aspect). The reason for this restriction, I suggested, was that when we speak about what is going on while we are speaking, then by and large only state descriptions (or imperfective aspect) are suited for what we are trying to do; and the standard use of the present tense is precisely to perform speech acts of this kind. This fact about the Present Tense in its standard use deserves to be reconsidered in the light of the two-dimensional characterization to which we want to commit ourselves, according to which the Present Tense is given by the pair $<\equiv, \equiv>$. Which part or parts of this characterization is responsible for this fact?

I doubt that it is possible to give a clinching answer to this question. But if we think of TPpts in the way that the term 'temporal perspective point' suggests, viz. as identifying a perspective or vantage point from which the described information is viewed, then it seems natural to seek the explanation of why the standard use of the Present Tense is restricted in the way it is, with the second member of the pair $<\equiv, \equiv>$, which says that the eventuality temporally coincides with the TP pt and thus that the described eventuality is going on at the time from which it is being viewed.

If this is right, then the first conclusion must be that the Simple Past of Simple Past event sentences, such as 'John arrived' or 'Mary wrote a letter', cannot be analyzed by a relation pair whose second member is $\equiv$; for whatever the first member of the pair may be, the second member indicates that relation between described eventuality and the perspective from which it is seen which is possible only if that eventuality is 'imperfective' (i.e. for us: if it is a state). Rather – this would appear to be the only possible alternative if the Simple Past of Simple Past event sentences is to be characterized by a pair of relations, between TPpt and n and between ev and TPpt – the pair characterizing these uses of the Simple Past ought to be $<\equiv, \prec>$.

This is as much motivation as I am able to give in support of the claim that not all instances of the Simple Past are analyzable as characterized by the pair $<\prec, \equiv>$. But if we accept it, and with it the claim it supports, and if at the same time we want to analyze Past Perfects as Simple Pasts that locate result states at past TPpts, then we are settled with two different characterizations for the Simple Past: $<\equiv, \prec>$ for the Simple Pasts of event sentences and $<\prec, \equiv>$ for Past Perfects. That still leaves with a range of uses of the Simple Past for which the characterization is still open, viz. Simple Pasts of imperfective sentences – that is, in our set up, Simple Pasts of state descriptions. What are we to say about these?

There is one subset of sentences belonging to this remaining category for which there are good reasons to assume that they too must be classified as instantiating the use of the Simple Past that is characterized by $<\prec, \equiv>$. These are sentences in which the past tense tense occurs in conjunction with the word *now* and in which *now* refers to some time in the past of n. Examples are provided by the discourses in (4.129).

(4.129)a. Fred lit a fire and sat down in front of it. It had been a hard week. But now he was able to relax/now he was enjoying himself at last.

b. Fred lit a fire and sat down in front of it. It had been a hard week. But now he had done everything they had asked him to take care of and he could relax.

c. Fred lit a fire and now he sat down in front of it.

In each of the two versions of the third sentence of (4.129.a) *now* refers to the time when Fred sat down in front of his fire, or perhaps just after it, when he was already sitting. The same is true for the occurrences of *now* in the Past Perfect sentence of (4.129.b). Both (4.129.a) and (4.129.b) seem perfectly felicitous (though perhaps they are more natural in writing than when spoken). But note that each of these sentences containing *now* is, by our criteria, the description of a state. In this regard (4.129.a) and (4.129.b) differ from (4.129.c), whose second sentence contains *now* but is clearly an event description. (4.129.c) also appears to be grammatical and the occurrence of *now* it contains is also understood as referring to the past time of the event that its second sentence describes. But there seems to be a distinct flavor to this discourse. It feels as if the second sentence is a kind of historical present, but with the 'present' aspect conveyed by *now* rather than by the use if a present tense. If this is right, then what we see in (4.129.c) involves a shift of the utterance time perspective in a more radical sense than the shifts that are captured in our account by past TPpts. (An additional mechanism would be needed to deal with cases like (4.129.c), hopefully (but not necessarily) the same one that would be needed to give a systematic formal account of historical presents. But this is a matter we set aside here, just as we did in Section 3.7.2.)

The difference we can observe between (4.129.c) on the one hand and (4.129.a) and (4.129.b) on the other is clearly correlated with the fact that the *now*-sentence of (4.129.c) is an event sentence, whereas the sentences in (4.129.a) and (4.129.b) that contain occurrences of *now* are state descriptions. So our conclusion is that past tense sentences in which *now* occurs and refers to a past time and which do not give rise to the special effects just observed in relation to (4.129.c) must be descriptions of states. We can account for this constraint by making the following assumptions:

(i) *now* always serves to locate the eventuality described in its clause; and

(ii) *now* may refer to times other than the utterance time; but when it does refer to some other time than n, then this time must always be the selected TPpt.

These assumptions have two relevant consequences: (a) given our earlier assumption that temporal coincidence with the TPpt requires the described eventuality to be a state, *now* is compatible only with descriptions of states, irrespective of whether it refers to n or to some time other than n; (b) the Simple Past tense of a Simple Past sentence containing *now* must be analyzed as characterized by $<\prec, \equiv>$. It is conclusion (b) that we have been after: Simple Past sentences with *now* are past tense sentences involving state descriptions whose Simple Pasts must be analyzed as characterized by the relation pair $<\equiv, \prec>$.[49]

[49]These considerations involving *now* were originally made in connection with its French counterpart, the word *maintenant*. In French, the formal equivalent of the Simple Past, the Passé Simplé, has been largely supplanted by the Passé Composé, the French form of the Present Perfect. But in combination with the Passé Composé *maintenant* can never refer to the time of the described event. Combinations of *maintenant* with the Passé Simple are marked to begin with because uses of the Passé Simple tend to be marked in any case (at least within contemporary spoken French). But I believe that even in cases where the Passé Simple is acceptable as such, adding *maintenant* is very awkward and perhaps fully unacceptable. [Check this!] If this is true, then French provides an even clearer case of what we have been trying to argue, viz. that except for discourses that produce the special effect observed in connection with (4.129.c) *now* in past tense sentences is compatible with state descriptions, but not with event descriptions: If our judgments about the possible

This settles the analysis of the Simple Past for some state describing sentences that are not perfects, but it still leaves a substantial contingent unaccounted for – Simple Past state describing sentences that are not perfects and that do not contain *now*. What are we to say about this remainder? I am not quite sure. On the one hand our observations about the role of *now* in past tense sentences provides some kind of clue. For many past tense state-describing sentences without *now* it is possible to add *now* without causing ungrammaticality (or the special effect noted in connection with (4.129.c)); the only clear change in meaning is that the sentence gains a contrastive dimension that it doesn't have without *now*. This is no proof that these sentences must be analyzed as involving $<\prec, \equiv>$, but it lends additional plausibility to such an analysis. On the other hand there also instances of which look like past tense state descriptions, but where such an analysis is much less plausible. Consider a discourse like the following.

(4.130)I am sorry not to have been in touch for so long. But we have had a dreadful time. First, our son Billy was ill and I had to be at home most of the time to look after him. Then I was ill myself, with the same nasty virus. Then my in-laws came to stay with us. ...

The last three sentences of (4.130) are all in the Simple Past and by the criteria we have been applying they are all state descriptions. But here adding *now* is not an option. *adding now* is out of the question, but that may be because as a general rule *now* and *then* are in complementary distribution. But in the context provided by the first two sentences of (4.130) *replacing then* by *now* is not an option either. Furthermore, the Simple Past sentences of (4.130) form the kind of temporal progression that is typical for narrative sequences of event sentences. So for these sentences an analysis of the Simple Past in terms of the pair $<\equiv, \prec>$ is not just possible; in the light of what we have been saying such an analysis suggests itself. Let us assume that the adverbs in these sentences – the adverb *first* and the successive occurrences of *then* that follow – should be interpreted as referring to successive times that are included in the 'awful time' referred to in the second sentence. Somewhat more precisely: *first* introduces a time within the temporal interval denoted by *an awful time* and the successive occurrences of *then* further times within that interval, temporally ordered in the same way as the occurrences of these adverbs in the text. (A proper account would have to spell out the contributions of adverbs like *first* and *then*, but to go into the necessary details

uses of *maintenant* are correct, then its incompatibility with event descriptions doesn't allow for any exceptions whatever.

of this would be counterproductive at this point.) Given these references by the temporal adverbs in (4.130) the states described by its successive clauses can then each be temporally located as surrounding the time denoted by the adverb of its clause.

The semantic representation for (4.130) that will result if its adverbs and the states described by its clauses are interpreted as indicated will be independent of the question whether the Simple Pasts of (4.130) are interpreted as instances of the relation pair $<\equiv, \prec>$ or as instances of $<\prec, \equiv>$. In this regard (4.130) is one of many examples: single clauses in the Simple Past or sequences of such clauses whose interpretation doesn't depend on whether their Simple Pasts are taken to be of the $<\equiv, \prec>$- or the $<\prec, \equiv>$-variety. This leaves us in a certain methodological quandary: Which Simple Pasts are to be treated as instances of which variety?

This is the situation we find ourselves in: On the one hand there are cases of the Simple Past sentences and discourses which ought to be treated as instances of $<\equiv, \prec>$. In these examples the clauses describe events. On the other hand there are examples whose Simple Pasts should be treated as instances of $<\prec, \equiv>$ and whose clauses describe states. And then are examples for which the choice between $<\equiv, \prec>$ and $<\prec, \equiv>$ doesn't seem to matter. Question: What policy should we adopt in those cases for which we lack clear evidence that the choice ought to be made this way or that? This is the solution I propose we adopt: Whenever a Simple Past clause is event describing, then its tense is to be treated as an instance of $<\equiv, \prec>$ and when it is state describing, then its tense is to be treated as an instance of $<\prec, \equiv>$. (This policy may eventually have to be adjusted in the light of further empirical observations or theoretical considerations, but that will not happen in those Notes.

4.3.0.4 A revised entry for the Simple Past and its use in the semantic representation of a sentence in the Past Perfect

The discussion above has led us to two different versions of the English Simple Past, one characterized by the combination $<\equiv, \prec>$ and one characterized by the combination $<\prec, \equiv>$. Our final decision was to link these two options to the aspect of the clause: When the input representation to the Simple Past is an event description, then this Simple Past is to be handled as an instance of $<\equiv, \prec>$, if it is an event description, then it is to be handled as an instance of $<\prec, \equiv>$.

For the lexical entry for the Simple Past this link makes things a little easier. As before, the entry must distinguish between the case where the input representation to the Simple Past operator is an event description and that where it is a state description. The difference with our earlier entry in Section 3.3 is that now there is for each of these two alternatives a TPpt presupposition that comes with it. These presuppositions differ for the two alternatives. At this point we are not yet in a position to say anything about the presuppositions that should come with the $<\equiv,\prec>$-version of the Simple Past. So for the time being we leave this part of the entry for the Simple Past tense open. (In the lexical entry below this is indicated by a question mark in angled brackets that acts as a place holder for the relevant information which will be supplied later on.)[50]

On the other hand we are in a position to be quite specific about the TPpt presupposition for the $<\prec,\equiv>$-version and this is what most significantly distinguishes the entry in (4.131) from the old entry in (3.22). Like the old entry the new entry is stated as an entry for the Tense feature past. Like presuppositions generated by other triggers, the resolution of the TPpt-presupposition fort the state description input of the entry is subject to special constraints. Just as for the other presuppositions we have so far considered, we need some way of marking TPpt-presuppositions as TPpt presuppositions (i.e. as presuppositions whose resolutions must obey the constraints on TPpt presupposition resolution). In the present case the condition 'TPpt(t)' can be used as a formal indicator that the presupposition to which it belongs is a TPpt presupposition. (So no subscript on the presupposition representation is needed in this case to record what kind of presupposition one is dealing with.)

[50]When discussing temporal quantification, in particular in Section 3.11.7, we found it necessary to impose perfect temporal coincidence between the quantification states created by quantifying temporal adverbs and the location times introduced by tense (expressed by Conditions of the form '$t = dur(s)$')). A similar need arises for past tense generic, habitual and dispositional sentences, involving the operators GEN, HAB and DIS. The lexical entry for past (and likewise those for fut) should be stated in a form that accounts for these possibilities. This ought to take the form of the entry providing a $\overset{!}{\vee}$ disjunction involving three rather than just two disjuncts, together with a regime that records for each input representation which of the three possibilities – $e \subseteq t$, $t \subseteq s$, $t = dur(s)$ – applies to it. To expand the entry right here also in such a way that it can take care of inputs that are descriptions of quantification states and the like would lead to a notational overload, which would serve no praticle purpose and only make structures harder to read. I have conscious;y desisted tho move.

(4.131) (revised lexical entry for the tense feature 'past')

past (tense feature)

Sel. Restr: eventuality description

Sem.Repr: $<ev_{ref}, ... \mid K> \leadsto$

$$<ev, ... \mid K \bigcup \begin{array}{|c|} \hline \\ <t \mid <\{ \begin{array}{|c|} \hline t? \\ \hline t \prec n \\ [??] \\ \hline \end{array} \}, \begin{array}{|c|} \hline \\ \hline \text{Event}(ev) \\ ev \subseteq t \\ \hline \end{array} >> \\ \\ \downarrow \\ \\ < t \mid <\{ \begin{array}{|c|} \hline t? \\ \hline t \prec n \\ \text{TPpt}(t) \\ \hline \end{array} \}, \begin{array}{|c|} \hline \\ \hline \text{State}(ev) \\ t \subseteq ev \\ \hline \end{array} >> \\ \\ \hline \end{array} >$$

Even in its present incomplete form (4.131) can be used for constructing the semantic representations of sentences in which the feature past receives a state description as input. As an example we go through the representation construction for the second sentence of example (3.1.b), which we repeat here one more time.

(3.1.b) John proved the theorem in twenty lines. Mary had proved it in ten lines.

The LF for the second sentence of (3.1.b) is largely determined by assumptions of which we have seen several applications before. There is just one complication, the role and syntactic status of the PP *in ten lines*. A deeper analysis than we are in a position to provide here would relate this PP to the nominal root $\sqrt{proof}$ from which the verb *prove* is presumably built: 'prove x' means, and has the structure of, 'provide x with a proof' and the PP serves to say more about the proof with which the event of proving the theorem 'supplies' the theorem: it is a proof of/in ten lines. A second question that is raised by the PP *in ten lines* is the syntactic and semantic status of the prepositional complement *ten lines*. A typical feature of PPs that serve to say more about the entities that are contributed by the nominal roots of verbs is that they say more about these entities, and do that in a form in which the prepositional complement is the predicate, whereas the preposition plays a case-like role, indicating how this predicate is related to the entity. (A‘ proof in ten lines’ is the same as a ‘proof of ten lines’ or a ‘ten line proof’.) This suggests that the prepositional complement *ten lines* of *in ten lines* isn’t a full argument phrase, which fills the second argument slot of the preposition *in* in the way that is by now familiar to us, but that it is a constituent of category NP, which ‘incorporates’ into this slot, with the semantic effect that the phrase *in ten lines* as a whole functions as a predicate, with *in* merely serving as an indicator of how this predicate relates to the phrase that it combines with.[51]

An analysis of *prove in ten lines* that does justice to the internal structure of *prove* and the special status of the ‘PP’ *in ten lines* is beyond the scope of these Notes. The best we can do is to treat the ‘PP’ as an ordinary PP which is adjoined to some projection of the verb and to treat its semantic contribution as that of providing an additional predication of the eventuality described by the verb (the referential argument of the representation of the PP’s adjunction site). We assume that the adjunction of the phrase *in ten lines* is low down, at the level of VP.

With this proviso the LF from which we are going to compute the semantics of the second sentence of (3.1.b) is as follows.

[51] For details see for instance (Rossdeutscher 2013).

(4.132)

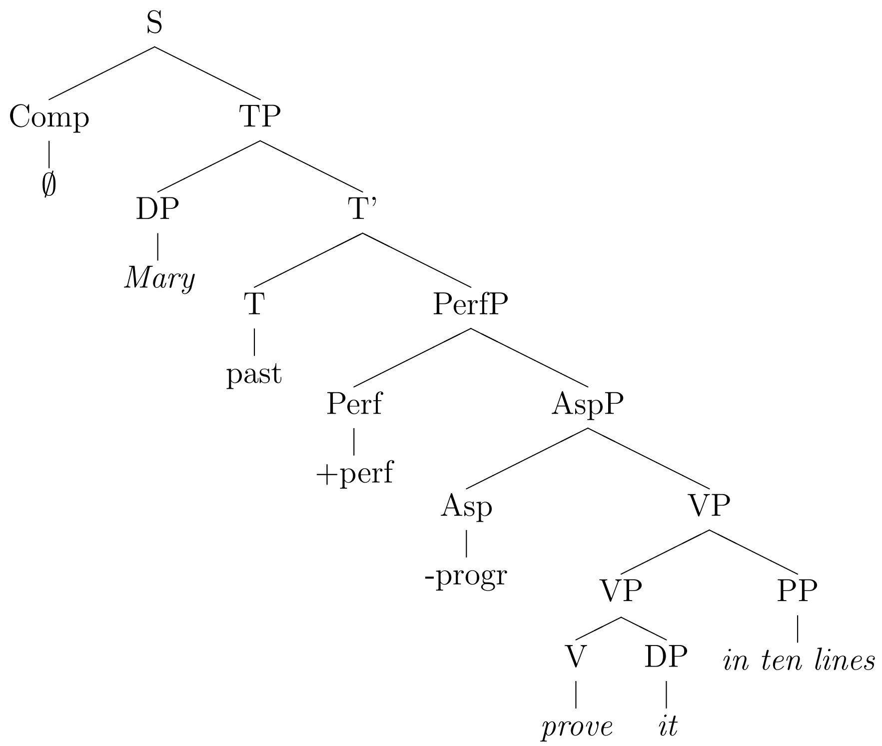

The representation for the lower VP of (4.132) can be computed using principles that have already been discussed and applied, including in particular the referential presupposition introduced by *it*. It is in the transition from the representation of the lower VP to that of the upper VP that implements our ad hoc solution for the contribution of *in ten lines*. We adopt a predicate 'in-ten-lines' ' as part of our logical form language to represent the contribution made by this PP. (4.133.a) gives the lower VP representation and (4.133.b) that of the upper VP.

(4.133)a. $$<e'_{ref}, y| < \{ \begin{array}{|c|}\hline y? \\ \hline \text{non-human}(y) \\ \hline \end{array}_{an.3spr.} \}, \begin{array}{|l|}\hline \\ \hline e'\text{:prove'}(\underline{x}_1,y) \\ \hline \end{array} >>$$

b. $$<e'_{ref}, y| < \{ \begin{array}{|c|}\hline y? \\ \hline \text{non-human}(y) \\ \hline \end{array}_{an.3spr.} \}, \begin{array}{|l|}\hline \\ \hline \text{in-ten-lines'}(e') \\ e'\text{: prove'}(\underline{x}_1,y) \\ \hline \end{array} >>$$

Nothing happens in the transition from the upper VP to AspP, and the transition from the AspP-representation to the representation of PerfP involves the by now familiar transformation of event descriptions into corresponding result state descriptions. I am assuming here that *prove* is a target state verb, with as result the existence of the proof whose 'creation' it is used to describe. Accordingly the relevant perfect operation is that involving the predicate Res, as shown for instance in (3.90) in Section 3.8. The resulting PerfP representation is shown in (4.134).

(4.134)

$$<s'_{ref}, y| <\{ \begin{array}{|c|}\hline y? \\ \hline \text{non-human}(y) \\ \hline \end{array}_{\,an.3spr.} \},$$

$$\begin{array}{|c|}\hline e' \\ \hline s'\text{: Res}(e', {}^\wedge e''. \begin{array}{|l|}\hline \\ \hline \text{in-ten-lines'}(e'') \\ e''\text{: prove'}(\underline{x}_1, y) \\ \hline \end{array}\ ) \\ \\ e' \supset\subset s' \\ \hline \end{array} >>$$

We now come to the contribution made by the feature past, with its new entry given in (4.131). (4.135) shows the result of this, the output representation of (4.131) when its input representation is instantiated to (4.134). Part of this application is the selection of the second $\overset{!}{\vee}$-disjunct from the output representation of (4.131) by the state description (4.134) that serves as input representation.

$$\text{(4.135)}\ <t', s'_{ref}, y \ |<\{ \begin{array}{|c|}\hline t'? \\ \hline t' \prec n \\ \text{TPpt}(t') \\ \hline \end{array}, \begin{array}{|c|}\hline y? \\ \hline \text{non-human}(y) \\ \hline \end{array}_{an.3spr.}\},$$

$$\begin{array}{|c|}\hline e' \\ \hline t' \subseteq s' \\ s'\text{: Res}(e',{}^\wedge e''. \begin{array}{|c|}\hline \\ \hline \text{in-ten-lines'}(e'') \\ e''\text{: prove'}(\underline{x}_1,y) \\ \hline \end{array}\ ) \\ e' \supset\subset s' \\ \hline \end{array} >>$$

The remaining construction steps needed to complete the construction of the preliminary DRS for (3.1.b) are familiar. The result is shown in (4.136).

(4.136)

$$<t'_{ref.pr}, y_{ref.pr} \ |<\{ \begin{array}{|c|}\hline t'? \\ \hline t' \prec n \\ \text{TPpt}(t') \\ \hline \end{array}, \begin{array}{|c|}\hline \\ \hline \text{non-h'n}(y) \\ \hline \end{array}_{an.3spr}, \begin{array}{|c|}\hline \\ \hline \text{Named}(x, Mary) \\ \hline \end{array}_{pr.na.}\},$$

$$\begin{array}{|c|}\hline s' \ \ e' \\ \hline t' \subseteq s' \\ s'\text{: Res}(e',{}^\wedge e''. \begin{array}{|c|}\hline \\ \hline \text{in-ten-lines'}(e'') \\ e''\text{: prove'}(m,y) \\ \hline \end{array}\ ) \\ e' \supset\subset s' \\ \hline \end{array} >>$$

To resolve the presuppositions in (4.136) we need the discourse context provided by the DRS for the first sentence. This DRS is given in (4.137). The presuppositions that are generated by the two definite noun phrases in this sentence, *John* and *the theorem*, have been accommodated, the one generated by *John* in the way typical of proper name presuppositions that is by now familiar to us, and the presupposition generated by the definite description by accommodating some not further specified context predicate C whose extension is assumed to contain a unique entity satisfying the predicate 'theorem'.

(4.137)
$$\begin{array}{|c|}\hline t \quad e \quad j \quad C \quad z \\ \hline t \prec n \quad e \subseteq t \quad \text{Named}(j,\textit{John}) \quad \text{theorem'}(z) \quad C(z) \\ \\ \begin{array}{|c|}\hline z' \\ \hline \text{theorem'}(z') \\ C(z') \\ \hline \end{array} \; \left\langle \begin{array}{c} \forall \\ z' \end{array} \right\rangle \; \begin{array}{|c|}\hline z' = z \\ \hline \end{array} \\ \\ \text{in-twenty-lines'}(e) \\ e\text{: prove'}(j,z) \\ \hline \end{array}$$

(4.137) is no help in the justification of the proper name presupposition in (4.136) that is triggered by *Mary*. This presupposition too can only be accommodated, with a similar effect on the sentence representation that is produced by the *John*-presupposition in (4.137). The *it*-presupposition can be resolved by identifying its referential argument $y_{ref.pr}$ with the dref y from the discourse context (4.137). What ram ins is the TPpt presupposition.

The question how TPpt presuppositions are to be resolved hasn't yet been explicitly addressed. The central principle that governs their resolution is that TPpt presuppositions are anaphoric presuppositions: their resolutions must involve the identification of the question-marked dref from the presupposition with an available dref from the discourse context that is 'of the right sort'. Being of the right sort means that it must be possible to 'identify' the dref with the question-marked time dref of the presupposition. This qualifies time drefs from the discoursed contexts, with which the question-marked time dref can be literally identified, but also eventuality drefs ev, with which the time dref can be 'temporally identified' by setting it equal to $dur(ev)$.

I believe that TPpt presupposition resolution is subject to even stronger constraints, to the effect that the resolving dref from the discourse context must be either an event dref introduced by the main verb of some preceding clause or the time dref introduced by some temporal locating adverbial. We leave the precise formulation of these further constraints to some other occasion.

For the case before us, the last informally specified constraints entail that the discourse context (4.137) contains only one possible resolver for the TPpt presupposition of (4.136), viz. the event dref e. Resolution via this dref leads to the new DRS condition $t' = dur(e)$. Resolving the three presuppositions of (4.136) in the ways described turns (4.136) into (4.138). Merging that representation with representation (4.137) for the first sentence of the two-sentence discourse (3.1) gives us the representation (4.141) for the two sentences together.

(4.138)
$$\begin{array}{|c|}
\hline
t' \quad s' \quad e' \quad m \quad y \\
\hline
\text{Named}(m,\mathit{Mary}) \quad y = z \quad t' = \text{dur}(e) \quad t' \subseteq s' \\
s'\text{: Res}(e',{}^\wedge e''. \begin{array}{|c|} \hline \\ \hline \text{in-ten-lines'}(e'') \\ e''\text{: prove'}(m,y) \\ \hline \end{array}) \\
e' \supset\subset s' \\
\hline
\end{array}$$

(4.139)

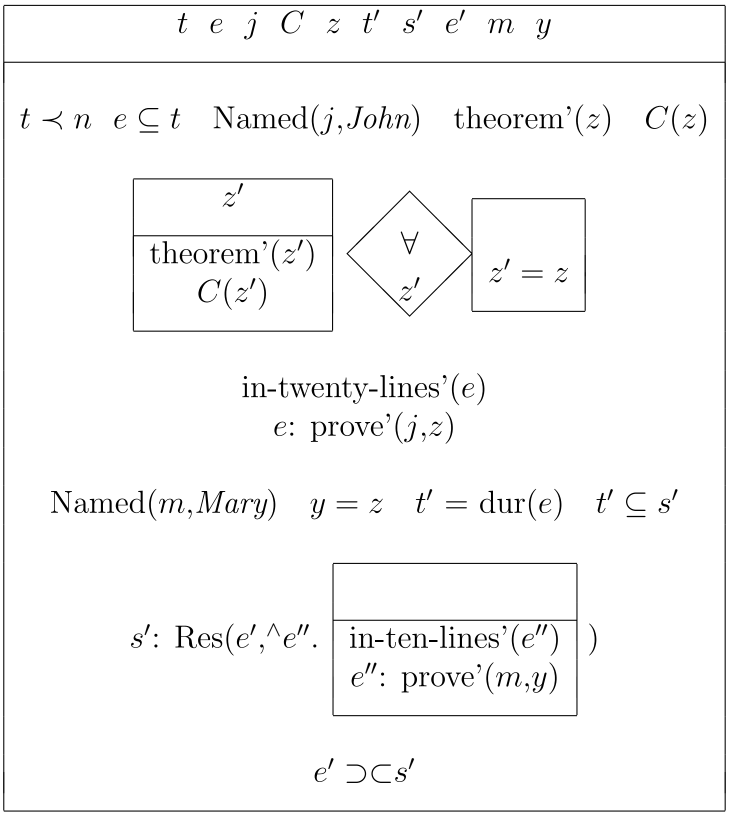

The conditions '$t' = \mathrm{dur}(e)$', '$t' \subseteq s'$' and '$e' \supset\subset s'$' entail that the event e' of Mary proving the theorem in ten lines preceded the event e of John's proving it in ten lines.

4.3.0.5 Context-dependent temporal locations of events

One of the examples discussed at the beginning of Section 3 was the contrasting pair of mini-discourses in (3.2).

(3.2) When Alan opened his eyes he saw his wife who was standing by his bedside.

(i) She smiled.

(ii) She was smiling.

Recall the difference between these two discourses. The final sentence of the first discourse, the Simple Past sentence 'She smiled', conveys that Alan's wife smiled in reaction to his opening his eyes; the Past Progressive sentence 'She was smiling' conveys that she was already smiling when he opened his eyes and that that presumably was the first thing he saw.

Much has happened in these Notes since this example was first brought up. In particular, our discussions have led to the assumption that clauses with progressive verb forms always describe states, whereas Simple Past tense clauses of event verbs describe events. This is relevant both to the DRS construction for the first sentence and to the construction for the second continuation in (3.2). For both progressives DRS construction for the clauses containing them resembles that for the Past Perfect in the second sentence of (3.1.b): the state description comes with a presupposition to the effect that a TPpt must be chosen (in accordance with the tense feature provided by the progressive form of the verb). In what comes next we will only look at how this world for the follow-up sentence 'She was smiling'. The semantics construction of the relative clause in the first sentence of (3.2), as part of a detailed construction for all of the first sentence, is left as an exercise to the reader.

In view of all the details provide in our account of the second sentence of (3.1.b) it will be possible to go through the DRS construction for continuation (ii) of the first sentence of (3.2) fairly quickly. Largely the construction will serve as a kind of consolidation of the construction that is just behind us. At the same time vote purpose is to make it as clear as possible how the construction for 'She was smiling' differs from the one for the non-progressive 'She smiled'.

(4.140) is the DRS for the first sentence of (3.2), which will serve as discourse context for the sentence 'She was smiling'. The syntactic structure of this last sentence is given in (4.141).

$$(4.140)\quad \begin{array}{|c|}\hline t_1\ \ e_1\ \ T_2\ \ e_2\ \ s_1\ \ t_2'\ \ a\ \ w\ \ z \\ \hline \\ t_1 < n\ \ e_1 \subseteq t_1\ \ t_2 < n\ \ e_2 \subseteq t_2\ \ t_2' = \mathrm{dur}(e_2)\quad e_2 \subseteq t_2' \\ \mathrm{Alan}(a)\quad \mathrm{wife'}(w,a)\quad a\text{'s bedside}(z) \\ \\ e_1\text{: open-one's-eyes'}(a) \\ e_2\text{: see'}(a,w) \\ \\ s_1\text{: PROG}(\hat{}e. \begin{array}{|c|}\hline e \\ \hline e\text{: stand-by'}(w,z) \\ \hline \end{array}) \\ \\ \hline \end{array}$$

(4.141)

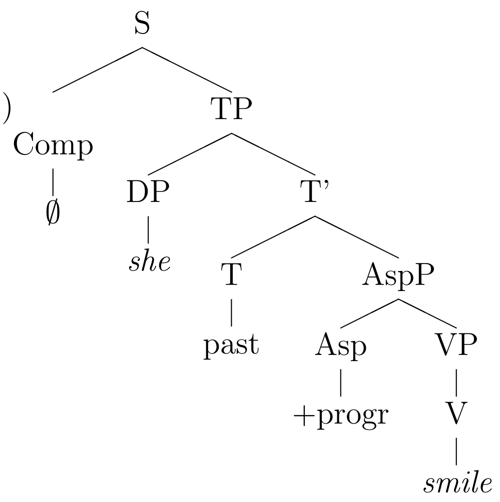

(4.142) gives the point of the construction where the representation has been constructed for the AspP.

(4.142)

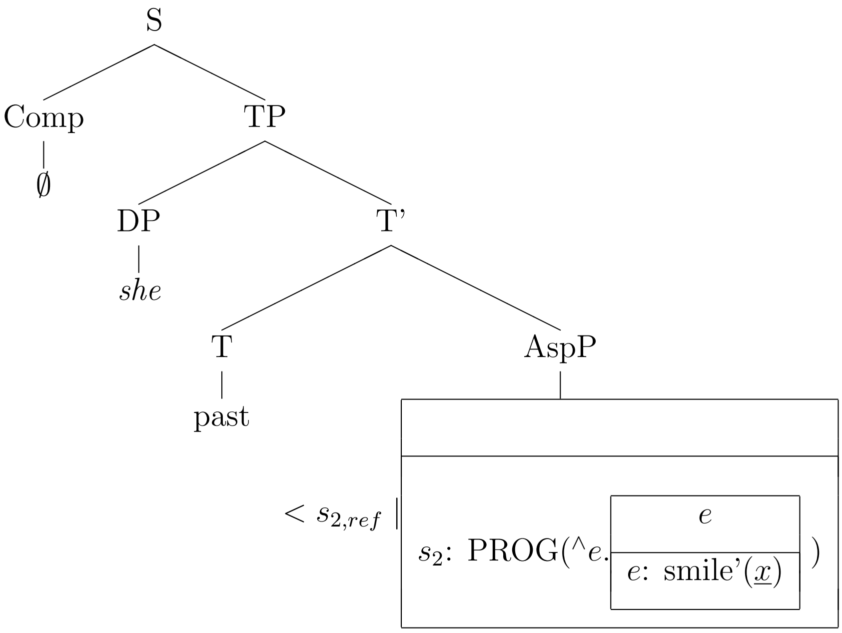

Applying the lexical entry (4.131) for past to the state description associated with AspP leads to the choice of the second disjunct of its $\overset{!}{\vee}$-disjunction. The result is given in (4.143).

(4.143)

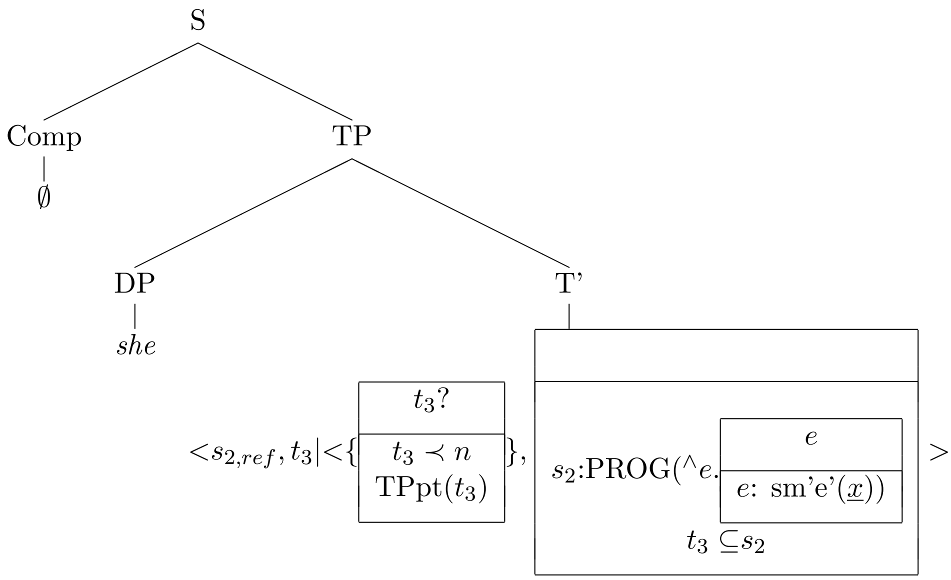

Combining this representation with the subject *she* introduces a presupposition for this pronoun. This presupposition can be resolved by using the dref w of the discourse context (4.140). Furthermore, resolving the TPpt presupposition to e_2 in (4.140) and then merging the DRS for the second sentence with that of the first leads to the DRS in (4.144).

(4.144)

$$\begin{array}{|c|}
\hline
t_1 \;\; e_1 \;\; T_2 \;\; e_2 \;\; s_1 \;\; t'_2 \;\; a \;\; w \;\; z \;\; s_2, t_3 \;\; v \\
\hline
t_1 < n \;\; e_1 \subseteq t_1 \;\; t_2 < n \;\; e_2 \subseteq t_2 \;\; t'_2 = \mathrm{dur}(e_2) \quad e_2 \subseteq t'_2 \\
\mathrm{Alan}(a) \quad \text{wife'}(w,a) \quad a\text{'s bedside}(z) \\
\\
e_1\text{: open-one's-eyes'}(a) \\
e_2\text{: see'}(a,w) \\
\\
s_1\text{: PROG}(^\wedge e. \begin{array}{|c|} \hline e \\ \hline e\text{: stand-by'}(w,z) \\ \hline \end{array}) \\
\\
s_2\text{: PROG}(^\wedge e. \begin{array}{|c|} \hline e \\ \hline e\text{:smile'}(\underline{x}) \\ \hline \end{array}) \\
\\
t_3 \subseteq s_2 \\
\hline
\end{array}$$

4.3.1 DRS Construction for 'She smiled'. The problem of Rhetorical Discourse Structure

In Section 3 we described the difference between the two discourses in (3.2). repeated in the last section, as that between a state that had already begun by the time the event described in the first sentence – that of Alan opening his eyes and seeing his wife – occurred (this is the progressive case of (3.2.ii)) and an event that is understood as following that of Alan opening his eyes and as a reaction to it, as in (3.2.i). The DRS construction for (3.2.ii) provides an account of the first half of this opposition. But what about the second half? When we first discussed this and some other examples in the introduction to Section 3, we noted that in pairs of sentences both of which are event sentences in the simple past, the relationship between the events described by the two sentences can vary, and that, evidently, the variation depends on other factors than the two tense forms. This is clearly shown by the three examples in (3.3), repeated here.

(3.3) a. Fred went to Rosie for dinner. He came home in a state of euphoria.
b. Fred went to Rosie for dinner. He put on clean trousers and his nicest shirt.
c. Fred went to Rosie for dinner. He bought flowers on the way.

When these examples were first presented we only noted that tense and aspect morphology – the fact that wet are dealing with two event sentences in the Simple Past – isn't enough to determine the temporal relation between the described events. That is an important difference between the sentence sequences (3.2.ii) and (3.1.b) on the one hand and (3.2.i) and (3.3.a,b,c) on the other (see Section 3.1). In the former examples the temporal relations between the eventualities that are described by the successive sentences are determined by temporal and aspectual morphology. In the latter cases they are not.

But how then *are* those relations determined for sentence combinations like those in (3.2.i) and (3.3)? That is a question for which there now exists convincing answers in principle. But it is an answer that leads us into a very different theoretical world from the one to which I have been at pains to restrict the discussion in these Notes. The crucial insight is that there are two fundamentally different aspects to the strong intuitions that we often have about temporal relations in cases of the sort illustrated in (3.2.i) and (3.3). On the one hand these intuitions depend on 'world knowledge': On what we

know about how things work – possibly or typically or necessarily – in the world in which we live. In view of such 'knowledge' we have certain expectations about what kinds of eventualities can follow, precede or overlap with other kinds of eventualities, and about which kinds typically follow or do not follow, precede or overlap with other kinds, and which do so necessarily. But that is only one part of it. The reason why our intuitions are so often so clear and so prominent is due to something different, a principle about the coherence of discourse interpretation that might be described as the *Principle of Rhetorical Completeness*. The Principle of Rhetorical Completeness says that it is an inalienable part of the interpretation a piece of coherent discourse or text that every sentence or clause (for multi-clausal sentences) is *rhetorically related* to at least one other sentence or clause. The substance of this claim resides in what forms the *rhetorical relations* between clauses or sentences can take. The first and main point on which 'rhetorical structure theories – theories which acknowledge the central importance of the Rhetorical Completeness Principle and which have undertaken to articulate its substance – differ from each other. There is a striking difference on this point in particular between what are arguably the two best-known accounts of rhetorical relations: *Rhetorical Structure Theory* ((Mann & Thompson n.d.), (Mann & Thompson 1987)) (the first Theory that attributed to Principle of Rhetorical Completeness the central importance it has and that offered a detailed and systematic discussion of it) and SDRT ((Lascarides & Asher 1993*b*), (Asher & Lascarides 2003)), which has been mentioned a couple of times before in these Notes. As observed earlier, SDRT is a theory that builds on DRT and shares many of its basic assumptions. So it is on SDRT that I will focus in the following few remarks on rhetorical structure. (Note well: hese are strictly general and informal remarks, which do not do justice by any stretch to the complexities of SDRT. SDRT has grown into a major theoretical enterprise, that has gone through various stages of development and that should be seen at this point (like DRT) as a family of related theories rather than a single one. (Asher & Lascarides 2003),a monograph of over 400 pages, is just one of a long list of publications in whig SDRT is explained or used.)
SDRT started from the observation that tense, and more particularly the simple past tense, is often unable to predict the temporal relations between the eventualities described by successive sentences in a discourse; what is responsible in such cases for the strong intuitions about temporal relations between events derives from our knowledge or expectations about how things can and cannot develop in time. (4.145) is a classic illustration of this observation.

(4.145)a. John fell. Bill helped him (to get back onto his feet).

b. John fell. Bill pushed him.

In (4.145.a) the event described in the second sentence is naturally understood as following the event described by the first sentence. In this regard the example is like (**??**). (The effect is particularly strong when the parenthetical addition 'to get back onto his feet' is part of the second sentence. That is so, because the presupposition carried by *back* is naturally resolved by assuming that the result state of the event e_2 described by the second sentence is a 'return' of the pre-state of the event e_1 described by the first sentence, and that is possible only if e_2 followed e_1. But even when the parenthetical part is absent from the second sentence of (4.145.a), e_2 is naturally understood as following e_1, because it is a natural and expected thing for someone to help another person who has just fallen and needs to get up.) In contrast with this the event described by the second sentence of (4.145.b) is most naturally understood as the cause of the event described by the first sentence, and therefore as preceding it.

When we compare (4.145.a) and (4.145.b), it is hard to escape the conclusion that the difference in temporal relation has to be explained in terms of the underlying causal connection between the events from first and second sentence. In (4.145.b) it is natural to understand the second event as the cause of the first event, and therefore as preceding it (or at the very least as starting before it); and when the context doesn't clearly point in a different direction, as it doesn't in the presentation of (4.145.b) here because no context has been provided, then this perception – that e_2 preceded e_1 in time – is well-neigh impossible to resist. In (4.145.a) the causal connection between the two events is evidently a different one. Inasmuch as there is any causal connection at all, it is an opposite one: John's fall provided the occasion for Bill's help; if John hadn't fallen, then the event described by 'Bill helped him' wouldn't have occurred either.

4.3.1.1 Rhetorical relations with temporal implications

SDRT (and other theories of rhetorical structure) account for the different temporal relations in (4.145.a) and (4.145.b) by taking it that different rhetorical relations between first and second sentence are involved in the interpretations of (4.145.a) and (4.145.b). The relation involved in (4.145.b) is referred to in SDRT and elsewhere in the rhetorical structure literature as *Causal Explanation*: this is the relation that holds between the first sentence S_1 of a succession 'S_1. S_2' of two single clause Simple Past tense sentences when the second sentence S_2 provides a causal explanation for the first sentence

S_1. More generally, causal explanation is the relation that holds between two sentences or clauses in a discourse or text iff the second term of the relation plays the role of providing a causal explanation for the first. The rhetorical relation that SDRT and other theories of rhetorical structure assume obtain in the case of (4.145.a) is usually referred to as *Narration*. It is the relation that holds between S_1 and S_2 in those cases where the event described by S_2 is understood as the kind that can be expected to be the next event after the event described by S_1.

Narration has a kind of default status as relation between successive sentences in a narrative piece of discourse or text because narrating events in the order in which they occurred is a default strategy for telling stories. Interpreters of narrative texts and discourses will therefore be disposed to taking Narration to be the rhetorical connection between successive sentences so long as that doesn't contravene strong prejudices about what types of events can and cannot follow other types. ((4.143.b) is one example where interpreters do have contravening intuitions and will assume the rhetorical relation to be causal explanation, thereby overriding the default relation Narration.)

Narration and Causal Explanation are two of the rhetorical relations that are acknowledged by, I believe, all substantive approaches to rhetorical structure (if perhaps not always under these names). But what other relations are there? On this point approaches differ, and they differ widely. Some of them, including the one advocated by Mann and Thomson, advocate a long list of relations, with fine discriminations between them. SDRT has been much more 'conservative' on this point, invoking as small a repertoire of relations as one can get by with to achieve the aims that the theory is designed to achieve. For SDRT the constraints on the set of rhetorical relations that a rhetorical structure theory should adopt are on the one hand that the repertoire of adopted relations must capture with sufficient precision how particular texts and discourses are in fact interpreted and on the other by the requirement that the theory must be able to account for how interpreters chose particular relations as the rhetorical relations between pairs of clauses in those texts and discourses. The first constraint imposes a kind of lower bound on the repertoire of rhetorical relations – without a minimum of rhetorical discriminations no intuitively plausible theory will work – whereas the second imposes a kind of upper bound: the task of determining what the rhetorical relation is between a given pair of clauses or sentences in a given discourse becomes more difficult as there are more relations to be chosen from. Since SDRT has focused especially on the problem of how rhetorical relations are chosen as an integral part of discourse interpretation, it has been under a strong pressure to keep the relation repertoire manageably small.

As I said, these remarks are not the preamble of anything like a proper presentation of SDRT. In particular I make no attempt to present a complete list of rhetorical relations and discuss their respective properties. But there are nevertheless a couple of general features of rhetorical relations that have become an integral part of SDRT – some adopted from earlier proposals by others, some discovered by the proponents of SDRT themselves – that must be mentioned. The first of these is the distinction between coordinating and subordinating rhetorical relations. SDRT treats Narration as a coordinating relation and Causal Explanation as a subordinating relation. Between them the coordinating and subordinating relations between sentences and clauses of a coherent text or discourse impose a structure on it that in first approximation can be represented in the form of a directed multi-graph, with two kinds of edges, 'horizontal' edges (or 'left-right' edges) and 'vertical' edges (or 'subordinating' edges). When a coordinating relation holds between S_1 and S_2 (in this order, i.e. S_1 is the first term of the relation and S_2 the second), then S_2 is to the right of S_1 in the discourse structure imposed by the rhetorical relations; when a subordinating relation holds between S_1 and S_2, that means that S_2 is a subordinate to S_1; this roughly means that S_2 is a way of going into more detail about the topic introduced by S_1, or is a constituent of a way of going into more detail about that topic. In this sense the two sentence discourses (4.145.a) and (4.145.b) have the graph structures shown in (4.146.a) and (4.146.b). More revealing is the structure in (4.147.b) for the discourse in (4.147.a).

(4.146)a. S_1——S_2

b. S_1
|
|
S_2

(4.147)a. John fell. Bill pushed him. He put his hands against his chest and gave him a mighty shove. John got up. He had hurt himself badly.

b.

$$\begin{tikzcd}
 & \mathrm{S}_1 \arrow[rr, dash] \arrow[d, dash] & & \mathrm{S}_5 \arrow[d, dash] \\
 & \mathrm{S}_2 \arrow[dl, dash] \arrow[dr, dash] & & \mathrm{S}_6 \\
S_3 \arrow[rr, dash] & & \mathrm{S}_4 &
\end{tikzcd}$$

From an SDRT perspective (4.147.a) consists of 5 'discourse units': the first sentence, the second sentence, the two conjuncts of the third sentence and the fourth sentence. Each of these units must be rhetorically connected to another one – more precisely, each unit u must be rhetorically connected with some preceding unit u', in the sense that there is a rhetorical relation R between these two units in which the preceding unit u' is the first argument and the unit u the second; in other words $R(u',u)$. This means that each unit except the very first one is connected in this way to at least one preceding unit, whereas the initial unit is connected to at least one unit following it by a relation in which the initial unit is the first argument.

I have repeatedly used the phrase 'at least' in the last paragraph because it is part of SDRT's concept of rhetorical structure that discourse units can be and often are rhetorically connected with more than one other unit. Not only will a non-initial unit u typically be connected both with a later unit u'' by a relationship $R'(u,u'')$ and with an earlier unit u' by a relationship $R(u',u)$, it is also not uncommon for a unit to partake in additional relations. The discourse structure for (4.147.a) shown in (4.147.b) provides examples of those possibilities. In this diagram the discourse units of (4.147.a) have been labeled as S_1, ..., S_6 in order of appearance. There are three rhetorical relations represented by the edges in (4.147.b), Narration, Causal Explanation and Elaboration, a relation not yet mentioned. Of these Narration is coordinating, as already noted, whereas the other two – Causal Explanation and also Elaboration – are subordinating. More specifically, the rhetorical connections between the discourse units of (4.147.a) that are depicted in (4.147.b) are as follows: (i) The rhetorical connection between S_1 and S_2 is given by $\mathrm{CaEx}(\mathrm{S}_1,\mathrm{S}_2)$ and the connection between S_1 and S_5 is given by $\mathrm{Nar}(\mathrm{S}_1,\mathrm{S}_5)$. (ii) The connections between S_2, S_3 and S_4 represent a more complicated set of rhetorical relations. The relation between S_2 and S_3 is Elaboration: S_3 tells us more about the event described by S_2 – it 'elaborates on' the description provided by S_2. In SDRT Elaboration is among the subordinating relations.[52]

The novel complexity arises in connection with S_4. S_4 is part of the elaboration of S_2 that was initiated by S_3. This description of what it does is reflected by a pair of rhetorical relations. The first is the relationship of Elaboration with S_2, in which S_4 stands in the same way that S_3 does. The second is a relation to S_3 (one that is compatible with the use of the conjunction *and*). It is usually assumed that this is an instance of Narration, which provides the rhetorical 'kit' of the 'sub-discourse' that consists of S_3 and S_4, the two sentences that together make up the elaboration of S_2. Since Elaboration is subordinating and Narration is coordinating, the two relations in which S_4 is involved according to this analysis give rise to the triangle-like structure in the lower half of (4.147.b).

The one but last sentence of (4.147.a) is S_5. This sentence is naturally understood as a narrative continuation of S_1. Hence the horizontal edge that connects it with S_1. The final sentence S_6 with its past perfect is linked to its predecessor in the manner we have seen in detail when dealing with the second sentence of (3.1.b). There is a sense in which such 'flashbacks' are like elaborations. They typically provide relevant background to the topical event, state or situation, by telling something about its 'prehistory'. Sometimes this relation is treated as a rhetorical relation in its own right, called 'Background' or 'Flashback', and sometimes as a case of Elaboration, which is what for reasons of simplicity, we will be assuming here. Either way, the relation is considered a subordinating one. In (4.147.b) S_6 is construed as standing in such a relation to its immediate predecessor S_5, with S_6 therefore appearing directly below S_5.

4.3.1.2 Rhetorical relations without causal implications

The rhetorical relations we have been talking about all deserve the name 'rhetorical' insofar as they are relations between bits of discourse or text

[52]The elaboration relations of (4.147.a) also illustrate a further point. I said that the sentences S_3 and S_4 tell us more about the event introduced by S_2. That is true enough. But saying more about a certain event can take different forms. In the case before us it seems plausible to say that S_3 and S_4 describe their own events e' and e'' and that these are parts of the event e described by S_2, or perhaps that they jointly make up e, in the sense that $e = e' \bigoplus e''$. But there also cases where elaboration cannot be analyzed in terms of mereological parts or sums. For instance S_2 of (4.147.a) could also have been followed by the sentence 'He did it quietly and from behind' (cf. the discussion of VP conjunctions in Section 3.10.2.)

that testify to the coherence of the discourse or text to which they belong, showing the relevance of each bit in relation to the interpretation of the part of the discourse that has already been put into place. It seems difficult, and is perhaps fundamentally impossible, to explicate what rhetorical relations are, and what the discourse coherence is that requires them and of which they are the local manifestations, in more basic terms: We have a conception of discourse and text coherence which requires that no parts of them are entirely unrelated to the rest – every part must be understandable as making a contribution to the discourse that is relevant to what has been said so far – and the right choice of rhetorical relations will make it possible to state what different forms such relevance of parts to the whole can take. There is however an important distinction among rhetorical relations to which I have not yet drawn attention. The few rhetorical relations that have been mentioned so far all have something to do with what is possible in the world that the discourse or text is talking about. This is clearest for Causal Explanation. The claim that Causal Explanation holds between discourse units u and u', i.e. that $\mathrm{CaEx}(u', u)$, entails the claim that the eventualities ev and ev' described by u and u' stand in a corresponding causal relation: ev' must be the or a cause of ev. Attributing the rhetorical relation to u and u' is therefore legitimate only to the extent that it is legitimate to attribute this causal relation to the eventualities described. It is for this reason that world knowledge is important to the attribution of this relation to pairs of discourse bits: the relation between the bits u and u' is plausible only to the extent that world knowledge justifies the causal relation on which the rhetorical relation builds. And conversely, adopting the relation of Clausal Explanation to the discourse bits u and u' entails a certain understanding of the episode or situation described: One cannot ascribe Clausal Explanation to u and u' without assuming that ev was the cause of ev'. And so, in particular, one cannot ascribe Clausal Explanation to u and u' without assuming that ev preceded (or at least: started no later than) ev'.

Narration is sensitive to world knowledge in a similar way and carries similar commitments about the described situation or episode. To assume $\mathrm{Nar}(u', u)$ is to assume that the event e described by u is a plausible next event after the event e' described by u' in the light of what we consider plausible developments of episodes in the real world. And so one cannot assume the rhetorical relation Narration without assuming that the event e described by u followed the event e' in time.

But not all rhetorical relations that a workable account of rhetorical structure have such implications for the situation, episode or history that the text describes. There are also relations that contribute coherence to a discourse and that carry no implications about any causal-like connections between the eventualities described. One such relation is *Contrast*. An example is the two sentence discourse (3.1.a), with which the introduction in Section 3 to PART II began. The two sentences contrast the proofs by John and Mary in terms of their length (20 lines as against 10 lines). That is enough to make the two sentences of (3.1.a) relevant to each other and the discourse consisting of the juxtaposition of these two sentences coherent. But nothing follows about whether there was any causal relation between the two proving events that these sentences refer to. In particular, nothing can be inferred, as we saw, about any temporal relation between them. That is a difference between a discourse like (3.1.a) and a discourse like (3.2.i), where a temporal inference is possible. And by the same token the difference between (3.1.a) and (3.1.b) indicates that in (3.1.b) it is just the tense form (the Past Perfect) that is responsible for the inference that Mary's proof came before John's proof.

There may be rhetorical relations for which it is not all that easy to decide whether or not they have real world implications in the way that Causal Explanation and Narration do, but Contrast doesn't. But the distinction is nevertheless a real one, with Contrast clearly belonging to the camp of the relations that do not require support from World Knowledge. When the rhetorical relation that connects two discourse units is of this second 'causation-free' kind, then discourse coherence as such requires no further justification in World Knowledge terms, although real world knowledge may still play a part in determining temporal relations.

In fact, as far as temporal relations are concerned, cases of Contrast differ. On the one hand there are those like (3.1.a), where no conclusion is possible. Then there are cases like (3.1.b), where it is the tense forms of the two sentences or clauses that tell us which event occurred before the other. And thirdly there are cases like those in (4.148.a,b), where there is a relation of Contrast between the two sentences – this is arguably enough to give coherence to the discourse – but where it is nevertheless, it is natural to think of the events the contrasting units describe as related to each other by virtue of being both parts of a larger complex event, such as that of Mary and John baking a cake together. And when the events are thought of in this way, then anyone who knows anything about baking cakes at all (including ignoramuses about cake making like myself) will infer that the dough making event came before the baking event. And that is so irrespective of the order in which the two part events are mentioned (see (4.148.b)).[53]

(4.148)a. John made the dough. Mary took charge of the actual baking.

b. Mary took charge of the actual baking. John made the dough.

c. Mary and John decided to make a cake for Sue's birthday. John made the dough. Mary took charge of the actual baking.

Note however in this connection that the natural pressure to interpret the events described by the two sentences of (4.148.a,b) as parts of a single enterprise can be reflected in the structure of the discourse. An example is (4.148.c), in which the joint enterprise is mentioned in the first sentence. The next two sentences of (4.148.c), the two sentences of (4.148.a), are then understood as a two sentence Elaboration of this first sentence. Perhaps the Elaboration relation between each of these sentences and the first sentence is all that is needed to justify the coherence of (4.148.c). But whether or not that is true, the conclusion that John's dough making came before Mary's baking will be forthcoming.

Summarizing the upshot of our discussion of Contrast: Among the relations through which discourse coherence can manifest itself there are those like Contrast which only have to do with the way in which information is structured in the discourse. In such cases there is no need to appeal to World Knowledge. One consequence of this is that in some cases, such as for instance (3.1.a), no conclusion needs and can be drawn about a temporal relation between the eventualities described by the two discourse units that are rhetorically linked by Contrast. But that of course does not mean that it is never possible to infer a temporal relation between the eventualities described by units linked in this way. In fact, there are two ways in which this can happen: either the temporal relation is made explicit by tense and aspect morphology, as in (3.1.a), or it can be inferred on the basis of World Knowledge, as in (4.148) or (3.3).

[53]It is a good question, which I will not try to answer here, whether in (4.148.a) and (4.148.b) Contrast is all that is needed as rhetorical justification of the two sentence discourse. Perhaps some additional relation between first and second sentence is part of the justification as well, and perhaps this second relation is different for (4.148.b) from what it is for (4.148.a). I will have no more to say about the possibility of more than one rhetorical relation between the same discourse units. I will assume for the remainder of this discussion that there can never be more than one relation between any two units, but I am sure that this will prove the right assumption to make in the end.

4.3.1.3 The power of tense and aspect

There is one further point that deserves to be stressed. The difference between (3.1.a) and (3.1.b) is that while (3.1.a) permits no inference to a temporal relation between the events described by its two sentences (but without producing an impression of incoherence), (3.1.b) conveys that the second of those events preceded the first. This is a feature of certain tense-and-aspect combinations – those where an event sentence in the Simple Past is followed by a the state describing sentence that can also be analyzed as a Simple Past, but where the tense form combines with an aspect operator that turns its input into a state description, such as the Progressive – in which case the tense form is often described as a 'Past Progressive' – or the Perfect, in which case the tense form is what is usually referred to as a 'Past Perfect'. Such tense-and-aspect combinations locate the state descried by the second sentence as temporally including the event described by the first, with different consequences for the event that is involved in the semantics of the second sentence. What deserves to be stress is that the implications carried by such tense-and-aspect combinations cannot be overruled by World Knowledge. Consider for instance the following two sentence discourse (4.149).

(4.149)John fell. Bill was helping him (to get back onto his feet).

The second sentence of (4.149) can only be understood as something that was going on when the event described by the first sentence took place. This is so even when the parenthetical part of the second sentence is included. With that inclusion (4.149) is a bit difficult to interpret; the addition of *again*, as in 'John fell again', would have been more natural, suggesting that John had fallen once already, that Bill was trying to get him back to his feet, but that John then had another fall, slipping from Bill's supporting arms perhaps. The point is that the combination of the Simple Past in the first sentence and the Past Progressive in the second *forces* an interpretation of the second sentence as describing a state that temporally included the event described by the first sentence. In (3.2.i) that is unproblematic, because world knowledge doesn't contradict such an interpretation or make it implausible. (4.149) confronts the interpreter with a problem because World Knowledge doesn't support the temporal relation that tense and aspect morphology imposes.

Discourse graphs like that in (4.147.b), which display rhetorical relationships of which some are coordinating and others subordinating, also impose an important constraint on a question about which we haven't so far said anything: When a new discourse unit has to be rhetorically connected with some earlier unit, which units from the preceding part of the discourse are possible candidates for such a connection? (Note: this is the same sort of question as the one which constituents in a discourse are potential antecedents for an anaphoric pronoun or description, but the answer is determined by quite different principles.) According to SDRT and other theories of rhetorical structure the main structural constraint on what can be the 'rhetorical anchor' for a new discourse unit is the so-called right frontier' constraint. The right frontier of a discourse graph is, at each stage of building the graph, the chain of nodes that starts with the rightmost node at the top level, looks if that node has nodes subordinate to it, chooses, in case that is so, the rightmost one of those, looks if that one has nodes subordinate to it, picks the rightmost one of those in case there are any, and so on. For the construction of (4.147.b) this means that after just S_1 has been introduced as graph node and no other nodes have been introduced yet, then (trivially) the right node frontier consists just of that one node; after S_2 has been added, the right node frontier consists of these two nodes; after addition of S_3 the right node frontier consists of all three nodes S_1, S_2 and S_3. Addition of S_4 – note that at this point both S_2 and S_2 are members of the frontier, so both can serve as anchors for S_4 – leads to a graph whose frontier consists of S_1, S_2 and S_4; S_3 is no longer accessible at this point. Adding S_5 as coordinated with S_1 eliminates S_1 from the right frontier and with it all members of the frontier subordinate to it. So at this point the frontier consists solely of S_5. Finally, after subordination of S_6 to S_5 the frontier consists of S_5 and S_6, though at this point the question is no further consequence since the discourse ends with S_6.

4.3.1.4 Computing rhetorical structure

These are only some of the most basic principles of SDRT, which give a first impression of what the rhetorical structure can be like of discourses and texts consisting of more than two discourse units. But it is enough to reveal some of the issues that have to be resolved by a theory which deals with rhetorical structure along the general lines we have sketched. First, as noted, one has to settle on a repertoire $\mathcal{R}$ of rhetorical relations: This choice means that for any coherent discourse D each non-initial discourse unit u belonging to it must be interpretable as standing in one of the relations from $\mathcal{R}$ to some other discourse unit that is 'accessible' from u (i.e. belongs to the right frontier at the point when u is added to the discourse graph). Second, a criterion has to be developed to decide what the discourse units of a discourse D are, which must be rhetorically connected with other units as part of a coherent interpretation of D. We have seen a glimpse of some of the difficulties of this issue, which have to do with the possibility that discourse units may be proper parts of other discourse units, such as when the discourse units S_3 and S_4 are treated as constituents of a larger discourse unit that consists of the two of them. It has also been suggested that discourse units are often smaller than full clauses.[54] Here we will assume for simplicity that all the discourse units are full clauses and that every clause is a discourse unit; on this assumption D is a succession of disjoint units.

The most important question, of course, to which a theory of rhetorical structure has to find an answer is this: Given a discourse D, a segmentation of it into discourse units and a non-initial discourse unit u of D, how is u to be connected by one or more rhetorical relations to one or more preceding discourse units of D?

Formally the problem can be described as follows. Suppose as given: (i) the set $\mathcal{R}$ of possible rhetorical relations, together with a specification for each $R \in \mathcal{R}$ whether R is coordinating or subordinating – $\mathcal{R}$ could be specified, for instance, as the union of a pair $<\mathcal{R}_c, \mathcal{R}_s>$ of a set $\mathcal{R}_c$ of coordinating relations and a set $\mathcal{R}_s$ of subordinating relations – and (ii) a segmentation DU(D) of D into its discourse units. Let us assume further that a semantic representation has already been put in place for the discourse up to u and that it is of the following form: A DRS $K_{u'}$ has been constructed for each of the discourse units u' in DU(D) that precede u in D as well as a DRS K_u for u itself. In addition a DRS K^u has been constructed for the part of D that precedes u. (Somewhat simplifying, we may assume that this DRS is the merge of the DRSs for the individual members of DU(D) that precede u.) In addition, we assume that rhetorical relations have been chosen for each of the non-initial discourse units in DU(D). It may be assumed that these are given as a set of relation specifications of the kind considered earlier i.e. as predications of the form '$R(u',u'')$' with $R \in DU$(D). From this set of clauses a multi-graph of the sort exemplified by (4.147.b) can be reconstructed and from that the current right frontier, which gives us the set of discourse units that can serve as potential rhetorical anchors for u. Let AN(D,u) be the set of discourse units of D on this right frontier.

Given these assumptions the problem of finding a rhetorical justification for u takes the form of choosing at least one of the pairs $<R, u'>$, where $R \in \mathcal{R}$ and $u' \in AN$(D,u) and justifying this choice. Choosing $<R, u'>$ amounts

[54] An example would be the PP *with a ratchet* in the sentences of (3.231) mentioned in footnote 60.

to adopting the predication '$R(u',u)$' and adding this clause to the set of rhetorical relationships that have already been adopted in the process of interpreting D.

What is involved in justifying the adoption of the relationship $R(u',u)$? There are at least two different ways to look at this question, both of which I will touch upon briefly. The first view has been integral to the logic-based approach to natural language semantics that has been central to all of formal semantics, including DRT and SDRT. According to this view, justifying the adoption of a clause '$R(u',u)$' must take the form of deriving it from relevant information. This information will include at a minimum K^u and K_u and in addition it will contain various kinds of 'world knowledge'. But in what form should world knowledge be available and what can the derivations be like that license the adoption of rhetorical relationships? Asher, the father of SDRT, has made significant contributions to both these questions, through his work on *Commonsense Entailment* (Asher & Morreau 1990) as an inference regime suitable for (among other things) the derivation of rhetorical relationships and his work on axiomatization of relevant bits of world knowledge in a form that is compatible with this inference regime (Lascarides & Asher 1993*a*), (Asher & Lascarides 2003). In addition, as discussed in detail in these publications, 'bridging principles' are needed that connect, for instance, causal relations between events of certain types with rhetorical relations between clauses that describe events of such types. Unfortunately, the correct definition of such inference systems turns out to be extremely difficult. First, how are we to make all the needed world knowledge available in advance, given how much potentially relevant world knowledge there is in principle? But what makes the task perhaps even more daunting is a problem that rears its head more or less from the start. (That is, it is a problem even for very small fragments of English with highly restricted vocabularies.) This is the problem that inference based on world knowledge tends to be defeasible: what counts as a valid derivation from given information – i.e. from given premises – may cease to count as valid in the presence of additional information, i.e. when the set of premises is extended. This problem has been acknowledged by SDRT from its beginnings: Common Sense Entailment, which is a system of defeasible inference, was in part developed to this end. But to formulate items of world knowledge in such a way that they can be used as premises in Common Sense Entailment derivations together with DRSs and principles that relate real world assumptions to predications expressing rhetorical relationships has proved a truly daunting task.[55]

[55]As I say this, I want to acknowledge explicitly how much has been accomplished in

I should hasten to add, though, that not all inferences to rhetorical relationships are problematic in this way: Not all of them involve the verification of cause-like relations between types of events. For instance, this problem does not arise for relations like Contrast. Establishing a relation of Contrast involves first and foremost an analysis of the formal structure of the clauses involved. Examples are the sentence pairs in (3.1) (repeated in the last section). Take for instance (3.1.a). Its two sentences form a contrasting pair in the sense that (i) their subjects *John* and *Mary* are non-coreferential (they refer to distinct individuals) and that the predicates expressed by remainders of the two sentences can be construed as being in some sense 'opposite'. The 'opposition' between the two predicates can also be identified by syntactically parallel constituents, the Prepositional Phrases *in twenty lines* and *in ten lines.*; sand here it is also important that these two phrases can be recognized as expressing distinct conditions (more precisely: incompatible properties of the events described by the two sentences). I am not claiming that all cases of Contrast can be established by such largely structure-based analyses of the clauses involved. But many cases follow the pattern just described. And at least in the example we are looking at very little if any world knowledge is involved. (Perhaps the knowledge that if a proof is of twenty lines it can not be a proof of ten lines and conversely is something like world knowledge. But the crucial point is that no world knowledge needs to be used – or for that matter can be used; see below - to establish a temporal relation between the events described by the two events. This is generally true of clauses that are rhetorically related by Contrast; and it is also generally true that Contrast can suffice to make a two sentence discourse coherent all by itself. This is the explanation of why a discourse like (3.1.a) is coherent. although nothing can be inferred about the temporal relations between the eventualities described by the contrastively related discourse units.

Contrast is not the only rhetorical relation that can be established on structural grounds alone. But for the point to be made here – that not all rhetorical relations carry implications about temporal relations – one example of such a relation suffices.

There is also another aspect of the computation of rhetorical relations that should be mentioned. The examples we have looked at so far provided little if any overt clues of what the rhetorical relationship between the relevant discourse units was. Often enough that is the situation – it isn't a spe-

this direction, especially in (Asher & Lascarides 2003).

cial, artificial feature of the examples chosen – and it is cases of this sort that constitute the most important challenges to theories of rhetorical discourse structure. But it is also common for sentences and clauses to contain overt markers which indicate how they are to be rhetorically linked to some other discourse unit (usually one that is adjacent to the one containing the marker). The most pervasive type of case is that where a subordinate clause is rhetorically related to 'its main clause' (i.e. the clause to which it is directly subordinated by the particle or phrase that heads it). Examples of such particles are *because*, *so that*, *due to the fact that*, ..., as well as certain uses of *although*, *in spite of the fact that* and others. *because* and *due to the fact that* explicitly express that the eventuality described by the clause they head is the (or a) cause of the eventuality described by the clause that serves as the subordinate clause's adjunction site. With *so that* the articulated causal relation is the reverse. *although*, *in spite of the fact that* express another, more complex causal relation between the eventuality they describe and the eventuality described by the adjunction site: Eventualities of the former type can in general be expected to prevent the occurrence of events of the latter type, but in the situation that is the topic of the given discourse an event of the latter type occurred nonetheless (i.e. in spite of the fact that an event of the former type occurred as well).

Not all heads of subordinate clauses indicate rhetorical relations that involve cause-like relations in the way that Causal Explanation and Narration do. Contrast, for instance, can also be indicated or made explicit – for instance by *whereas* and also by *while*in one of the uses of that word. And there are also such 'heavy' phrases as *in contrast* or *in oppositional to this*. (Interestingly, these sound odd more often than not when they are added to the second of two sentences which can be recognized as standing in a Contrast relation; which goes to show that 'Contrast' as a technical concept of rhetorical structure theories is not quite what we normally describe when we use the word *contrast*.)

Overt clues to rhetorical relationships can also be found in main clauses. *nonetheless* and its close equivalent *nevertheless* are cases in point. As demonstrated in the last sentence of the preceding paragraph, these adverbs perform the same kind of function as the subordinate clause conjunctions *although* and *in spite of the fact that*: they too contribute the information that the main clause event occurred although an occurrence of an event of the type described by the rhetorical anchor – typically in such cases the main clause immediately preceding the clause with *nonetheless/nevertheless* – also occurred.

On the whole natural languages tend to have substantial repertoires of words and more complex expressions to indicate rhetorical relations; and to some extent these repertoires are open-ended: speakers can come up with new expressions to clarify their rhetorical intentions, and occasionally they do. The study of such discourse particles is an important and challenging part of linguistics.[56] But the challenge is, for all we can tell at the current time, a modest one when compared with that of developing a system of default inference that is capable of deriving, for substantive fragments of English and other natural languages, what rhetorical relations hold between different discourse units of coherent discourses and texts.

4.3.1.5 The second perspective

The second perspective I mentioned is part of an approach to the theory of linguistic meaning that is very different from the methodology of everything that we have been doing in these Notes. So it doesn't really belong here. But it has gained so much prominence over the past couple of decades (and is gaining still) that that by itself would justify mentioning it in an introduction to semantics, even if that introduction belongs, you might say, almost to a different scientific universe.

There are two distinct sources of motivation for this perspective, one cognitive, one computational. Both originated, I believe, at least in part from frustration with the snail-paced progress that 'formal semantics' approaches like the one pursued in these Notes were perceived to makeover the decades since formal semantics began with the work of Montague. On the one hand the apparently slow progress confirmed cognitively oriented linguists in their conviction that there was something fundamentally wrong with these approaches; and on the other hand there were those who had been waiting for 'formal semantics' to deliver the foundations on which they could build implementations capable of dealing effectively with computational tasks (such as machine translation or automated question answering) at acceptable performance levels. The perceived – and justly perceived – failure of formal

[56]The repertoire of fixed expressions for indicating rhetorical relations in English is modest in comparison with what can be found in some other languages. German is an example of a language with a remarkably rich ecology of 'discourse particles', some with what look like quite exotic properties to linguists who previously only concerned themselves with the discourse particles of English. A good deal of progress has been made over the past two decades or so with these exotica.

semantics to deliver what one was waiting for. induced them to look for what they needed elsewhere

The cognitively motivated doubts that 'formal semantics' could be the right way of going about developing a viable account of meaning in natural language, was based on an intuition that the workings of the language processing mind are far more intricate and complex than could be captured by the comparatively simple and 'superficial' descriptions that formal semantics was seen to be working with, and that the only hope for significant advances would lie in buillding artificial systems (on paper, simulated in current digital computers or implemented with the help of some kind of electronic hardware) which incorporate enough of the anatomy and physiology of the networks of cells in the human brain to be able to mimic its performances.

The architecture of the 'cell systems' – or 'neural networks' as they are standardly referred to – used in work inspired by this idea has been very simple and abstract when compared with the actual anatomy of the brain and the physiological processes that drive its processing activities. But nevertheless these seemingly simply structured networks have proved to be an astoundingly powerful instrument for the performance of many sophisticated mental tasks. The way this has been shown is to simulate the workings of such networks on standard digital computers. Indeed, simulation on digital computing machinery has become the standard way of building network-based systems, and that is true in particular of the network-based systems used in computational linguistics.

There are three important features of such network implementations. First, it is de facto impossible to understand in detail how their outputs are related to their inputs. (When their input-output relations appear to be rule-governed, in the sense that all (or a large majority) of the individual input-output pairs satisfy the rule, then there is no way of looking into the system to gain insight into how the internal workings of the system produce this regularity effect.) Second, in order that the system succeeds in performing well, it needs to be trained on inputs of for which the task it is designed to perform is defined. The system's performance will therefore depend crucially on what kinds of 'training data' it is given. Linguistic training data typically consist of 'corpora', collections of texts that have been produced for the purposes and in the contexts that ordinary people – and *not* investigating linguists! – have produced for whatever purposes and in whatever contexts people do use language: various forms of literary texts, newspapers, Wikipedia articles, internet chats, transcripts of telephone conversations etc. In order for these trainings data to be effective, especially when the tasks involve semantics, they (a) need to be copious, so that the system which they are used to train can encounter sufficient numbers of instances of the patterns that matter to the task it is meant to learn to perform and (b) usually the data have to be *annotated* in the right way, providing the system with clues as to what it has to look out for. About what kind of information annotations should provide, and why such annotations tend to be indispensable, can be made clearer in a few words for classificatory systems, systems whose task is to classify items of a certain kind as being of one of two or more distinct types (or as belonging to one of two or more classes). For instance, the task could be that of classifying discourse units in terms of the rhetorical relations that connect them with other discourse units in their discourses. To enable such a system to learn what it should, it will be useful, and almost certainly necessary, for the data to be collections of texts whose discourse units are annotated with specifications of what, according to some independently determined criteria, are the rhetorical relations that connect them to other units.[57]

Third, such systems are, in an essential way, probabilistic. Take again the case of a system whose task it is to classify items of a certain kind as belonging to one of a finite number of possible 'types' P_1,..., P_n. What the system should (after some suitable amount of training) deliver will consist (a) of detecting in the material it receives as input the items that need classification, and (b) for each such item u it should assign to u a probability distribution over the types P_1,..., P_n: for each P_i the system assigns a probability p_i that u is of type P_i with all these probabilities p_i summing to 1. The performance quality of the system is to be assessed in terms of (a) how many items in its inputs it identifies as cases of the kind for which it is has been designed and (b) how 'steeply correct' the probability distributions are that it assigns to individual items, where a probability distribution has a high degree of being 'steeply correct' if one probability value p_i dominates by a large amount all the others, and if that the type P_i for that highest value p_i is the one to which u does in fact belong (according to the independent criteria). (The optimum for 'steep correctness' is of course a probability distribution in which the correct classification gets probability 1 and all the other possible classifications get the value 0.)

In particular, we can think of such a system as dealing with the 'computation'

[57]Note that when a system is trained with such annotated data, then the input-output relations it learns may display rule-like regularities because those very regularities have been imposed on the training data by the annotators, who may follow their own rule-based understanding of the information they have been interacted to annotate.

of predications '$R(u',u)$' for each non-initial discourse unit u from a given discourse D in the following way. Let us assume that the system is given its first task, that of identifying the discourse units of the texts D that it receives as inputs, for free, in that it receives each input D with a segmentation into discourse units ready made. The system's output will then be, for each unit u of D a probability distribution over the finite set of possible predications '$R(u',u)$'.[58] If this is what the system will return for each discourse unit u of D, then the system can also be used to determine a probability distribution over sets of such relationships, one for each discourse unit in D. How the value for such a set is composed out of the values of the predications $R(u', u)$ for the individual u's, is a matter to be addressed separately, and one that we won't explore here. But let's assume, for the sake of argument, one way in which probability distributions about the possible rhetorical anchoring for the individual non-initial discourse units of a discourse D may be turned into a distribution over the possible rhetorical discourse structures for all of D. According to this assumption each such rhetorical structure is identified as a set of predications '$R(u',u)$' for the different non-initial units u of D, with one such predication for each u. According to the assumption the value that the system assigns to each such set is the product of the values it assigns to the individual predications in the set.

Let us assume that a system of the kind alluded to in the paragraphs above assigns values to possible 'rhetorical discourse structures' for texts D. Here again we can assess the 'goodness' of the predication that the system makes about what the rhetorical structure of a given input D is in terms of the 'steep correctness' of the distribution it associates with D. But more generally we may stipulate that the rhetorical structure that the system 'assigns' to D is the set of predications $R(u', u)$ to which the system assigns the highest value. We can then, working backwards as it were, take the system to determine for each u in D as its rhetorical anchoring the predication $R(u', u)$ that belongs to the rhetorical structure assigned to D by the criterion just mentioned.

This is of course no more than a sketch of how a network-based approach

[58] I am assuming here that the different possible rhetorical relationships for any given discourse unit u are mutually exclusive. If this assumption is not made, and non-singleton sets of rhetorical relation predications '$R(u',u)$' for the same u are possible too, then the story is somewhat more complicated – we now have to assume that the system provides a probability distribution over the subsets of the set of all the possible predications of this form, i. e. over the power set of the set of those predications (or perhaps over some subset of the power set, consisting only of those predication sets that satisfy certain coherence constraints).

might handle the assignment of rhetorical relations to discourse units. But even if many of the assumptions we have been making along the way may be challenged, and even if I do not know of any system of the kind discussed that makes reasonably plausible predictions about rhetorical relations of discourse units from actual texts, the general point should be clear. When information about rhetorical relationships is needed in an account of semantic processing such as the one we have developed in these Notes, where such information is sometimes required for dealing with temporal relations between events, then the best we can do in the absence of a proper treatment of rhetorical relations is to adopt the stopgap measure of appealing to an oracle that will tells what the rhetorical relation is in those cases where our construction algorithm needs the information.

4.3.1.6 Back to 'She smiled'

We now almost have all that we need in connection with the second sentence 'She smiled' of (3.2.i). Our oracle $RJ_{\mathbf{D}}(u)$ gives us for the case where D is (3.2.i) and u is the discourse unit provide by its second sentence 'She smiled' a rhetorical justification $R(u',u)$, where u' is the relevant unit that is supplied by the first sentence of (3.2.i). It isn't immediately clear what the rhetorical anchor is in this case. What has been said about (3.2.i) did not really address that question. According to our present set-up, however, this is also something that will be provided by the oracle function RJ. Let us make what seems to me the most reasonable assumption of what the anchor is in this case: the main clause *he saw his wife* of the first sentence of (3.2.i). Moreover, the rhetorical relation is, according to the discussion above, Narration. In other words, the output of the oracle in this case is: 'Narration(*he saw his wife*, *she smiled*)'.

The one thing that is still missing at this point is an articulation of the temporal relations determined by the different causality-related rhetorical relations. Of those relations three have been mentioned explicitly in the present section: Causal Explanation, Narration and Elaboration. As argued earlier, these relations entail the following temporal relations between the events e and e' described by the unit u rhetorically justified by the relation and its rhetorical anchor u': (i) Narration entails '$e' \prec e$; (ii) Elaboration entails '$e \subseteq e'$; (iii) what temporal relation is determined by Causal Explanation is a slightly more difficult question. The example of Causal Explanation we have discussed – (4.145.b), repeated below – is useful for illustrating what the difficulty is.

(4.145.b) John fell. Bill pushed him.

Among the scenarios that (4.145.b) might be taken to describe there are on the one hand those in which Bill gave John a push that caused him to fall down, but where the pushing and the falling are temporally separate events, separated by, say, a brief period during which John tried to recover his balance, but failed. But there are also scenarios in which Bill's pushing isn't over by the time John's falling has begun – John is still just about on his feet but already tottering. In such cases cause and effect temporally overlap. Common between all such cases is that the onset of the cause is no later than the onset of the effect. I will represent this relation between the first and the second event as $\subset\!|$. So if e is the event expressed by the second argument u of the rhetorical predication 'CaEx(u',u)' and e' the event described by u', then 'CaEx(u',u)' entails '$e \subset\!| \; e'$'.

At long last we are ready to return to the DRS construction for (3.2.i). We recall the DRS for the first sentence, given earlier as (4.141).

(4.141)

$$
\begin{array}{|c|}
\hline
t_1 \;\; e_1 \;\; T_2 \;\; e_2 \;\; s_1 \;\; t'_2 \;\; a \;\; w \;\; z \;\; s_2, t_3 \;\; v \\
\hline
t_1 < n \quad e_1 \subseteq t_1 \quad t_2 < n \quad e_2 \subseteq t_2 \quad t'_2 = \mathrm{dur}(e_2) \quad e_2 \subseteq t'_2 \\
\text{Alan}(a) \quad \text{wife'}(w,a) \quad a\text{'s bedside}(z) \\
\\
e_1\text{: open-one's-eyes'}(a) \\
e_2\text{: see'}(a,w) \\
\\
s_1\text{: PROG}(^\wedge e. \begin{array}{|c|} \hline e \\ \hline e\text{: stand-by'}(w,z) \\ \hline \end{array}\,) \\
\\
s_2\text{: PROG}(^\wedge e. \begin{array}{|c|} \hline e \\ \hline e\text{:smile'}(\underline{x}) \\ \hline \end{array}\,) \\
\\
t_3 \subseteq s_2 \\
\hline
\end{array}
$$

We already noted that for the representation construction of the second sentence of (3.2.i) our general assumption about what is delivered by the rhetorical relation oracle must come to this: Narration(*he saw his wife*, *she* *smiled*). In view of the module that spells out the consequences of rhetorical relationships to temporal relations between the described eventualities this rhetorical relationship entails the Condition '$e_2 \prec e_3$', where e_3 is the dref for the event described by 'She smiled'. This last Condition will be added to the DRS for the second sentence, and with that to the DRS for (3.2.i) as a whole.

What remains for us to spell out is a way of bringing the relevant information about rhetorical structure, as it is provided by the oracle function RJ, to bear on the construction of the second sentence of (3.2.i) and on other cases of discourse interpretation where temporal relations are determined by rhetorical relations. This is what I propose: In all cases where information about rhetorical structure has implications for temporal relations, the need for this information is to be expressed in the form of a presupposition that is triggered by the relevant clues from the sentence or clause for which the semantic representation is being constructed. One problem here is which part of parts of the syntactic structure of the sentence or clause should be made responsible for the triggering of this presupposition. Here we take advantage of the simplifying assumption made earlier that the discourse units of a discourse or text D are its full clauses. According to the assumptions we have been using all along in the Notes about syntax, all full clauses are of category S; that is, the top node of their syntactic structure is labeled 'S'. Note that when we identify the discourse units of D with the expressions that are part of D whose syntactic structure has an 'S'-labeled top node, then we do admit discourse units that contain other discourse units as parts For our present purposes this is unproblematic so long as we can distinguish for each such structure its 'main eventuality'. (For instance, in an *if*-clause-main clause sentence the main eventuality is the eventuality described by the main clause.) But that notion is easily defined on the basis of our definitions of DRS construction. (I forego the soehwta awkward but uninformative details.)

Rhetorical relation presuppositions, then, are triggered by S-nodes. But what form should they be given? I propose that we stick as closely as possible to the forms of presuppositions for which provisions have already been made in our semantic representation constructions. That is, we will assume that the presupposition asks for a rhetorical relation R and an earlier discourse unit u' by decorating representing drefs with question marks. As things stand, there is no need to put anything into the Condition Sets of the rhetorical relation presuppositions, since no constraints on presupposition resolution are seriously formalized anyway. In our stopgap treatment of rhetorical relations the resolution of rhetorical relation presuppositions is per oracle – it is in the resolution of these presuppositions that the oracle function RL enters into

DRS construction. Resolution of such a presupposition consists in consulting the oracle for the given D and u and adding its verdict in the form of some suitable DRS Condition to the non-presuppositional DRS directly following the presupposition in the given preliminary representation.

In order to determine a suitable form for this DRS Condition we need to say something that I set aside earlier when RL was introduced:What is the exact form in which it must be supposed to deliver its pronunciations. Up to now we have been assuming that the information which the pronunciations the RJ oracle makes are predication '$R(u',u)$'. But in precisely in what form should we assume that the oracle presents the arguments u and u'? Here the natural choice seems to be this: (i) the unit u for which '$R(u',u)$' is the rhetorical justification should be given in the form of its preliminary representation; and (ii) the unit u' which anchors the rhetorical justification of u should be given by its full DRS. This second assumption presupposes that the DRS for u' is recoverable at the point where u is being processed. In the light of how we have defined the construction of discourse representations thus far that looks like a potential problem in that the DRS for u' may have been dissolved when it was integrated into the DRS for the entire part of the discourse preceding u. But that problem is easily resolved, by keeping the DRSs that have so far been constructed for individual S structures in the discourse in a separate component of the discourse context, with pointers to the S structures that they have been constructed from. (Here too I forego spelling out the necessary details formally.)

For the case under discussion – the second sentence 'She smiled' of (3.2.i) – these proposals come to this. The representation of the TP node of this sentence is given in (4.150). This representation can be constructed according to the by now familiar principles. Note that the TPpt presupposition, triggered by T, is already part of this representation and that in case, where the triggering feature is pres, it is the trivial presupposition that requires the question-marked dref t_3 to be set equal to n. The difference between this TPpt presupposition and the one triggered by the second sentence of (3.2.ii) has so far only been discussed informally. I repeat here only what I said earlier: When the tense feature is past and the eventuality is a state description 'for reasons of construction', i.e. if the dref for the described state is introduced either by the feature prog or by the feature +perf, then a past TPpt is to be chosen at which the described state is to be located. When the tense feature is past and the described eventuality is an event, then the TPpt is to be identified with n and temporal location must involve the rhetorical anchor that is supplied by the rhetorical structure module of the construction algorithm, which for us reduces to the *RJ* oracle.

(4.150)

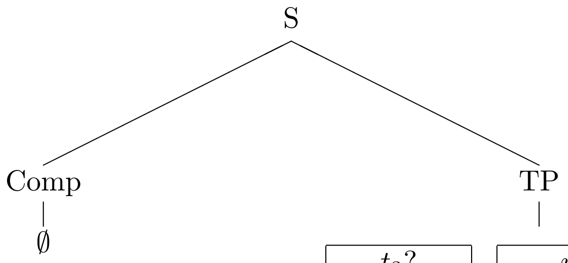

$$<e_{3,ref}, t_3, w'|<\{ \begin{array}{|c|}\hline t_3?\\\hline t_3 = n\\ \text{TPpt}(t_3)\\\hline\end{array}, \begin{array}{|c|}\hline w'?\\\hline \text{human}(w')\\ \text{female}(w')\\\hline\end{array} \}, \begin{array}{|c|}\hline \\\hline e_3\text{: smile'}(w')\\ e_3 \subseteq t_3\\\hline\end{array} >>$$

Completing the construction of the preliminary DRS leads from (4.150) to (4.151). Processing the S-node adds the rhetorical relation presupposition.

$$(4.151) <t_3, w'|<\{ \begin{array}{|c|}\hline t_3?\\\hline t_3 = n\\ \text{TPpt}(t_3)\\\hline\end{array}_{TPpt}, \begin{array}{|c|}\hline w'?\\\hline \text{human}(w')\\ \text{female}(w')\\\hline\end{array}_{3d.p.pron}, \begin{array}{|c|}\hline R?\quad u'?\\\hline \\\hline\end{array}_{rh.rel} \}, \begin{array}{|c|}\hline e_3\\\hline e_3\text{: smile'}(w')\\ e_3 \subseteq t_3\\\hline\end{array} >>$$

How are the presuppositions of this preliminary representation resolved? As we noted, the TPpt presupposition is trivial in this case, involving setting t_3 equal to n and adding this condition to the non-presuppositional DRS. The pronoun presupposition requires the discourse context in (4.141), which provides w as the antecedent for w'. Lastly, resolution of the rhetorical relation presupposition is by consultation of the oracle *RL*. We already stipulated that in this case the oracle's verdict is the predication 'Narration((4.141), (4.151))', with (4.151) the preliminary DRS for the current discourse unit and (4.141) the DRS for its rhetorical anchor. Short of a proper treatment of rhetorical relations and their representation we can record the effect of this resolution by simp[ly adding this predication to the Condition Set if the non-presuppositional DRS. Together with the resolutions of the other two presuppositions this leads to the DRS in (4.152).

(4.152)

$t_3 \;\; e_3 \;\; w'$
$t_3 = n \;\; w' = w$ Narration((4.141), (4.151)) e_3: smile'($\underline{x}$) $e_3 \subseteq t_3$

Crucial for our purposes is the temporal implication carried by the Condition 'Narration((4.141),(4.151))'. As explained, this Condition implies a relation of temporal succession between the main eventuality of (4.141), which is the seeing event represented by the dref e_2, and the main eventuality of the current discourse unit, represented by the dref e_3 of the preliminary representation (4.151): in other words, '$e_2 \prec e_3$'. The presence of 'Narration((4.141), (4.151))' in (4.152) licenses the addition of this temporal Condition, leading from (4.152) to (4.153).

(4.153)

$t_3 \;\; e_3 \;\; w'$
$t_3 = n \;\; w' = w$ Narration((4.141), (4.151)) $e_2 \prec e_3$ e_3: smile'($\underline{x}$) $e_3 \subseteq t_3$

Merging (4.153) with (4.141) then gives us the semantic representation for (3.2.i) that we were after; which concludes the treatment of this example.

Even though what has been said about rhetorical structure in this section has been extremely superficial from the perspective of someone who is seriously interested in this aspect of discourse semantics, its discussion may nevertheless have seemed quite out of proportion with the very moderate use that we ended up making of it. If I have gone into as much detail about rhetorical relations and rhetorical structure seven so, this has been to put a strong emphasis on the point that the principles governing temporal relations in a discourse like (3.2.i) are quite radically different from those that are involved in a discourse like (3.2.ii). Superficially the two discourses may look like they are on a par, and that is the way they have been initially discussed in these Notes (as well as in earlier work of my own, in which the pair consisting of (3.2.i) and (3.2.ii) and the striking contrast in meaning which they illustrate has more than once played a central role). But when one goes down into the details of where the temporal relations that interpreters consistently infer from these two discourses precisely come from, then it becomes clear that what in the one case can be treated just as a matter of form (as in (3.2.ii)) requires what are considerations of a wholly different nature in the other case (as in (3.2.i)). This by the way is not an uncommon phenomenon: Other examples are known as well of two linguistic structures looking much the same and getting the same kind of interpretation, or where they make the same kind of contribution to the interpretation of the whole. But even if such give or take between mere form and form-external sources of information is common enough, for algorithmically conceived accounts like ours such differences may nonetheless produce major headaches. What is possible, and even easy, to deal with a given machinery for the one case may be a different a vastly more challenging task for the other cases, even if at first sight there is little that seems to set the two kinds of cases apart. The difference between formally and rhetorically determined temporal relations is just one example of this, but it is a particularly striking one.

It is well to mention in this connection a point that might have been brought up a little earlier. Part of this section has been devoted to pointing out how central rhetorical structure is to discourse interpretation generally. From this perspective the consequences that certain rhetorical relationships have for temporal relations may appear as little more than a moderately interesting side effect. Eventually a theory of discourse representation such as the one of which bits have been developed in these Notes will have to incorporate an account of this aspect of discourse meaning independently of such temporal side effects. In this connection, and more specifically in relation to the DRS construction that has just been sketched for (3.2.i), the following question may have occurred to the reader: When rhetorical relation presuppositions are made part of the preliminary representation of a sentence like ‘She smiled’ from (3.2.i), shouldn’t they not be made part of other preliminary representations as well? In particular, should we not see the proposal made in this section about the treatment of the second sentence of (3.2.i) as implying the plea for a kind of global revision in this spirit of pretty much all the preliminary representations presented in Section 4? The proper answer to this question can only be a confirmative one. But the way in which the case presented by ‘She smiled’ in (3.2.i) differs from all other cases considered in Section 4 is that this has been the only case where the rhetorical relation presupposition is needed in order to construct a temporal relation between the eventualities described by successive sentences of a discourse. As long as our focus is on temporal relations, adding rhetorical relation presuppositions in these other cases will be harmless but also pointless. Given the aims we have been pursuing in these Notes, adding rhetorical relation presuppositions in the other cases would just be useless clutter.

With this we do not only come to the end of the present section, but to the end of these Notes. Quite a bit could have been added, in the first place more about the temporal and aspectual properties of discourse, but also about topics that have been no more than touched on here and there, such as information structure and propositional attitudes. But at least we have reached a conclusion of sorts to the task we set ourselves at the beginning of Section 3, when a number of examples were discussed that illustrate what was the very first impetus to DRT: the sentence-transcending effects of tenses and aspect operators. Not all of those examples have been dealt with in Section 4.4. But the tools developed in Section 4 should suffice to deal with the remaining ones too. Nothing would make me happier than when readers who have made it to this point will go beyond it by tracking the representation constructions of these examples and others of their own invention.

Bibliography

Abusch, D. (1993), 'The scope of indefinites', *Natural Language Semantics* pp. 83–135.

Alchourron, C., Gärdenfors, P. & Makinson, D. (1985), 'On the logic of theory change', *Journal of Symbolic Logic* **50**, 510–530.

Anscombe, E. (1964), 'Before and after', *The Phil Review* **73**, 3–24.

Asher, N. (1992), 'A default, truth conditional semantics for the progressive', *Linguistics and Philosophy* **15 (5)**, 463–508.

Asher, N. (1993), *Reference to Abstract Objects in Discourse*, Kluwer, Dordrecht, Dordrecht.

Asher, N. (2011), *Lexical Meaning in Context, a Web of Words*, CUP.

Asher, N. & Lascarides, A. (2003), *Logics of Conversation*, Cambridge University Press.

Asher, N. & Lascarides, A. (2007), Segmented discourse representation theory: Dynamic semantics with discourse structure, *in* H. Bunt & R. Muskens, eds, 'Computing Meaning', Kluwer Academic Publishers, pp. 141–162.

Asher, N. & Morreau, M. (1990), Commonsense entailment: A modal theory of nonmonotonic reasoning., *in* J. van Eijck, ed., 'Logics in AI, Proceedings of JELIA 1990', pp. 1–30.

Baader, F., Calvanese, D., McGuinness, D., Nardi, D. & Patel-Schneider, P. (2003), *Description Logic Handbook*, Cambridge University Press.

Bach, E. (1986), 'The algebra of events', *Linguistics and Philosophy* pp. 5–16.

Beaver, D. (1992), The kinematics of presupposition, *in* P. Dekker & M. Stokhof, eds, 'Proceedings of the Eighth Amsterdam Colloquium', ILLC, University of Amsterdam, pp. 17–36.

Beaver, D. (1997), Presupposition, *in* J. van Benthem & A. ter Meulen, eds, 'Handbook of Logic and Language', Elsevier, Amsterdam, pp. 939–1008. **URL:** *citeseer.nj.nec.com/beaver96presupposition.html*

Beaver, D. (2001), *Presupposition and Assertion in Dynamic Semantics*, CSLI.

Beaver, D. & Clark, B. (2008), *Sense and Sensitivity: How Focus Determines Meaning*, Wiley-Blackwell.

Beck, S. & Snyders, W. (2001), The resultative parameter and restitutive again, *in* C. Fery & S. W, eds, 'Audiatur Vox Sapientiae', Akademie-Verlag, Berlin.

Berman, S. R. (1991), On the Semantics and Logical form of WH-Clauses, PhD thesis, UMass, Amherst.

Blackburn, P., de Rijke, M. & Venema, Y., eds (2001), *Modal Logic*, Cambridge University Press, New York.

Bos, J. (2008), Wide-coverage semantic analysis with boxer, *in* J. Bos & R. Delmonte, eds, 'Semantics in Text Processing. STEP 2008 Conference Proceedings', Research in Computational Semantics, College Publications, Tilburg, pp. 277–286.

Brasoveanu, A. (2007), Structured Nominal and Modal Reference, PhD thesis, Rutgers Unviersity.

Bybee, J. & Dahl, O. (1989), 'The creation of tense and aspect systems in the languages of the world', *Studies in Language* **13**, 51–103.

Carlson, G. & Pelletier, F. (1995*a*), *The Generic Book*, University of Chicago Press, Chicago.

Carlson, G. & Pelletier, J. e. (1995*b*), *The Generic Book*, University of Chicago Press, Chicago.

Champollion, L. (2014), 'The interaction of compositional semantics and event semantics', *Linguistics and Philosophy* **38**, 3166.

Chierchia, G. (1991), 'Anaphora and dynamic binding', *Linguistics and Philosophy* **15**(2), 111–183.

Chierchia, G. (1995), *Dynamics of Meaning: Anaphora, Presupposition, and the Theory of Grammar*, University of Chicago Press, Chicago.

Chierchia, G. (2001), A puzzle about indefinites, *in* C. Cecchetto, G. Chierchia & M. Guasti, eds, 'Semantic Interfaces: Reference, Anaphora, and Aspect', CSLI, Stanford.

Chomsky, N. (1981), *Lectures on Government and Binding*, Foris, Dordrecht.

Chomsky, N. (1982), *Some Concepts and Consequences of the Theory of Government and Binding*, MIT, Cambridge.

Chomsky, N. (1995), *The Minimalist Program*, MIT, Cambridge.

Cocchiarella, N. (1965), Tense and Modal Logic, PhD thesis, UCLA.

Cooper, R. (1979), The interpretation of pronouns, *in* 'Syntax & Semantics, vol 10', Academic Press, pp. 61–92.

Cooper, R. (1983), *Quantification and Syntactic Theory*, Reidel, Dordrecht, Dordrecht.

Coppock, E. & Beaver, D. (2014), Definiteness and detrminacy. Manuscript, Gothenburg University/University of Texas.

Dahl, O. (1985), *Tense and aspect systems*, Blackwell, Oxford.

Davidson, D. (1967), 'Causal relations', *The Journal of Philosophy* **LXIV**, 691–703.

Davidson, D. (1970), Mental events, *in* L. Foster & J. Swanson, eds, 'Experience and Theory', Duckworth, London, pp. 81–95.

de Swart, G. & Molendijk, A. (1999), 'Negation and the temporal structure of narrative discourse', *Journal of Semantics* **16**, 1–42.

Dowty, D. (1979), *Word Meaning and Montague Grammar*, Reidel, Dordrecht, Dordrecht.

Elbourne, P. (2005), *Situations and Individuals*, MIT Press, Cambridge Mass.

Endriss, C. (2009), 'Quantificational topics – a scopal treatment of exceptional wide scope phenomena', *Studies in Linguistics and Philosophy* .

Evans, G. (1977), 'Pronouns, quantifiers and relative clauses', *Canadian Journal of Philosophy* **7**, 467–536.

Evans, G. (1980), 'Pronouns', *Linguistic Inquiry* **11**, 337–362.

Farkas, D. (2002), 'Specificity distinctions', *Journal of Semantics* **19**, 1–31.

Fodor, J. & Sag, I. (1982), 'Referential and quantificational indefinites', *Linguistics and Philosophy* .

Frege, G. (1892), Ueber sinn und bedeutung. Translated as 'On Sense and Reference' in Geach and Black (eds.), Translations from the Philosophical Writings of Gottlob Frege, Blackwell, Oxford (1952).

Geach, P. (1962), *Reference and Generality : An Examination of Some Medieval and Modern Theories*, Cornell University Press, Ithaca. Third revised edition: 1980.

Geurts, B. (1997), 'Good news about the description theory of names', *Journal of Semantics* **14**, 319–348.

Geurts, B. (1999), *Presuppositions and Pronouns*, Elsevier Science.

Geurts, B. & Maier, E. (2010), Layered drt. manuscript.

Graff-Fara, D. (2015), 'Names are predicates', *Philosophical Review* **124**, 59–117.

Groenendijk, J. & Stokhof, M. (1990), Dynamic montague grammar, *in* L. Kalman & L. Polos, eds, 'Papers from the Second Symposium on Logic and Language', Akademiai Kiadoo, Budapest, pp. 3–48.

Grosz, B., Joshi, A. K. & Weinstein, S. (1995), 'Centering: A framework for modeling the local coherence of discourse', *Computational Linguistics* **21**, 203–226.

Heim, I. (1982,1988), *The Semantics of Definite and Indefinite Noun Phrases*, Garland Press.

Heim, I. (1983), 'On the projection problem for presuppositions', *Proceedings of the West Coast Conference on Formal Linguistics* **2**, 223–248.

Heim, I. (1990), 'E-type pronouns and donkey anaphora', *Linguistics and Philosophy* **13**(2), 137–177.

Heim, I. & Kratzer, A. (1998), *Semantics in Generative Grammar*, Blackwell, Oxford.

Hinrichs, E. (1986), 'Temporal anaphora in discourses of english', *Linguistics and Philosophy* **9**(1), 63–81.

Johansson, L., Österberg, J., Sliwinski, R. & Sobel, H. (2009), *Logic, Ethics and All That Jazz: Essays in Honour of Jordan Howard Sobel*, Dept. of Philosophy, Uppsala University.

Kamp, H. (1981*a*), 'Evénements, représentations discursives et référence temporelle', *Langages* **64**, 39–64. For the original English ms. from which the paper in *Langages* is the French translation, see *Semantics and Pragmatics*, Vol 10 (2017.

Kamp, H. (1981*b*), A theory of truth and semantic representation, *in* J. Groenendijk et al., eds, 'Formal Methods in the Study of Language', Mathematisch Centrum, Amsterdam. Reprinted in: K. von Heusinger and A. Ter Meulen, 'Meaning and the Dynamics of Interpretation. Selected Papers of Hans Kamp'. Brill, 2013.

Kamp, H. (1990), Prolegomena to a structural account of belief and other attitudes, *in* C. A. Anderson & J. Owens, eds, 'Propositional Attitudes—The Role of Content in Logic, Language, and Mind', University of Chicago Press and CSLI, Stanford, chapter 2, pp. 27–90. Reprinted in: K. von Heusinger and A. Ter Meulen, 'Meaning and the Dynamics of Interpretation. Selected Papers of Hans Kamp'. Brill, 2013.

Kamp, H. (2001,2011), Articulated contexts. Manuscript.

Kamp, H. (2001*a*), The importance of presupposition, *in* C. Rohrer, A. Rossdeutscher & H. Kamp, eds, 'Linguistic Form and its Computation', CSLI, Standord. Reprinted in: K. von Heusinger and A. Ter Meulen, 'Meaning and the Dynamics of Interpretation. Selected Papers of Hans Kamp'. Brill, 2013.

Kamp, H. (2001*b*), Presupposition computation and presupposition justification: One aspect of the interpretation of multi-sentence discourse, *in* M. Bras & L. Vieu, eds, 'Pragmatic and Semantic Issues in Discourse and Dialogue', Elsevier. Reprinted in: K. von Heusinger and A. Ter Meulen, 'Meaning and the Dynamics of Interpretation. Selected Papers of Hans Kamp'. Brill, 2013.

Kamp, H. (2003), Einstellungszustände und einstellungszuschreibungen in der diskursrepräsentationtheorie, *in* U. Haas-Spohn, ed., 'Intentionalität zwischen Subjektivit"at und Weltbezug', mentis, Paderborn, pp. 209–289.

Kamp, H. (2015), 'Using proper names as intermediaries between labelled entity representations', *Erkenntnis* **80**, 263–312.

Kamp, H. (2016), Entity representations and articulated contexts. to appear in a collection in honor of David Kaplan.

Kamp, H. (2017), 'Events, discourse representations and temporal reference', *Semantics and Pragmatics* **17**.

Kamp, H. (2019), Verbal communication, mental representation and deixis: How the things we say ar connected with mand and world. in prparation.

Kamp, H. & Bende-Farkas, A. (2019), Epistemic specificity from a communication-theoretic perspective. to appear in Journal of Semantics.

Kamp, H. & Reyle, U. (1993), *From Discourse to Logic*, Kluwer.

Kamp, H. & Rohrer, C. (1983*a*), Temporal reference in french. ms. Stuttgart.

Kamp, H. & Rohrer, C. (1983*b*), Tense in texts, *in* Bäuerle, Schwarze & V. Stechow, eds, 'Meaning, Use and Interpretation of Language', De Gruyter, Berlin, pp. 250–269.

Kamp, H., van Genabith, J. & Reyle, U. (2011), Discourse representation theory: An updated survey, *in* D. Gabbay, ed., 'Handbook of Philosophical Logic, Vol. XV', Elsevier, pp. 125 – 394.

Kaplan, D. (1989), Demonstratives, *in* H. W. J. Almog, J Perry, ed., 'Themes from Kaplan', Oxford University Press.

Karttunen, L. (1973), 'Presuppositions of compound sentences', *Linguistic Inquiry* **4**, 169–193.

Karttunen, L. (1974), 'Presupposition and linguistic context', *Theoretical Linguistics* pp. 181–194.

Kratzer, A. (1989), 'An investigation of the lumps of thought', *LInguistics and Philosophy* **12**, 607–653.

Kratzer, A. (1998), Scope or pseudoscope? are there wide-scope indefinites, *in* S. Rothstein, ed., 'Events and Grammar', Kluwer, Dordrecht.

Krifka, M. (2002), Be brief and vague! and how bidirectional optimality theory allows for verbosity and precision., *in* D. Restle & S. Zaefferer, D., eds, 'Sounds and Systems. Studies in Structure and Change. A Festschrift for Theo Vennemann', Berlin: Mouton de Gruyter.

Kripke, S. (1980), *Naming and Necessity*, Harvard University Press, Cambridge. First appeared in: D. Davidson and G. Harman(eds), *Semantics of Natural Language*, Reidel, 1972.

Kripke, S. (2009), ‘Presupposition and anaphora: Remarks on the formulation of the projection problem’, *Linguistic Inquiry* **40**, 367–386.

Landman, F. (1989*a*), ‘‘Groups’, Part I, II’, *Linguistics and Philosophy* **12**, 559–605, 723–744.

Landman, F. (1989*b*), ‘The progressive’, *Natural Language Semantics* **1**, 1–32.

Landman, F. (2000), *Events and plurality: the Jerusalem lectures*, Kluwer, Dordrecht.

Langendoen, D. & Savin, H. (1971), The projection problem for presuppositions, *in* C. Filmore & D. Langendoen, eds, ‘Studies in Linguistic Semantics’, Holt, New York, pp. 55–62.

Lascarides, A. & Asher, N. (1993*a*), A semantics and pragmatics for the pluperfect, *in* ‘EACL 93’, Utrecht, pp. 250–259.

Lascarides, A. & Asher, N. (1993*b*), ‘Temporal interpretation, discourse relations, and common sense entailment’, *Linguistics and Philosophy* **16**, 437–49.

Lasersohn, P. (2003), ‘Pragmatic halos’, *Language* **75(3)**, 522â551.

Lewis, D. (1973), *Counterfactuals*, Basil Blackwell, Oxford.

Lewis, D. K. (1975), Adverbs of quantification, *in* E. Keenan, ed., ‘Formal Semantics of Natural Language’, Cambridge University Press, Cambridge.

Link, G. (1986), Generalized quantifiers and plurals, Report CSLI-86-67, Stanford University, CA.

Longobardi, A. (1994), ‘Reference and proper names: A theory of n-movement in syntax and logical form’, *Linguistic Inquiry* **25**, 609–655.

Mann, W. & Thompson, S. (1987), ‘Rhetorical structure theory: A framework for the analysis of texts’, *International Pragmatics Association Papers in Pragmatics, I, 79-105* **1**, 79–105.

Mann, W. & Thompson, S. (n.d.), Rhetorical structure theory: Description and construction of text structure, *in* 'Natural Language Generation: New Results in Artificial Intelligence', pp. 279–300.

May, R. (1977), The Grammar of Quantification, PhD thesis, MIT.

May, R. (1985), *Logical Form*, MIT Press, Cambridge Mass.

Moens, M. & Steedman, M. (1988), Temporal ontology in natural language, *in* 'Proceedings 25-th annual meeting of the Association for Computational Linguistics', pp. 7–14.

Neale, S. (1990), *Descriptions*, Bradford Books, MIT Press.

Nouwen, R. (2015), Plurality, *in* M. Aloni & P. Dekker, eds, 'Cambridge Handbook of Semantics', Cambridge University Press, Cambridge, U.K.

Parsons, T. (1990), *Events in the Semantics of English*, The MIT Press, Cambridge, MA.

Partee, B. (1973), 'Some structural analogies between tenses and pronouns in english', *Journal of Philosophy* **70**, 601–9.

Partee, B. (1984), 'Nominal and temporal anaphora', *Linguistics and Philosophy* **7**.

Partee, B. H. (1972), Noun phrase interpretation and type-shifting principles, *in* J. Groenendijk, D. de Jong & M. Stokhof, eds, 'Studies in Discourse Representation Theory and the Theory of Generalized Quantifiers', D. Reidel, Dordrecht, Holland, p. 115â143.

Partee, B. H. & Rooth, M. (1983), Generalized conjunction and type ambiguity, *in* R. Bauerle, C. Schwarze & A. von Stechow, eds, 'Meaning, Use and Interpretation of Language', Walter de Gruyter, Berlin, West Germany.

Peters, S. & Westerstahl, D. (2006), *Quantifiers in language and logic*, OUP, New York.

Portner, P. (2003), 'The (temporal) semantics and the (modal) pragmatics of the present perfect', *Linguistics and Pilosophy* pp. 459–510.

Prior, A. (1967), *Past, Present and Future*, Clarendon Press, Oxford.

Reichenbach, H. (1947), *Elements of Symbolic Logic*, Macmillan, London.

Reyle, U. (1993), 'Dealing with ambiguities by underspecification: Construction, representation and deduction', *Journal of Semantics* **10**, 123–179.

Reyle, U., Rossdeutscher, A. & Kamp, H. (2007), 'Ups and downs in the theory of temporal reference', *Linguistics and Philosophy* **30**(5), 565–635.

Roberts, C. (1987), Modal Subordination, Anaphora, and Distributivity, PhD thesis, University of Massaschusetts.

Roberts, C. (1989), 'Modal subordination and pronominal anaphora in discourse', *Linguistics and Philosophy* **12**, 683–721.

Rooth, M. (1985), Association with Focus, PhD thesis, UMass, Amherstt.

Rooth, M. (1992), 'A theory of focus interpretation', *Natural Language Semantics* **1**, 75–116.

Rossdeutscher, A. (2013), Denominal spatial verbs revisted, *in* A. Rossdeutscher, ed., 'Sub-lexical investigations: German particles, prefixes and prepositions', Vol. 11 of *SinSpec (Working Papers of the SFB 732)*, University of Stuttgart, pp. 58–88.

Russell, B. (1905), 'On denoting', *Mind* **14**, 479–493.

Scha, R. (1981), Distributive, collective and cumulative quantification, *in* J. e. a. Groenendijk, ed., 'Formal Methods in the Study of Language', Mathematisch Centrum, Amsterdam, pp. 483–512.

Schwarz, B. (2001), Two kinds of long distance indefinites, *in* R. van Rooy & m. Stokhof, eds, 'Proceedings of the Thirteenth Amsterdam Colloquium, Amsterdam University', University of Amsterdam, Amsterdam, pp. 192–197.

Schwarzschild, R. (1996), *Pluralities*, Kluwer.

Smith, C. (1991), *The Parameter of Aspect*, Kluwer.

Staab, S. & R.Studer (2003), *Handbook on Ontology*, Springer.

Stalnaker, R. (1968), A theory of conditionals, *in* N. Rescher, ed., 'Studies in Logical Theory', American Philosophical Quarterly Monograph Series 2, Basil Blackwell, Oxford.

Stalnaker, R. (1972), Pragmatics, *in* D. Davidson & G. Harman, eds, 'Semantics of Natural Language', Reidel, pp. 380–397.

Stalnaker, R. (1974), Pragmatic presuppositions, *in* M. Munitz & P. Unger, eds, 'Semantics and Philosophy', New York University Press,.

Stalnaker, R. (1979), Assertion, *in* P. Cole, ed., 'Syntax and Semantics Vol.9: Pragmatics', Academic Press, New York, pp. 315–332.

Strawson, P. (1950), 'On referring', *Mind* **59**, 320–344.

Thomason, R. (2002), Combinations of tense and modality, *in* D. Gabbay, ed., 'Handbook of Philosophical Logic VII', Kluwer, Dordrecht.

T.Reinhart (1997), 'Quantifier scope: How labor is divided between qr and choice functions', *Linguistics and Philosophy* **20**, 335–397.

Van Der Sandt, R. (1992), 'Presupposition projection as anaphora resolution', *Journal of Semantics* **9**, 333–377. Special Issue: Presupposition, Part 2.

Van Der Sandt, R. & Geurts, B. (1991), Presupposition, anaphora, and lexical content, *in* O. Herzog & C.-R. Rollinger, eds, 'Text Understanding in LILOG', Springer.

Vendler, Z. (1967), Verbs and times, *in* Vendler, ed., 'Linguistics in Philosophy', Cornell University Press, pp. 97–121.

Vlach, F. (1992), The semantics of the progressive, *in* P. Tedeschi & A. Zaenen, eds, 'Syntax and Semantics. Vol. 14: Tense and Aspect', Academic Press, pp. 271–292.

Von Fintel, K. (2004), Would you believe it? the king of france is back! presuppositions and truth-value intuitions., *in* M. Reiner & A. Bezuidenhout, eds, 'Descriptions and Beyond', OUP.

Walker, M. A., Joshi, A. & Prince, E. (1998), *Centering Theory in Discourse*, Clarendon Press, Oxford.

Webber, B. (1988), 'Tense as discourse anaphor', *Computational Linguistics* **14**(2), 61–73.

Winter, Y. (1997), 'Choice functions and the scopal semantics of indefinites', *Linguistic and Philosophy* **20**, 399–467.

Zanardo, A. (1985), 'A finite axiomatization of the set of strongly valid ockamist formulae', *Linguistics* **14**, 447–468.

Zeevat, H. (1989), 'A compositional approach to discourse representation theory', *Linguistics and Philosophy* **12**, 95–131.

Zweig, E. (2008), Dependent Plurals and Plural Meaning, PhD thesis, New York University.