Entity Representations and Articulated Contexts

An Exploration of the Semantics and Pragmatics of Definite Noun Phrases

Hans Kamp
University of Stuttgart
University of Texas at Austin*

*This paper has gone through several phases and it has taken its time to do so. A first version was presented at the Congrès International des Linguistes, Prague, 2003. Since then versions of the paper have been discussed on various occasions, most recently at a workshop on reference that took place at the Centre for Mind and Nature of Oslo University in June of 2015 and in a lecture at ESSLLI 2015 in Barcelona. Thanks to all audiences who have helped me to focus on various aspects of the proposals I will be making below. I know that their remarks have led to improvements in what follows; if they do not like what they see, they should know that things would have been worse, hadn't been for them. The greatest help I have had in getting this document into the shape that it is in at the moment has come from my colleagues at the IMS of the University of Stuttgart. I would like to thank in particular Arndt Riester, who has made use of some of the main ideas that are developed below in his own work on the classification of definite noun phrase occurrences in texts and in the annotation schemes he has developed on the basis of that work – see Riester (2008), which gives an account of Riester's use and elaboration of the ideas of the present paper. Arndt's comments have been especially helpful. Others who have helped improve the ms. significantly with corrections and suggestions are Gemma Boleda and Casey Woolwine. Charles Yee converted an earlier draft of this paper into Latech and thereby put me on the right foot for the corrections that I made to the ms. after that. Part of the work that led to this document was carried out while the author was a member of the research group 'Meaning and understanding across languages' (director: C. Fabricius-Hansen) which met during the academic year 2010-2011 at the Center for Advanced Study of the Norwegian Academy of Science. [For Prepublication only: The paper is planned to appear as part of a tribute to David Kaplan, with the preliminary title: 'Sense, Reference and Use– Afterthoughts on Kaplan'.]

For David Kaplan,
in admiration and
gratitude

1. Introduction

1.1 Motivation

The starting point for this paper was an attempt to develop a notion of context which encompasses all information that may be needed in the interpretation of definite noun phrases.[1] Such contexts must include information of various kinds and coming from different sources. Since the interpretation of some NP types can make use of some parts of this information but not of other parts, while other information can be used to interpret other types[2], the different kinds of information must be kept separate, by the human interpreter and therefore also by a theory that is to tell a credible story about how human interpreters deal with definite NPs. I will assume in what follows that these different kinds of information are being kept separate by virtue of being allocated to separate 'contexts'. However, the separation between these contexts is not absolute. There are interactions between them. In particular, certain interpretational moves lead to the transfer of information from one 'context' to another. Because of this the different 'contexts' are best treated as distinct but interacting components of a single overarching context. I will refer to such overarching contexts as *Articulated Contexts*. An Articulated Context can be thought of as a kind of articulated lorry – a vehicle composed of several loosely connected parts, each with its own content and its own measure of kinetic independence, but always moving in tandem, and with the occasional reloading.

Some of the components of Articulated Contexts are made up of *Entity Representations*. The notion of an Entity Representation that I will be using is one of the ingredients of MSDRT (Mental State Discourse Representation Theory), a DRT-based theory of the structure of mental states that was developed originally as a framework for the semantics of simple and complex attitude

[1] I am using the old-fashioned term "noun phrase" (or "NP") to refer to what nowadays (since Abney (1987)) is more often referred to as "D(eterminer) P(hrase)"

[2] For instance, information that can be used in the interpretation of definite descriptions is not accessible to the interpretation of third person pronouns.

reports. According to MSDRT, Entity Representations are constituents of mental states. This entails that adopting them as constituents of Articulated Contexts introduces a psychological dimension to the notion of an Articulated Context. This then more or less inevitably leads to a conceptualization of Articulated Contexts as psychological constructs in their entirety and of interpretations that make use of information from Articulated Contexts as psychological processes. From the anti-psychologistic perspective that has long dominated formal semantics such a turn towards the mentalistic cannot but seem suspect. For those of such an anti-psychologistic persuasion there is little I can do to assuage their misgivings or objections. All I can offer is a partial reassurance: Towards the end of the paper it is argued that the proposals I will have made preserve much of the user-independent way of analyzing natural language meaning that the true anti-psychologist considers non-negotiable.) However – this is perhaps the most central message of this paper – there are some aspects of natural language semantics that cannot be accounted for unless the mental is brought into play. In particular, when it comes to NP reference, non-mentalistic accounts show their limits.

The turn I will propose is not just a turn from a user-independent semantics to a theory of text and utterance interpretation in which the mental states of the language users play a prominent part. It is, more than that, the turn to a full-blown communication-theoretic perspective, in which speech acts are analyzed as vehicles for the transfer of thoughts from producer to recipient. More specifically – since the focus of this paper is on the semantics of definite noun phrases – we will be concerned not only with the *interpretation* of NPs but also with the *NP choices* that speech producers make when they put their thoughts into words. How producers choose their NPs depends on the interpretational possibilities they attribute to their audience. So something needs to be said about speaker-hearer coordination and related notions like Common Ground. The final parts of the paper will touch on these matters.

1.2 How the paper is set up, and why

The first part of the paper, consisting of Section 2, deals with an issue that is preliminary to our development of the notion of an Articulated Context that follows in Section 3. It presents a way of unifying two notions of context that have been prominent in the semantics literature and for which detailed,

formally precise definitions have been given. These are: (i) the notion of an *utterance context* as it is understood in the work of, most notably, Montague, Kaplan, Lewis and Creswell[3]; and (ii) the notion of a *discourse context*, as it is found in Dynamic Semantics[4]. The first part differs from the second part in conforming to a widely if mostly tacitly accepted methodological principle of formal semantics: the semantic properties of natural languages are properties of autonomous formal systems, whose existence and structure is independent of the uses that can be made of them, and whose properties can therefore be described without any reference to the psychology of their users.

In fact, this first part could in principle be detached from the longer second part. But I am not sure that that would be a good thing. From my own perspective the first part is mostly of interest as a first step towards the notion of an Articulated Context, and therewith as a prelude to a theory of reference in which mental representation is essential. A further reason for not wanting to turn it into a document of its own is that in my own perception it suffers from a certain pedestrian ponderousness. I am aware that the going is slow and the surprise rate is low. But I haven't found the right way to quicken or lighten the pace. So long as this is part one of something longer, and as a kind off preliminary to what follows, readers may be willing to carry on, hoping that things will brighten up once they have left prt one behind them.

One reservation about the present document that has been conveyed to me is that it is rather short on examples, which make visible what its theoretical proposals come to when they are applied to particular cases. I take the point. I myself see the example shortage in particular as a problem for Section 3. The reader of this Section may well feel a growing need for concrete illustrations of proposals that are often stated in quite general and abstract terms. Unfortunately, worked-out examples don't come for free. More often than not they require close attention to details that have nothing to do with the what they are intended to illustrate, but that need to be addressed nonetheless. Spelling these details out with sufficient precision tends to take up considerable extra space; and, worse, it has a tendency to detract from the matter at issue. If I have been sparing with examples, than this has been primarily for these connected reasons.

[3]Montague (1970), Kaplan (1989), Lewis (1970),Cresswell (1973)

[4]See among others Heim (1982,1988), Kamp (1981b)

Almost all the examples in Section 3 are examples of deictic uses of demonstrative NPs (i.e. of NPs that in the singular begin with *this* or *that*) and in the plural with *these* or *those*). One reason for focussing on examples of this kind is that it is here that the proposals of the present paper differ most markedly from Kaplan's theory of Demonstratives, developed in Kaplan (1989) and related work. For Kaplan deictic uses of demonstratives and standard uses of indexicals like *I*, *you*, *now* etc are all instances of what is a single phenomenon at heart: that of referring directly through exploitation of the utterance context. On the account proposed in Section 3 below the similarity is much more tenuous. According to this account the interpretation of deictically used demonstrative NPs is unique in that it typically involves two distinct channels of access to the NP's referent, one as the referent of the NP used by the speaker and one via the interpreter's direct perception. The processes involved are quite different from what we find with the typical interpretations of indexicals (or, for that matter, with any other uses of definite NPs.

The story of how the two access channels interact in the interpretation of deictic demonstratives I see as one the most interesting spin-offs from the communication-theoretic approach adopted in Section 3. It is also a story that as far as I can see couldn't be told without making use of the Entity Representations that are a central feature of MSDRT and most of its applications. Detailed examples of deictic uses of demonstrative NPs are therefore also useful because they help to appreciate the roles that Entity Representations play in the present approach to definite Noun Phrase interpretation.

None of this is a justification for skimping on examples of other types. But it should explain why examples of deictically used demonstratives came to be seen as a having first priority.

2. The Unification of Utterance Context and Discourse Context

In this section it is shown how utterance context and discourse context can be united into a single formally explicit notion of context. The section starts

with an informal summary in Section 2.1 of the roles that utterance context and discourse context play in accounts of indexicals, anaphoric pronouns and other expressions; it also reminds us that sometimes utterance context and discourse context are needed at the same time. Sections 2.2 and 2.3 give condensed presentations of the notion of utterance context that can be found in the formal semantics literature[5] (Section 2.2) and the notion of discourse context that is an integral part of Discourse Representation Theory (Section 2.3).

The remaining parts of Section 2 are devoted to the unification of these two notions. Section 2.4 shows that as far as time is concerned a combination of utterance context and discourse context was de facto already achieved in the earliest versions of DRT (e.g. in Kamp (1981a) and in Kamp and Reyle (1993), Ch. 5). Section 2.5 shows, by way of an example, how the approach of which Section 2.4 reminds can be extended to components of the utterance context other than the utterance time. Furthermore, this section introduces the notion of a singular proposition (needed in connection with all utterance context components) and that of a partial proposition (needed in particular when components such as the speaker and the addressee are brought into play).

Section 2.6 deals with the last preliminary — the definition of the *concept of an utterance context* (first introduced in Section 2.2), as distinct from the individual utterance contexts that had been the primary topic of discussion so far. This definition is needed for the definitive presentation of the merge of utterance context and discourse context that is given in Section 2.7. Section 2.7 can be seen as the central part of Section 2. It is here that the unification of utterance context and discourse context receives its general formal definition.

Sections 2.8 and 2.9 wind up Section 2 with some observations that will be relevant to the discussions of Section 3. Section 2.8 elaborates on and qualifies the impressionistic observation made above — that utterance contexts are static and discourse contexts are dynamic. Section 2.9 discusses the difference between the indexical discourse referents (representing utterance time, speaker and other component of the utterance context) that are central

[5]See in particular Kaplan (1989)

to the unification of utterance and discourse context that is proposed in this paper and the strictly mentalistic indexical discourse referents of MSDRT, such as the discourse referent i, which represents the self in MSDRT's descriptions of mental representations.

2.1 Utterance Context and Discourse Context

Intuitively, utterance context and discourse context correspond to different kinds of contextual information. The utterance context captures information pertaining to the utterance as individual act and to the circumstances in which that act is performed. Core examples of such information are: the time at which the utterance is made, the place in which it is made, the agent who made it and the one or ones to whom it is addressed. For some (including Kaplan) the utterance context will also include certain kinds of information about the locale in which the utterance is made, e.g. what entities from the local environment are perceptually accessible to the discourse participants. All such information is, you might say, *content-external*: it is what it is, irrespective of what content is expressed by the given utterance and by the larger discourse of which the utterance is part.

The information captured by the discourse content is the opposite of this. It is *provided by* the content of the discourse. More precisely, it is given by the content of the antecedent part of the discourse – that part which precedes the utterance in question. The information provided by the discourse context is thus (by definition, so to speak) *content-internal*. ('Content-external context', or simply 'external context', and '(content-)internal context' might have been good names for utterance and discourse context; but I don't know that anyone ever used this terminology and we won't either.)
Not only are utterance context and discourse context meant to account for different phenomena, they also have quite different properties. Most notably, discourse contexts are *dynamic* in a way that utterance contexts are not: the discourse context changes as interpretation of the discourse progresses, growing with each new sentence that adds information to what the discourse is telling. Utterance contexts do not change in this way.

The large majority of studies known to me that deal with context-dependent aspects of interpretation focus on one of these two context types while ignoring the other. That may be acceptable so long as it is the avowed aim of the

study to investigate the properties of just one kind of context dependence. But it won't do when the interpretation of an utterance depends on utterance context and discourse context alike. And that is an extremely common situation: Many utterances contain both expressions that depend on the utterance context and expressions that depend on the discourse context. In such cases the utterance as a whole will evidently depend on both utterance context and discourse context. This means that if we want a notion of context which supplies all the information that the interpretation of such utterances requires, then contexts of this new sort must subsume, in some way or another, the utterance context on the one hand and the discourse context on the other.

In fact, the need for such a more comprehensive notion of context is even more urgent than this consideration implies. There exist simple expressions – expressions that are not built in morphologically tractable ways from smaller parts – whose semantics involves both utterance context and discourse context. This may be either because some occurrences of the expression depend on the utterance context and others on the discourse context, or – this is the most telling case for us – because single occurrences of the expression involve both context types at once.

Here are some prominent examples:

i. Third person pronouns. Much attention has been paid to the ways in which the interpretation of third person pronouns can depend on the discourse context. This is true in particular of discussions of pronouns within dynamic semantics: In many early presentations of dynamic systems third person pronouns are the paradigms of anaphoric utterance constituents, for whose interpretation the discourse context is needed. But pronouns also have a deictic use; when used that way they behave like demonstratives in the sense of Kaplan, and do not depend on the discourse context, but rather on the utterance context. Deictic uses of third person pronouns are found in conversation rather than in written texts. But pronoun occurrences in conversation can be anaphoric as well as deictic. Such occurrences are thus in principle ambiguous between a deictic and an anaphoric interpretation. To my knowledge the literature has been silent on the question what principles govern the resolution of this kind of ambiguity. The more comprehensive notion of context developed here makes it easier to address this question.

ii. Demonstrative phrases. The literature on demonstrative noun phrases — *this, that, this book, that bird perched on the roof of the house over there* — shows a bias opposite to the one we find in most discussions of third person pronouns: almost all discussions of demonstratives focus on their deictic use. But demonstrative phrases can also get their interpretation from the discourse context, even if such uses are comparatively rare. An example that illustrates this:

(1) If one Texan steals the cattle of another Texan, then that Texan will be very cross.

Here *that Texan* is anaphoric to *another Texan* and thus depends for its interpretation on the (local) discourse context provided by the *if*-clause.)[6]

iii. A third type of noun phrase whose interpretations can involve the utterance context as well as the discourse context are definite descriptions — phrases that begin with the definite article *the*. The history of theorizing about definite descriptions is quite different from that pertaining to the expressions mentioned so far and, at least within formal semantics and logic, it goes back a good deal farther in time. Early treatments of definite descriptions, including in particular that of Russell's *Theory of Descriptions* and the one suggested by Frege to which it was offered as an alternative, largely ignore dependence on context. More recent accounts do pay attention to context dependence, but differ in the kind of context they focus on. Strawson (and perhaps others before him) noted the role of the utterance context. (For instance, a use of the phrase *the table* can succeed in referring to the unique table in the room because the discourse participants are in the room and the discourse is about what is and happens inside the room.) Dependence of definite descriptions on discourse context has become a prominent concern

[6]Note that the indefinite NP *another Texan*, although not anaphoric as a whole, also contains an anaphoric element in the adjective *other*. *other* is anaphoric to the NP *one Texan* in the sense that it restricts the extension of the occurrence of the noun *Texan* that follows it to entities that are different from the semantic value of its anaphoric antecedent *one Texan*. This is an example of an anaphoric link between two noun phrases that does not amount to identity. In this regard *other* differs from the paradigms of anaphoricitty, the third person singular personal pronouns, which are always coreferential with their antecedents, and also from non-pronominal anaphoric noun phrases like the *that Texan* of (1), which is coreferential with its anaphoric antecedent *another Texan*.

within Dynamic Semantics, especially in the work of Heim and that of Van Der Sandt. (Heim (1983), Van Der Sandt (1992), Van Der Sandt and Geurts (1991)).

Not only can the interpretation of a definite description depend on the utterance context or the discourse context, complex definite descriptions will often depend on both of them at the same time because they contain one constituent that depends on the utterance context and another that depends on the discourse context. An example, randomly chosen, is the description (*the one thing that he and I have in common*, where *I* depends on the utterance context, while *he* may require an anaphoric interpretation.[7]

iv. Fourth, and most important of all, there are single expressions whose interpretation typically depend on utterance context and discourse context at once, in the sense that the interpretation of single expression *tokens* involves both utterance context and discourse context. Prime examples of this are the tenses of the verb. Take the simple past tense in English. The point is best made by looking at an example. Consider the following two-sentence report on some past episode:

(2) Mary went to the doctor. She was ill.

Focus on the past tense of *was* of the second sentence. On the one hand this tense relates the state described by this sentence – that of Mary being ill – in a certain way to the utterance context: the state must have held at some time or times preceding the utterance time. (In this the past tense is obviously different from the present tense. Note for instance what happens when the *was* of the second sentence is replaced by *is*.) On the other hand, the only natural interpretation of (2) is one according to which the state of Mary being ill held at the time when she went to the doctor (whenever that

[7] We note in passing that definite descriptions also give rise to yet a further issue relating to context. So-called 'discourse-new' occurrences of definite descriptions (Gundel et al. (1993)) depend for their interpretation on contextual information that is neither part of the utterance context nor of the discourse context. (See Fraurud (1990), Poesio and Vieira (1998), Spenader (2002) among many others.) This is one reason why contexts which accommodate all the different types of contextual information that may be needed for the interpretation of the various possible uses of definite NPs must be more than mere combinations of utterance contexts and discourse contexts. Discourse-new definites were a crucial motivation the motivation for the more encompassing notion of an Articulated Context that will be developed in Section 3.

may have been). That is, the past tense of *was* is sensitive to the discourse context, which is provided by the first sentence of (2). A proper semantic treatment of tenses therefore requires simultaneous consideration of both utterance context and discourse context. Discourse Representation Theory – the theory that will serve as general setting for the proposals of this paper – later on as the core of MSDRT, but for now in the earlier form presented in, for instance, Kamp and Reyle (1993), Beaver et al. (2007, 2015) or Kamp and Reyle (2011) – adopted more or less from the start a way of building the information provided by the utterance context that is needed for the interpretation of tenses and certain other temporal expressions – viz. the utterance time – into its representations of discourse contexts. This way will be our point of departure for merging utterance content information and discourse context information later on in Section 2.

Noun phrases that involve both utterance context and discourse context at the same time are not all that easy to find, but there are some. One example is the plural first person personal pronoun *we*. Standard uses of *we* involve the utterance context in that the set of individuals an occurrence of *we* denotes must always contain contain the speaker. But the other members of the set can be determined in different ways, and often it is the discourse context that is needed to determine which they are. A well known example is the following mini-discourse due to Partee:

(3) When John comes alone, we play duets. But when he brings a cellist, we play trios.

This paper does not look at plurals and examples of type (iv) will not be explicitly considered.

2.2 Utterance Contexts

The motivation behind the notion of utterance context is explained with exemplary clarity in the work of Kaplan (1989). Certain expressions get their semantic values from the context in which they are used; and, once assigned, these values show a remarkable robustness – or 'rigidity', as philosophical terminology has it. This robustness is most easily demonstrated for one type of expressions that get their values from the utterance context, the *indexicals*. Some indexicals are NPs, viz. the first and second person pronouns, and for these it is perhaps easiest to explain what is meant by saying that they are

indexicals. Consider the first person singular pronoun *I*. The referential robustness of *I* can be demonstrated by comparing *I* with definite descriptions like *the speaker* or *the speaker of this utterance*, which are coreferential with *I* in simple sentences, as shown by those in (4).[8]

(4) a. I am hungry.
b. The speaker (of this utterance) is hungry.

However, the intersubstitutability *salva veritate* of the first person singular personal pronoun *I/me/my* and *the speaker (of this very utterance)* fails for occurrences in more complex sentences, in which the NP occurs within the scope of a modal or intensional operator. An example is the pair of sentences in (5).

(5) a. If Estelle had spoken these (last) words, you would have done everything you could to please the speaker.
b. If Estelle had spoken these (last) words, you would have done everything you could to please me.[9]

[8](15b) would be an odd way of referring to oneself. We normally do not refer to ourselves in the third person. Nevertheless, and this is the point at issue here, its truth conditions are the same as those of (15a): the same person has to be hungry for the sentence to be true.

[9]Some of the examples in the literature, including some of Kaplan's own, involve a kind of self reference at the level of the utterance: the utterance speaks of a counterfactual situation in which it itself is made by some other speaker than the one who actually did make it. Such examples raise some questions about utterance identity that, I feel, unnecessarily complicate the point that the examples are intended to make about the nature of indexicals like *I*. As far as I can see, examples such as (5a) and (5b), in which self-reference plays no part – 'these words' can be interpreted as referring to the last thing the speaker said before she uttered (5a) or (5b), and adding 'last' makes this fully explicit – are just as suited to make the point that indexicals are immune to shifts like those produced by the antecedents of counterfactual conditionals and thus obey semantic principles that are crucially different from those governing the definite descriptions with which they are intersubstitutable *salva veritate* in non-embedded positions.
One of the aims of Kaplan's self-referential counterfactuals was to refute Reichenbach (1947)'s account of indexical expressions as 'token-reflexive'. According to this account – or at any rate, according to it on the most plausible way of taking what Reichenbach says – it claims that words like *I* and *you* are 'token-reflexive' in the sense that their meanings can be given by the expressions *the speaker of this utterance* and *the addressee of this utterance*, where the phrase *this utterance* is meant to refer to the very utterance as part of which the phrase occurs. Examples that show the untenability of this claim should take the form of pairs of sentences that differ only in that one of them has, say, *I* or *me* in a

The intended setting for the utterance in (5a) is that of a speech act in which the speaker is referring to something she has just said to the addressee, but without the impact she had hoped for. And this prompts her to vent her frustration over the addressee's devotion to Estelle: Estelle only needs to say a word to the addressee and he will try his best to please her. But clearly (5b), in which *the speaker* has been replaced by *me*, cannot be understood as expressing the same content. It says that if Estelle had spoken then the addressee would have done everything he could to please the actual speaker, not Estelle. The reason for the difference between (5a) and (5b) is obvious: whereas *the speaker* in (5a) is naturally understood as referring to the person who has been speaking 'those words' in the situation described by the *if*-clause of the sentence (viz. Estelle), the NP replacing it in (5b), i.e. the non-nominative form *me* of the first person pronoun, can only be interpreted as referring to the actual user of *me* and thus to the *actual* speaker of those words.

Examples like those in (5) show that the mechanisms involved in the interpretation of indexicals cannot be the very same as those involved in the

---

position where the other has *the speaker of this utterance* and where the constituent *this utterance* of the second phrase is given the 'self-referential' interpretation just described. Sentence pairs like those in (5) don't quite fill that bill, but (1) is a pair that does.

(1) a. If Estelle had spoken these very words, you would have been paying the closest possible attention to the speaker of this utterance.

b. If Estelle had spoken these very words, you would have been paying the closest possible attention to me.

The pair of sentences in (1) provide a clear counterexample to the specific token-reflexive analysis that Reichenbach offered for the word *I*, and not only to the claim that the semantics of *I* differs from that of certain definite descriptions with which it corefers in extensional positions. But the second point is a different one from the first; as (5) shows, it can be illustrated also by comparing *I* with descriptions that corefer with it in extensional positions but are not token-reflexive in Reichenbach's specific sense. (Note by the way that it is obviously impossible to demonstrate by example that *I* differs semantically from all coreferring definite descriptions, since there is an open-ended, and in principle infinite, number of descriptions that might be offered as semantic equivalents, and there is of course no hope of going through all those one by one. That there won't be any generally intersubstitutable definite descriptions can only be established by providing (i) a theory of how *I* works, (ii) providing a theory of how definite descriptions work and (iii) deriving from the combination of these two theories that the reference conditions of *I* differ from those of any definite description.)

interpretation of definite descriptions such as *the speaker* etc. The reference mechanisms for indexicals are insensitive to the situational shifts that can be effected by the antecedents of subjunctive conditionals, whereas the mechanisms applicable to definite descriptions can be sensitive to such shifts. An obvious and natural explanation of this is that indexicals like *I/me* get their referents directly from the situations in which the utterances containing them are made. This renders them insensitive to situation shifts like those that are effected by, for instance, the antecedents of conditionals.

Another way to look at the difference between *I/me* and *the speaker* is by considering the truth values to which simple sentences containing indexicals like *I* evaluate in different possible situations, or different *possible worlds*, including the non-actual situations or worlds which have to be considered in the evaluation of counterfactual conditionals. For instance, take a particular utterance of (15a), the sentence "I am hungry." Possible situations may differ from each other in that different people are hungry. So the extension of the word *hungry* - the set of individuals that are hungry - will vary from one situation to the next. As a consequence, the truth value of the given utterance of (15a) may also vary from one situation to the next — in one situation the referent of *I* will belong to the extension of *hungry* and the sentence will be true while in another the referent will not belong to the extension and the sentence will be false. But the referent of *I* will be the same no matter which situation we take; it will always be the actual speaker of the actual utterance. In other words, the actual speaker will be a fixed constituent of the *proposition* expressed by the uttered sentence: the proposition is that function from possible situations or worlds to truth values which assigns to any possible situation or world the truth value 'true' iff the actual speaker belongs to the extension of *hungry* in that situation or world and the value 'false' otherwise.

Such considerations naturally lead to Kaplan's "three level" semantics, in which the meanings of sentences and other expressions are given by "two-tier" functions, which are first applied to "utterance contexts" — these fix, among other things, the referents of the indexicals occurring in an uttered sentence — and then to particular worlds or situations which determine the extensions of predicates and non-indexical NPs in the sentence and, via those extensions, a truth value for the sentence as a whole. Kaplan calls these two-tier functions *characters*, and the results of applying them to utterance contexts *intensions*. *Extensions* are the result of applying intensions to a

situation or world.[10]. (In case the expression is a sentence, its intensions are propositions in the sense defined above, and its extensions are truth values.) Indexicals are unlike most other expressions, exemplified by predicate words like *hungry*, in that their semantic values depend on the utterance context; they are also unlike words such as *hungry* in that the values that utterance contexts assign to them are independent of the possible situation or world of evaluation. In other words, their intensions are *constant* functions. In particular, the intension of *I* in a given utterance context c is the constant function that maps each situation or world of evaluation to the speaker of c.

The general architecture of the semantic theory that we will be assuming in this paper is not quite like the three level architecture proposed by Kaplan. But we adopt his notion of utterance context (which for present purposes doesn't differ in any essential respects from the utterance contexts one finds in the work of Montague, Lewis, Cresswell and others)[11] and also the principle that it is the utterance context which fixes the reference of indexicals like *I*. It is part of that principle that the referent does not depend on any other factors. This entails that the referent of *I* remains constant throughout the range of the situations or worlds that provide the second arguments of the two-tier character functions.

One of the tasks of a theory of indexical reference of a language L is to specify which expressions are the indexicals of L. For all I know, there are no reliably complete specifications of the full spectrum of indexical expressions for any natural language – not even for English, the language for which indexicality must have been studied more deeply and intensively than it has for any other. But it has been widely assumed that apart from the pronouns *I* and *you* the list includes also the adverbs *now* and *here*, and it is often assumed as well that the list of temporal indexicals includes besides *now* the adverbs *today*, *tomorrow*, *yesterday*, *nowadays* and also compounds like *last week*, *next week*, *three days ago* and a range of other compound phrases. Whether any of these

[10]In Montague (1970) a similar three-way distinction is drawn. Montague uses for the functions themselves simply the term "meaning", but otherwise his terminology is the same.

[11]More accurately, what we assume is a more restricted notion of utterance context than is found in Kaplan's work, one that is suitable for the semantics of indexicals, but not for that of demonstrative NPs (in the traditional morphological sense of the term, that of noun phrases which begin with *this* or *that*).

expressions strictly conform to the criteria for indexicality has been a topic for debate almost from the time when Kaplan's proposal became known to a wider public, and by now there is, as far as I can tell, a growing perception that few if any of them behave fully in the way in which a perfect Kaplanian indexical ought to behave. This is perhaps most obvious for the presumed temporal indexicals, for which *now* is often taken as a paradigm. *now* can refer to the utterance time (the time provided by the utterance context), but in story telling it often doesn't, referring instead to some other time – to the psychological 'now' of some protagonist of the narrative, or, more abstractly, to a temporal perspective time determined by the stage which the narrative has reached at the point where the given instance of *now* occurs in it. Much the same is true for the other time-related candidates of our list, though there can be subtle differences between individual items (Kamp and Rohrer (1983)). *Here* shows similar deviations from strictly indexical behaviour. As far as English is concerned, the pronouns *I* and *you* appear to fit the standard closely, although even they have properties that an indexical would not be expected to have (Nunberg (1993), Nunberg (1999)).[12] And there are other languages in which the behaviour of the first and second person pronouns does not match the largely indexical behaviour of English *I* and *you* (see in particular Schlenker (2003)).

Once it has been recognized that a purported indexical does not quite live up to the strict indexicality canon, questions arise as to exactly when and how it deviates. There are many questions here. But that does not alter the

[12] According to the exceptionally perceptive account of indexicality that can be found in the cited papers by Nunberg, all indexicals, including the robust English indexicals *I* and *you*, allow for interpretations which take them as shorthands for certain definite descriptions. These descriptions must be uniquely instantiated by the referent of the indexical in the actual situation, but may pick out different referents in other possible situations. An example of such an interpretation of *I* is provided by the statement "I could have been a burglar", made by the daughter of an elderly widow living in a dangerous part of town who has just responded to the ringing of the door bell by opening the door without first looking through the spy hole. The intuitively correct interpretation of this statement involves the description 'the person who has rung the door bell'. This description is instantiated by the actual referent of *I* in the daughter's statement, viz. the daughter herself, but by individuals distinct from the daughter in the counterfactual worlds brought into play by the modal *could* (e.g. professional burglars who have been making the neighborhood unsafe). This important aspect of the semantics and pragmatics of indexicals is orthogonal to the issues discussed in the present paper and so it has been said aside here.

fact that expressions of which it has been claimed that they are indexicals often do behave like true indexicals, getting their referents from the utterance context and then holding on to these at any world or situation of evaluation. In those cases the utterance context does play the reference-determining role for which the work of Kaplan and others has accounted. It is to those occurrences of "indexical" expressions that we will be referring when speaking of 'indexicals' in this paper. How the occurrences of expressions that do behave like pure indexicals can be distinguished from the ones that don't is a non-trivial, intriguing and, to my knowledge, unsolved question as well. But it too is one I will ignore.

One can formally define utterance contexts as sequences consisting of a number of 'actual' entities: the actual speaker who makes the utterance, the actual time at which the utterance is made, the actual place of the utterance and (when defined) the utterance's actual addressee or audience. Since each of these entities is separately determined by the utterance, we can also think of the components of utterance contexts as the results of applying certain functions to the utterance in question: a speaker function, which assigns to each utterance the actual speaker of that utterance, an utterance time function, which assigns to each utterance the time at which this utterance is made (or treated as made by the interpreter), and so on[13]. (The addressee/audience function would be a partial function, the others we assume to be defined for all utterances.)

It is this *concept of an utterance context*, as a bundle of functions which map utterances to certain entities associated with them, that I see as the crucial ingredient to the accounts of indexicality that Kaplan and others have contributed to the theory of meaning. I will refer to this bundle as the 'Utterance Context Concept' or, more briefly, as 'UCC'. I assume that some fixed bundle of such functions is given, though I do not want to commit myself here to a complete list of such functions. But I will assume that at least the following functions are included: f_{sp}, the function which assigns to each utterance its speaker or writer; f_{n}, the function that assigns each utterance its utterance time; and f_{ad}, the partial function which is defined only for those utterances

[13]There are various problems connected with these different functions. For instance, what is "the" time at which a given utterance occurs? These problems are non-trivial and for the analysis of some utterances and discourses they can be important. (See eg. Reyle et al. (2007), Sn. 6). But we won't worry about such problems here.

that are addressed to an audience or to a single addressee and which returns that audience or addressee when it is applied to an utterance of that sort.

Kaplan includes among the expressions that get their semantic values from the utterance context not only the indexicals but also (deictic uses of) demonstrative phrases like *this*, *that*, *that bird over there* and so on and extends the tuples that are his utterance contexts accordingly (so that they include the referents for all the deictically used demonstrative phrases that the given utterance contains). It is not immediately clear how room can be made for such entities at the level of an Utterance Context Concept. But presumably it would make it necessary to make UCC a concept that varies as a function of what utterance is under consideration, with each utterance giving rise to a UCC that has functions for each of the deictic demonstratives that it contains. Extending the notion 'Utterance Context Concept' would be possible in principle. But we will not follow Kaplan in adopting as broad a notion of Utterance Context as he favours. Our notion of Utterance Context Concept will be limited to functions of the kind mentioned above (such as f_{sp}). The referents of deictically used demonstrative NPs will be treated as constituents of another context – the 'Environment Context' – which will be introduced in Section 3.

2.3 Discourse Contexts

The distinction between the general concept of an utterance context and particular utterance contexts is important for us, since it is the concept of an utterance context that will be instrumental in the unification of utterance context and discourse context we are aiming for. The reason why there is a potential difficulty here has to do with the notion of discourse context that we will be using, which is part of the version of dynamic semantics with which I will be working, viz. Discourse Representation Theory. DRT differs from other versions of dynamic semantics in that it describes language interpretation as a process that creates and operates on semantic representations. In particular, the processes involved in linking a sentence interpretation to the interpretation of the preceding discourse are analyzed as operating on (i) a representation derived from the syntactic form of the sentence and (ii) a representation of the preceding discourse. This second representation plays the part of discourse context. Since it is a representation, it does not involve real entities directly, although its constituents can be 'anchored' to real entities. These "anchors" to real entities will play an essential part in Section

3, where they will be mediated by Entity Representations.

Entity Representations belong to MSDRT, to which we will turn later. Here I limit myself to a brief review of those aspects of DRT that will be relevant for the purposes of this paper by looking in some detail at an example. Most importantly for present purposes, this example can be used to show how the representational setting of DRT makes it possible to unify utterance context and discourse context into a single context representation[14]. The example consists of a two-sentence "discourse", given in (6):

(6) i. Last week Fred bought a donkey.
ii. He sold it the next day.

In the version of DRT we will be using, interpretation of this two sentence discourse takes the form of:

(a) constructing a semantic representation of (6i),
(b) constructing a preliminary semantic representation of (6ii), and
(c) incorporating this preliminary representation into the representation of (6i) (which at this point serves as representation of the discourse context for (6ii)).

The result is the semantic representation shown in (11) below for the two sentences of (6) taken together.[15]

The DRS for (6i) is given in (7).

[14]The example should also give readers who are unfamiliar with DRT some general sense of how the theory works and what its representations look like. I must stress, however, that it is not my aim to make this paper fully accessible to readers without any previous exposure to DRT. Much of the paper should be understandable even without such exposure, but probably not all of it.

[15]DRT assumes that the construction of a semantic representation for a sentence S proceeds from a syntactic analysis of S, which is provided by some parser that has already done its work when the representation construction starts. (This is a notoriously unrealistic idealization, but one which the theory shares with most other formal approaches to the syntax-semantics interface.) Here I do not show the syntactic trees for (6i) and (6ii) nor do I say anything about the construction principles which convert such trees into semantic representations.

(7) (Discourse Representation Structure for (6i))

n t_1' t_1 e_1 f d
"week-before-the-week-of"(n)(t_1') $t_1 < n$ $t_1 \subseteq t_1$' $e_1 \subseteq t_1$
Fred(f) donkey(d)
e_1: buy(f, d)

Like any other DRS, (7) consists of two components, (i) its Universe and (ii) its Condition Set. The Universe is a set of discourse referents. These function as representations of entities. The Condition Set consists of DRS conditions, which attribute properties and relations to the entities represented by the discourse referents in the Universe. Thus "donkey(d)" states that the entity represented by d is a donkey, "Fred(f)" that the entity represented by f is the bearer of the name *Fred* (more exactly: the individual that the speaker has used the name *Fred* to refer to in her utterance of (6)). "e_1: buy(f,d)" states that the entity represented by e_1 is an event of f buying d. The remaining discourse referents and conditions have to do with the temporal location of e_1. These elements are contributed by (i) the past tense of (6i) and (ii) the adverbial *last week*. t_1 represents the location time of e_1. (This is expressed by means of the condition "$e_1 \subseteq t_1$", which says that e_1 is temporally included within t_1.) The past tense of (6i) contributes the information that t_1 is in the past of the utterance time, which is represented by n. The temporal location of e_1 is also, and more narrowly, characterized by the adverb *last week*; this second temporal constraint is expressed by the condition "$t_1 \subseteq t'_1$", where t'_1 represents the time denoted by the adverbial. The condition expressing that t'_1 is the denotation of *last week* has been abbreviated as "week-before-the-week-of'(n)(t')"[16]. The intention should be clear: the condition fixes the entity represented by t'_1 to be the week immediately preceding the one which

[16]The semantics of *last week* can be computed from the semantics of the word *last* and the word *week*. A proper treatment of the semantics of *week* should be part of a general semantics for "calendar terms" and of the temporal ontology underlying this semantics. Such an ontology must consist of (i) an account of the logical structure of time, and (ii) the largely conventional partition of the time axis into calendar-related intervals, determined according to the principles of the relevant calendar. (In general the relevant calendar will be that of the culture within which the represented utterance or discourse is embedded. For current Western culture this is the Gregorian calendar. For some of the details for such an account see e.g. Kamp and Schiehlen (2002).)
As noted earlier, phrases like *last week*, consisting of the word *last* followed by a calendar term like *week*, are indexical expressions. For instance, *last week* normally denotes the last

contains the utterance time that is represented as n.[17]

DRSs can be regarded as the formulas of some formal language. Current DRT offers a range of such "DRS-languages", designed for the analysis of different linguistic phenomena (and thus to be used as representation formalisms for language fragments within which these phenomena occur). Each DRS-language is specified by way of explicit definitions of its syntax and its model-theoretic semantics; in this regard DRT abides by the standards in formal logic for the design and definition of formal languages. The truth definitions for DRS languages, which are the core of their semantics, take a form which exploits the structure of DRSs. As noted above, DRSs are pairs consisting of a *universe* (a set of discourse referents) and a set of so-called DRS conditions. A DRS K = <U,Con> is defined as true in a model M iff there exists an "embedding" function from the universe U of K into the universe of M under which all conditions of Con are satisfied in M. One consequence of this way of defining truth is that as far as the truth conditions of K are concerned the discourse referents in U play the part of existentially bound variables. DRT does not acknowledge a separate category of individual constants, but discourse referents can play a constant-like role. More specifically, discourse referents can play the part of individual constants when they are anchored to particular entities. It is part of the logic of discourse referent anchoring that anchors are respected by all permissible embedding functions. (An embedding function f respects the anchoring of x to $\mathbf{d}$ if $f(x) = \mathbf{d}$.)

Anchors can come about in various ways. One of these is the interpretation of indexical expressions. When a discourse referent is introduced to represent the referent of an indexical expression, this always carries with it the introduction of an anchoring relation or function that links the discourse referent to the referent. Discourse referents that arise in this way, and that are anchored by such a relation or function to the entities they represent, are called *indexical discourse referents*. The DRS (7) contains one example of an indexical discourse referent, viz. n.

---

interval satisfying the extension of the noun *week* which lies entirely before the utterance time.

[17] Since n is always available in the construction of DRSs there is no need to mention it explicitly as part of the Universe of any DRS. So in practice occurrences of n in DRS Universes are often omitted. From here on I will follow this practice.

To explain how occurrences of n get anchored to the times of the utterances represented by the DRSs containing them it is necessary to say more about the models for the intended DRS language (of which the DRS in (7) is one of the formulas). Three properties of these models are important in the present context:

(i) Each model $\mathbb{M}$ must have a range W of different possible worlds, with respect to which each DRS from the given DRS language determines a (non-degenerate) *intension*. The intension of a DRS K in model $\mathbb{M}$ is composed form the truth values that K has in each of the possible worlds of $\mathbb{M}$ (or, more formally, the function which maps each world w in W to the truth value of K in w in $\mathbb{M}$).

(ii) In order that every DRS K has a well-defined truth value at each of the worlds w in W, $\mathbb{M}$ must determine for each w in W an extensional model $\underline{\mathrm{M}}_w$. In each such model $\underline{\mathrm{M}}_w$ an extensional DRS K will be either true or false, according to the truth definition familiar from standard DRT: K is true in the model iff there exists a verifying embedding of it in $\underline{\mathrm{M}}_w$ and false if there isn't such an embedding.[18] The *intension* of K in $\mathbb{M}$ — or, as it is often also called, the *proposition expressed by* K in $\mathbb{M}$ — can then be defined as the function which maps each w ϵ W to the truth value of K in $\underline{\mathrm{M}}_w$.

(iii) Our models must have a sufficiently rich ontology. Since our DRSs contain discourse referents for times and events:[19] this ontology must include entities of those sorts; otherwise there could be no verifying embeddings in these models of DRSs that contain such discourse referents. In fact, the models M_w, whose ontology includes both times and eventualities, should specify among other things which of their eventualities go on or hold at which times, and thereby trace, so to speak, the histories of the worlds w with which they are associated: through their connections with times the eventualities of $\underline{\mathrm{M}}_w$ are temporally ordered, laying out the world w from its temporal beginning to its end.[20].

[18] An extensional DRS K is a DRS which contains no intensional operators. All DRSs we will consider in this section are extensional in this sense.

[19] And also states. As it happens, (7) provides no examples of state discourse referents, but we will encounter instances of state discourse referents later on.

[20] Here I am cutting a long story short. In particular nothing will be said in this paper about the ways in which the time structures of different worlds may be connected with

When two possible worlds are very different from each other, their respective time structures may be different as well. However, for the considerations of this paper there is no need for worlds that differ from the actual world to such an extent. On the other hand, interpretations of many intensional constructions, some of which will be important in what follows, are premised on the assumption that all worlds involved share the same time structure. We will therefore make the simplifying assumption that for each model $\mathbb{M}$ there is a single time structure $T_{\mathbb{M}}$ that is common to all the extensional models $\underline{\mathrm{M}}_w$ associated with the worlds w of $\mathbb{M}$.

These considerations lead to models $\mathbb{M}$ of the following form:

(8) (Definition of intensional models)

By a *model* we understand a tuple $\mathbb{M}$ = <W, T, $\underline{\mathrm{M}}$>, where:

i. W is a non-empty set of "possible worlds";

ii. T is a time structure <T,<>, where T is a set of temporal instants and < is the "earlier-later" relation between instants;

iii. $\underline{\mathrm{M}}$ is a function from worlds w to "histories unfolding in T", that is, to models $\underline{\mathrm{M}}_w$ each of which tracks the development of its world through time, where time is given by the structure T.[21]

each other. For an early discussion of the interactions between worlds and times see eg. (Thomason and Gupta (1980). For some aspects of the DRT treatment of the relations between worlds and times see Kamp et al. (2011) and Kamp and Reyle (2011)).

[21] The way in which a model $\underline{\mathrm{M}}_w$ tracks a possible history is different from the formalization of the notion of a 'history' in what may be the most familiar form, viz. that provided by the model theory of systems of Priorean tense logic. In that model theory the models are 'temporally intensional' models $\mathbb{M}$, which for each t from a given time structure T provide an extensional model $\mathbb{M}_t$ that determines the extensions of the non-logical constants of the language in $\mathbb{M}$ at t. In DRS languages like the one exemplified in (7) temporal dependence at the atomic level is found only in those predicates that are used to represent verbs. An example is the transitive event verb *buy*, whose counterpart in the DRS language of (7) is the three-place predicate 'buy' ', which is found in 'atomic predications' of the form 'e: buy'(x,y)'; and another is the stative verb *love*, with the predicate 'love' as its counterpart in our DRS language, which occurs in atomic predications of the form 's: love'(x,y)'. The extension of the predicate 'buy' ' in $\underline{\mathrm{M}}_w$ will be a set of triples <e,a,b> such that e is an event of a buying b in the world w described by $\underline{\mathrm{M}}_w$, and the extension of 'love' ' in $\underline{\mathrm{M}}_w$ will be a set of triples <s,a,b> such that s is a state of a loving b. The time dependence of the extensions of these predicates is given through the location of the events and states

Above I claimed that in an extensional model $\underline{\mathrm{M}}_w$ any extensional DRS, such as for instance (7), is either true or false. But now that I have said what these models are actually like it might look as if that could not be right. $\underline{\mathrm{M}}_w$ presents the development through time of the world w. A DRS like (7) can be true, it would seem, at one time t of this model while failing to be true at some other time t': there could be a verifying embedding of (7) in $\underline{\mathrm{M}}_w$ which maps n to t, but no verifying embedding which maps n to t'. This apparent inconsistency can be resolved as follows. (7) is to be understood as the representation of *some particular utterance* of (6i), made at some particular time t (rather than as a representation just of the sentence as type). When a DRS plays this role, then the time of the represented utterance fixes the denotation of the occurrences of n contained in this DRS. That is, for the DRS *as representation of an utterance made at t* only embeddings f are admissible which obey the constraint that $f(n) = \mathrm{t}$. If there is a verifying embedding into $\underline{\mathrm{M}}_w$ that answers this constraint, then (7), as representation of the given utterance of (6i), counts as true in $\underline{\mathrm{M}}_w$ (and therewith as establishing the truth in $\underline{\mathrm{M}}_w$ of the utterance it is being used to represent); if there is no verifying embedding satisfying the constraint, then the DRS and the utterance it represents are false in $\underline{\mathrm{M}}_w$.

---

involved. That is, for each event e and state s in the universe of $\underline{\mathrm{M}}_w$ the model specifies their *duration*, τ(e) or τ(s), as an interval of T. We can construct models $\underline{\mathrm{M}}_{w,t}$ from the model $\underline{\mathrm{M}}_w$ which resemble the models $\mathbb{M}_t$ of the model theory for tense logic. These are models for a predicate logic which differs from our DRS language in that each n+1-place predicate P' is replaced by an n-place predicate P". (For instance, the 3-place predicate 'buy' ' is replaced by a 2-place predicate 'buy" '.) The extension of such a predicate P" in $\underline{\mathrm{M}}_{w,t}$ consists of all n-tuples <a_1,...,a_n> such that for some eventuality (that is: event or state) ev <ev,a_1,...,a_n> belongs to the extension of P' in $\underline{\mathrm{M}}_w$ and t belongs to τ(ev). (Thus the extension of 'buy" ' in $\underline{\mathrm{M}}_{w,t}$ will consist of all pairs <a,b> for which there is an event e such that <e,a,b> belongs to the extension of 'buy' ' in $\underline{\mathrm{M}}_w$ and t belongs to τ(e).) As noted, DRS-languages like that instantiated in (7) only account for the temporal dependence of verbal predicates (i.e. those that are used to represent verbs). Non-verbal predicates (used to represent nouns, adjectives, prepositions) are treated as temporally constant. This is a simplifying assumption that is of course blatantly incorrect, notwithstanding the fact that it is very often made in formal semantics. In more recent versions of DRT this is being corrected. (See Tonhauser (2002) for discussion and the forthcoming Kamp and Roßdeutscher (2019) for an account of the temporal dependence of non-verbal predicates as part of an up-to-date, wide coverage DRS construction algorithm.)
In what follows the models $\underline{\mathrm{M}}_{w,t}$ are not quite what is needed. This is because verification of DRSs in models presupposes that a time of evaluation (e.g. the time of the utterance whose content is represented by the DRS) is given. verification of the DRS in $\underline{\mathrm{M}}_w$ then entails that its n is mapped to t. See the main body of the text.

So, for the time being at least[22], let us assume that the DRSs play the part of representations of particular utterances. In order to make formal sense of this we need to assume that the utterances represented by our DRSs are themselves events in the models in which the DRSs are assigned their extensions and intensions. Such an utterance, made in some world w of a given intensional model $\mathbb{M}$, will determine a time of $\mathbb{M}$ as the utterance time and thereby fix the embedding target of the occurrences of n in its representing DRS.

To simplify what follows I will assume some particular intensional model $\mathbb{M}$ as given and also that one of its worlds is the actual world w_0. Moreover, we will, for the time being, be thinking of the represented utterances as actual utterances, i.e. as utterances made in w_0, and thus as events occurring in $\underline{M}_{w_0}$. None of this is essential, but it helps to focus on what does matter.

(9) repeats what has just been said about the temporal constraint that the utterance represented by a DRS imposes on the denotation of the occurrences of n it contains, and therewith on its embedding conditions in any of the models $\underline{M}_w$.

(9) An occurrence of n in a DRS K that represents the content of an utterance u represents the utterance time t_u of u.

(9) exemplifies a general restriction on the semantics of indexical discourse referents. Such discourse referents must always get their values from some utterance context. This condition is satisfied when an indexical discourse referent occurs in a DRS that represents some particular utterance. As implied above (see footnote 22), this is not the only way in which such DRSs can be interpreted. Section 2.5 describes an alternative.

After this outline of the semantics of DRSs like (7) we return to the DRS construction for (6). Our next step is the construction of the preliminary DRS for the second sentence of (6). It is given in (10).

[22] An alternative option is discussed in Section 2.6

(10) (Preliminary DRS for (6ii))

t_2' t_2 e_2 x y
"the-day-after-that-of"(?)(t_2') $t_2 < n$ $t_2 \subseteq t_2$' $e_2 \subseteq t_2$
x = ? y = ?
e_2: sell(x, y)

Most of what needs to be said about (10) has already been covered by the comments to (7). Two points remain. The first concerns the question marks in the conditions " "the-day-after-that-of"(?)(t_2')", "x = ?" and "y = ?". This use of question marks is a device that is sometimes resorted to in DRT as a shorthand for representations of *anaphoric presuppositions*. Pronouns and other anaphoric expressions are assumed to give rise to "anaphoric presuppositions" — presuppositions that can be resolved by identifying a semantic value for the anaphoric expression in the context in which it is used. Thus, the conditions "x = ?" and "y = ?" represent the anaphoric presuppositions that come with the pronouns *it* and *his*. Resolution of these presuppositions consists in finding discourse referents in the discourse context to serve as (representatives of) antecedents for these pronouns; these discourse referents are then substituted for the question marks, thereby turning the conditions into equations that express the relation of 'coreference' between the pronoun and the constituent that the resolution identifies as its antecedent.

In " "the-day-after-that-of"(?)(t_2')" the question mark indicates the implicit argument of the word *next*. This brings us to the second point. Observe that an NP like *the next day*, just as *the next week*, is not an indexical like *next week* but an anaphoric expression, which must be interpreted as referring to the day which follows a day that is represented (and salient) in the context. (Why it is the presence of the definite article *the* that makes this difference between an indexical and an anaphoric semantics is something I am not trying to explain here.[23]) Hence (10) also has a condition with a question mark, to be replaced by some time- or eventuality-representing discourse referent

[23]There is a further feature that distinguishes, say, *the next week* from *next week*. It isn't just that *the next week* is anaphoric rather than indexical. It is in fact *anti-indexical* in the sense that it cannot be used to refer to the week following the week of the utterance time (that is, not *as* the week following the week of the utterance time, in the way that *next week* is (and must be) used. (That is, the utterance time n may not be used as anaphoric antecedent for *the next week*; accidental coreference with the week following that of the

α that is available in the discourse context. Substitution of α for '?' yields a condition which says that t_2 represents the day following the day that includes the time represented by α (in case α represents a time) or the duration of α (in case α represents an eventuality).)

Combining (10) with (7), which acts as discourse context for the interpretation of (6ii), requires resolution of the question marks in (10). As we just noted, this takes the form of finding discourse referents in the discourse context, provided by (7), that can serve as anaphoric antecedents for the discourse referents co-occurring with the question marks.[24]

Resolution of the presuppositions introduced by *it*, *his* and *the next day* to, respectively, the discourse referents x, y and t_1 from the discourse context (7), takes the form of substituting these discourse referents for the relevant question marks in (10). The DRS this yields can then be merged with (7). The result is (11), a DRS representing the utterance (6) consisting of the sentences (6i) and (6ii) together.

---

utterance time aren't excluded.) The same is true of *the next day*, which cannot be used in the way we use *tomorrow*. Anti-indexicality is a widespread phenomenon, which is also found, for instance with the indexical pronouns (1st and 2nd person) and the non-indexical ones (3rd person). You cannot uses a third person pronoun in the normal way to refer to yourself, nor can you use it to refer to your addressee. It seems something of a general principle that indexical expressions and anaphoric/deictic expressions came as complementary pairs, with distinct, non-overlapping uses.

[24]Exactly how antecedents for anaphoric noun phrases are identified is a notoriously hairy problem, but it is one we set aside here (as has been the usual practice for presentations of DRT in which there is no special emphasis on the compositional aspects of DRS construction.) One of the complications that a comprehensive theory of anaphora resolution has to deal with is that different anaphoric elements are governed by different resolution principles. We already noted one example of such a difference: the principle which governs anaphora resolution for pronouns differs from that governing resolution of anaphoric definite descriptions in that the constraints on pronoun resolution appear to be more restrictive than the constraints on resolving the presuppositions of anaphoric descriptions. (For discussion see Heim (1990), Van Der Sandt (1992), Roberts (2012)).

(11)

t_1' t_1 e_1 f d	t_2' t_2 e_2 x y
"week-before-that-of"(n)(t_1')	"the-day-after-that-of"(e_1)(t_2')
$t_1 < n$ $t_1 \subseteq t_1$' $e_1 \subseteq t_1$	$t_2 < n$ $t_2 \subseteq t_2$' $e_2 \subseteq t_2$
Fred(f) donkey(d)	x = f y = d
e_1: buy(f,d)	e_2: sell(x, y)

2.4 Utterance Context as Input to the Semantics of Discourse Representations

Both (7) and (11) are 'proper' DRSs, which represent well-defined propositional contents (of (6.i) and (6.i,ii), respectively). But what exactly are these contents, and exactly how do DRSs like (7) and (11) determine them? As noted, the classical truth definition for DRT (Kamp and Reyle (1993)) has it that a DRS K is true in an extensional model M iff there is a verifying embedding of K into M. We have indicated how this definition must be adapted so that it can take proper account of occurrences of the indexical discourse referent n. It is time to see how this adaptation can be made explicit.

This time we focus on (11). Assume that (11) is the representation of an utterance u of (6.i,ii) at a time t_u in a world w_u belonging to some given model $\mathbb{M}$. For each world w of $\mathbb{M}$, (11) will, as representation of u, get the value 'true' iff there exists a verifying embedding f of (11) into $\underline{M}_w$ *such that f(n)* $= t_u$. Furthermore, we can, as before, obtain the propositional content of (11) relative to $\mathbb{M}$ by collecting the truth values thus defined for the different worlds of $\mathbb{M}$ into a single object, the function which maps each world w of $\mathbb{M}$ to the truth value in M_w of (11) as representation of u.

More generally, let K be a DRS for an utterance u of some discourse D and suppose that n is the only indexical discourse referent in K. Let w_u be the world of $\mathbb{M}$ and t_u the time of $\mathbb{M}$ such that u occurred in w_u at t_u. Then we can define the *propositional content of* K *as representation of* u *relative to* $\mathbb{M}$ to be the function which assigns to each world w of $\mathbb{M}$ the truth value in M_w of K as representation of u; and here, once again, the truth value of K is determined using only embedding functions f such that $f(n) = t_u$. We denote the truth value of K as representation of u in M_w as '$[K]_{u,\mathbb{M},w}$'. The propositional content of K in $\mathbb{M}$ as representation of u — also called '*the*

proposition expressed by K *in* $\mathbb{M}$ *as representation of* u' or, alternatively, '*the proposition expressed by* u *in* $\mathbb{M}$ *according to its representation* K — is then the function λw.$[K]_{u,\mathbb{M},w}$.[25]

2.5 Indexical Discourse Referents other than n. Singular Propositions and Partial Propositions

The anchor t_u of the occurrences of n in K can be thought of as a kind of constituent of the propositional content that K determines when it plays the part of representation of the particular utterance u. For it is always this time t_u that is used in the evaluation of K (i.e. t_u enters into the evaluation of K at each world w of $\mathbb{M}$). Familiar philosophical terminology expresses this by saying that the propositional content is a *singular proposition with respect to* the time t_u. Propositional singularity is an extremely common phenomenon. For one thing, utterances of by and large all natural language sentences anchor some or all of the predications occurring in the sentences temporally to the utterance time t_u (either directly, as holding at t_u, or indirectly, as holding at some time that stands in a certain temporal relation to t_u). Thus the proposition expressed by a DRS K representing any such utterance will be singular at least with respect to n. Moreover, singularity is by no means restricted to occurrences of n. Every anchored discourse referent will be a source of singularity. That includes all discourse referents that are introduced as representatives of indexical expressions. (We will see later that anchoring of discourse referents can arise in many other ways as well.)

Indexicals other than those that are represented with the help of n also give rise to singularity. For example consider an utterance u, by a speaker d_u at a time t_u in a world w_u, of the sentence

(12) I am hungry.

[25]Note that when the proposition expressed by the utterance u in w_{u} is defined in this way, it is true in some other world w' of $\mathbb{M}$ iff K is verifiable in $\underline{\mathrm{M}}_{w'}$ at the time t_u at which the represented sentence or discourse is uttered in w_{u}. According to the model theory for DRT that we have adopted this is also true for worlds w' in which u was made at some time different from t_u (e.g. shortly before or shortly after that time). I think this is intuitively plausible, and as far as I know it is consistent with the views of others who have expressed the view that an utterance fixes the time at which the proposition it expresses must be true in non-actual worlds.

This sentence contains two indexical elements, the present tense and the pronoun *I*. The sentence can be represented by means of the DRS in (13), in which n represents the utterance time and *sp* the speaker.

(13)

t	s
n ⊆ t	t ⊆ s
s: hungry(sp)	

(13) thus has two indexical discourse referents, n and *sp*. Each introduces an element of singularity into the propositional content that (13) represents. We have already seen in what way this is done by n: At any world w (13) is true iff it has a verifying embedding f in M_w such that $f(n)$ is the actual utterance time t_u. *sp*, which is introduced as representative for the referent of the pronoun *I*, is responsible for adding a second dimension of singularity. It too is anchored, to the actual speaker d_u of the represented utterance. So evaluation of (13) at any world w only admits embedding functions f such that $f(sp) = \mathrm{d}_u$.[26]

Anchoring *sp* to d_u, however, creates a potential problem: in all likelihood there will be worlds in our given model $\mathbb{M}$ in which d_u does not exist. In such a world evaluation of (13) by a function f for which $f(sp) = \mathrm{d}_u$ makes little intuitive sense. What would it mean for d_u to be hungry – or for that matter, for d_u not to be hungry – in a world in which d_u doesn't exist? Such worlds are better excluded from the propositional content of the utterance, for if we do include them, then some arbitrary decision has to be made about the truth or falsity of the intuitively meaningless predications in which d_u is nominally involved; and that can only lead to confusion.[27]

However, by excluding worlds from the domain of the propositional content of u, we abandon the traditional concept of a proposition as a function that is defined for all worlds. That is, we can no longer insist that all proposi-

[26]The Universe of this DRS neither contains an occurrence of n nor of *sp*. In omitting *sp* we have applied the same convention that we adopted earlier for n in footnote 17. This convention can be applied to all indexical discourse referents.

[27]Settling such intuitively meaningless predications by stipulation doesn't seem to work very well. At least, all global conventions for doing this that I have seen have counterintuitive consequences for the logic that comes with a semantics which includes such stipulations.

tions are total functions, and have to settle in general for *partial* propositions – functions from worlds to truth values that need not be defined for every world. As a matter of fact, the need for partial propositions arises for various other reasons as well; non-existence of the entities purportedly represented by indexical discourse referents is only one of them. Partial propositions, and more generally partial intensions, are an inescapable reality of natural language semantics, one that we are well-advised to accept for what it is and with which we must learn to cope as well as we can.

One problem we are faced with when we allow for partial intensions as semantic values of expressions concerns the domains of those intensions: What more can we say about their domains than that they aren't always the set of all worlds? In general this question appears to be quite involved. But here we will assume that existence failure is the only source of failure, in other words that the intension is defined in all worlds in which all the anchors for indexical discourse referents (and other discourse referents that require an anchor) exist.

From a traditional perspective in which it is taken for granted that propositions (and intensions generally) are total functions, partiality is bound to look like an imperfection. But when we remind ourselves that language is used for communication, that communication always takes place against the background of some given set of assumptions and that its function is to update this set of assumptions by adding new information to it, then we can see that partiality need not be much of a problem. What is more, it is something that is only to be expected: As long as the expressions that the speaker uses all have intensions that are defined for every world compatible with the background assumptions, communication will function smoothly, for these are the only worlds that matter. Of course, partiality *can* be a problem, viz. when an intension fails to be defined for a world that is admitted by the background assumptions. Definedness for these worlds is a general presupposition of proper language use. If this presupposition fails, then updating the given assumptions with the information conveyed by the current sentence will be in jeopardy, and some repair strategy will be needed (e.g. by an accommodation of the background assumptions such that these now ensure that the presupposition is satisfied after all).

2.6 DRSs as Representations of Content and Representations of Character

So far we have looked at DRSs like (7) and (11) as representations of particular utterances. But this is only one of the ways we can look at such DRSs. Contrary to what I implied earlier, they can also be taken as representations of sentence or discourse *types* – as representations of linguistic expressions as such. Looking upon a DRS like (7) in this second way is tantamount to abstracting over the different contexts on which the contents of the different utterances of (6i) depend. And this move – of abstracting out the impact of utterance context on represented content – is closely reminiscent of the move that Kaplan makes when he distinguishes character from content. Kaplan's characters are functions that map utterance contexts to intensions; these are functions in their turn, which yield extensions when applied to 'circumstances of evaluation. In a DRS like (7) the dependence on utterance context is represented through the presence of n. We can see this DRS – let us refer to it as 'K' for the remainder of this discussion – as specifying a function from any possible value t_u for n to the proposition that K determines when regarded as representation of an utterance made at t_u.

So that we can be a little more explicit, let once more $\mathbb{M}$ be some given model. Then, as discussed in section 2.3, the proposition expressed in $\mathbb{M}$ by K as representation of an utterance u (made at the time t_u of $\mathbb{M}$) is the function $\lambda w.[K]_{u,\mathbb{M},w}$, the function which maps each world w from $\mathbb{M}$ to the truth value $\underline{\mathrm{M}}_w$ of K as a representation of u, in.

When Kaplan's distinction between character and content is transferred to our set-up, it is the worlds w from $\mathbb{M}$ which now play the part of circumstances of evaluation. Furthermore, for the case we are considering the only way in which the content $\lambda w.[K]_{u,\mathbb{M},w}$ of K depends on the utterance context is via t_u. So we can also represent this content as $\lambda w.[K]_{t_u,\mathbb{M},w}$.

The parallel is now perfect: For any time t_u of $\mathbb{M}$ $\lambda w.[K]_{t_u,\mathbb{M},w}$ denotes the proposition that the sentence represented by K would express, relative to $\mathbb{M}$, if it was uttered at t_u. Thus $\lambda t_u.\lambda w.[K]_{t_u,\mathbb{M},w}$ denotes the character of the represented sentence: the function that, when applied to any (real or possible) utterance time t_u yields the propositional content $\lambda w.[K]_{t_u,\mathbb{M},w}$. This content yields its extensions (= truth values) when applied to the 'circum-

stances of evaluation' $w \in W_{\mathbb{M}}$.

In the last few paragraphs I have been speaking of 'possible utterance times'. But what is a possible utterance time? Are all times of $\mathbb{M}$ possible utterance times of $\mathbb{M}$? Or should utterance times be subject to some restrictions? For what follows this question isn't of direct importance. But it is important in connection with a complex issue that should at least be mentioned and it also offers me an opportunity to adjust a remark I made earlier. To begin with the adjustment: In Section 2.3 I intimated that it mattered little for our discussion whether we would confine our attention to actual utterances or take possible utterances into account as well. That show of indifference was defensible when I made it, but it isn't any longer. For now that we have introduced the notion of the character of an expression type, we have therewith also raised the question what the domains are on which characters are defined. And whatever character domains may be, it seems quite clear that they should not consist exclusively of the utterance contexts of actual utterances. This is a point that Kaplan has argued emphatically and conclusively. To mention just one of his points – perhaps the most decisive one of all – so far most sentences of English have never been uttered by anyone. And for all we can guess – and the guess is a pretty safe one – most sentences of our language will never be uttered at all. But unuttered sentences are just as meaningful as those that make it at some point or other into actual utterances. All of them, actualized or not, should be assigned non-vacuous character, which determine the contents that would be expressed through any utterances of them that have been, or will be or could have been made. And we get non-vacuous characters for unuttered sentences only when their domains contain possible as well as actual utterance contexts.

With this rectification behind us, let us go back to the question that prompted it: What is a possible utterance time? One reason why answers to this question matter is that they may have an impact on the logic of temporal indexical expressions that the semantics will generate. To give just one example – the example is somewhat strained, but in this respect no worse than many others that can be found in the philosophical logic literature – the sentence 'Someone currently exists.' would come out as 'logically true' if a time qualifies as a possible utterance time only if there is someone around at that time who is capable of producing an utterance (though irrespective of whether any utterance was actually produced at that time by anybody); but without this

constraint the sentence would not come out as logically true. To be a little more precise, we get these verdicts if logical truth is defined as follows: a sentence is logically true if every possible utterance of it expresses a proposition that cannot be false; or, using the more technical vocabulary of character and content: the sentence is logically true iff any application of its character to an element in its domain will invariably yield a proposition that is true.

There may be no single true answer to the question what should count as a possible utterance time, and in that case the best a logic of temporal indexicals can deliver is an analysis of how the set of logical truths depends on the different ways in which 'possible utterance time' might be defined. (This would be a situation much like the one we find in relation to Priorean Tense Logic.)[28]

So far in this section we focused on the DRS (7). But the same considerations can also be applied to the DRS (11) for the succession of the two sentences in (6). This DRS too can be looked at either as the representation of a particular utterance of (6) or alternatively as the representation of the character of (6), seen as the succession of two sentence types. Note that in this second capacity (11) is the representation of the character of (6) *given a certain resolution of the anaphoric constituents of its second sentence (6ii).* On this view resolution of anaphoric relations has priority over the exploitation of the utterance context in fixing the values of indexicals. I am not sure that this priority issue has ever been raised. But it would seem plausible enough that there should be such a priority. To resolve the anaphoric presuppositions that need to be dealt with as part of interpreting (6) we do not need to know anything about actual or possible utterances and their times. Discourse-internal anaphora resolution is part of the interpretation process

[28] An important part of the Logic of Demonstratives as Kaplan has formulated it has to do with the question what distinguishes potential utterance times from times in general, and, more generally, what distinguishes potential utterance contexts from arbitrary situations. (See in particular Section XIX of Kaplan (1989).) It should be noted that the model theory we have adopted adds a further complication to these logical matters. Our model theory admits partial propositions. Logical consequence relations can still be defined forv accounts that admit partial propositions (e.g. by using versions of supervaluation (van Fraassen (1975))). And the impact that different definitions of 'possible utterance time' or 'possible utterance context' will have on the set of logical truths will remain. The technicalities involved in formulating logics of indexicals in such a partial setting are non-trivial, but this is not the place to delve deeper into this.

that yields a representation of character, and, via that, an identification of utterance content. Indeed, if we want to apply the character-content distinction to discourses consisting of multiple sentences as well as to single sentences then the priority is inescapable. This is an important moral for accounts that cover both indexicality and discourse anaphora.

2.7 Embedding Utterance Contexts within Discourse Context Representations: Generalising from the Case of Time

In the last section we saw how one component of the utterance context, viz. the utterance time, can be captured in semantic representations provided by DRSs through the use of the indexical discourse referent n. We have also seen, at the hand of example (13) in Section 2.5, how representations like that of the utterance time by means of n are equally possible for other components of the utterance context. It is not hard to see how these two lines of thought can be combined to yield a complete incorporation of utterance contexts into our discourse contexts.

Recall our decision in Section 2.2 to distinguish between individual utterance contexts and the *concept* of an utterance context (UCC). The UCC was identified as a bundle of functions that map possible utterances to the different components of their contexts. So far we have been concentrating on one of these functions, the function f_n which maps utterances to their utterance times, and in our semantic representations we have been making use of the indexical discourse referent n, which corresponds to this function in that in any utterance representation it represents the time that f_n determines for the represented utterance. We now extend our DRS language by assuming that there is an indexical discourse referent corresponding to each of the functions in UCC. In particular, the discourse referent corresponding to the speaker function (which assigns to each utterance its speaker or author) will be *sp* and that corresponding to the addressee function will be *ad*. Introduction of a token of *sp* into a sentence DRS is the result of interpreting the pronoun *I*; tokens of *ad* are mostly the result of interpreting the singular pronoun *you*. What we have said in the preceding sections about the semantics of DRSs containing n can be generalized straightforwardly to the DRS-language that has a larger contingent of indexical discourse referents. Suppose that UCC is the set of functions $\{f_1,..,f_k\}$, that $inddref_1,..,inddref_k$ are the indexical

discourse referents corresponding to these and suppose that K is a DRS that contains occurrences of any of these indexical discourse referents. Such a DRS can also be seen in the two ways mentioned: (i) as representation of a particular (actual or possible) utterance u, K determines the propositional content obtained by anchoring the occurrences of each indexical discourse referent $inddref_i$ it contains to the value that the corresponding function f_i returns when applied to u; (ii) as representation of the sentence or discourse type from which it has been constructed, K determines the function from possible utterances to the corresponding contents. One example: As representation of the character of (12) (13) determines relative to $\mathbb{M}$ the function $\lambda t.\lambda x.\lambda w.[[(13)]]_{\mathbb{M},t,x}$. As representation of the content of some particular utterance u of (12) (13) determines the result of applying this function to the entities f_t(u) and f_{sp}(u).

At this point we have shown how the impact of utterance contexts on the semantic content of utterances can be captured through the use of indexical discourse referents. And we have also seen how DRSs containing indexical discourse referents have a kind of double status: they can be looked upon as the semantic representations of particular utterances, which determine their propositional contents, but also as identifiers of the characters of sentence and discourse types. This completes the first main task of the paper.

And with that we also approach the end of Section 2. But there are two further issues I want to bring up before we reach beyond utterance and discourse context to the further components of which Articulated Contexts will be made up. In section 2.8 we have a brief look at the dynamic aspects of the utterance context. It was observed earlier that one important difference between utterance contexts and discourse contexts is that the latter are dynamic – they change as discourse interpretation progresses – whereas utterance contexts are typically static – they don't change in the course of discourse interpretation. But this isn't always so; in some cases utterance contexts also show a certain kind of discourse-internal dynamics.

Second, some readers may have wondered about my use of the term 'indexical', which I have applied on the one hand in relation to certain natural language expressions, such as English *I*, *you*, *now* and the tenses of the verb, and on the other to the discourse referents n, sp, ad, etc. There isn't anything wrong with this 'overloading' of the term. But it may be helpful to say

a little more about the relationship between indexical expressions of natural languages like English and indexical discourse referents. Moreover, in other work[29] I have applied the term 'indexical discourse referent' in a sense that differs from the one in which I have been using it here. Since this second notion will be important in Section 3, it is important to distinguish it from 'indexical discourse referents' spoken of so far. Section 2.9 will deal with these terminological matters and with what lies behind them.

2.8 How static is the Utterance Context?

Perhaps the most salient difference between utterance contexts and discourse contexts, I have been saying, is that by and large utterance contexts are static whereas discourse contexts are dynamic. As we have seen in our discussion of example (6) discourse contexts change with the interpretation of each new sentence of an ongoing discourse. In contrast, the components of the utterance context tend to remain the same throughout the interpretation of a single discourse. The way in which this difference is captured at the formal level of DRS construction is this: When the interpretation of a non-indexical constituent (i.e. one that does not depend for its interpretation on the utterance context) involves the introduction of a discourse referent, the requirement is that always a 'fresh' discourse referent is being used, one that does not yet occur either in the DRS under construction or in the DRS representing the context. (This requirement must be observed even for anaphoric expressions such as anaphoric pronouns; these involve introduction of a fresh discourse referent together with a presuppositional requirement that an antecedent is to be found for it in the context, to which the fresh discourse referent can then be set equal.) For constituents whose interpretation involves the utterance context the processing principles are different. In such cases the connection with the utterance context is established by introducing into the DRS a *new token* of *the one* indexical discourse referent that represents the relevant component of the utterance context. In this way many copies of the same indexical discourse referent can be introduced at different points of the construction of a sentence or discourse representation and all of them represent the same entity. For instance, each occurrence of *I* in an extended monologue will be represented by a token of the discourse referent *sp*. And every interpretation of a main clause tense occurring somewhere in a

[29]See Kamp (1990), Kamp (2003), Kamp et al. (2011)

given discourse will involve the use of a token of n. In the completed DRS all these tokens of sp will represent the same speaker and all the tokens of n will represent the same time. Thus in the course of interpreting a multi-sentence discourse the utterance context is preserved while the context DRS changes with the interpretation of each new sentence, reflecting the dynamic nature of the discourse context.

But the stability of utterance contexts during discourse interpretation isn't absolute. There are discourses that involve shifts in the values of one or more UCC functions. One example are the running commentaries on sports events and other occasions where things happen in rapid succession. Commentators will use the simple present tense to describe the successive events as they occur. Here the production of each new event-describing sentence or clause counts as a separate utterance, with its own utterance time that anchors the particular event which it reports. In such commentaries the utterance time keeps changing. In a DRS for such a discourse the utterance times of the successive present tense statements must be represented by distinct discourse referents, and not by copies of a single indexical discourse referent n. (For some remarks on the dynamics of utterance time see Reyle et al. (2007).)

Another aspect of utterance context dynamics can be observed in conversations, in which the discourse participants alternate as speakers and listeners. Here we need distinct discourse referents sp_1, ..., sp_m and ad_1, ..., ad_m for each of the m participants in the conversation. The correct treatment of dialogue and other forms of conversation within a DRT framework presents a range of challenges, many of which have not yet been met.[30] Dealing with the value shifts of sp and ad that mark each transition between turns is the least of theses challenges.

So much for the dynamics of the indexical dimension of discourse interpretation.

2.9 *sp*, *ad*, *i*, *I* and *you*.

We have spoken on the one hand of 'indexical expressions' – among them are the first and second person pronouns and the tenses of the verb – and

[30]For an approach to dialogue that is akin in sprit to DRT see van Leusen and Muskens (2003).

on the other of 'indexical discourse referents', among them *sp*, *ad* and *n*. We have seen that there is an obvious connection between these two uses of 'indexical'. But let us state the point once more.[31] Consider for instance the pronoun *I* and the discourse referent *sp*. For both reference is determined by the speaker function f_{sp}, one of the bundle that make up the utterance context concept UCC. More precisely, when a DRS is constructed for an utterance containing *I*, then *sp* is introduced in the argument positions that *I* occupies in the uttered sentence or discourse, and when this DRS is evaluated for content, then the function f_{sp} will provide the value for *sp* just as it would provide the value for *I* in a semantics that assigns content directly to the utterance on the basis of the natural language syntax of the expression uttered. Similar connections hold for other pairs of indexical expressions and discourse referents, such as *you* and *ad*, the present tense and *n* and so on.

More interesting is the relation between indexical discourse referents as we have been using the term here and another use that I have made of this term elsewhere. The primary example of an 'indexical' discourse referent in this second sense is the symbol '*i*' as it is used in MSDRT. The '*i*' of MSDRT is used as a constituent of descriptions of mental states, where it functions as a representation of the self of the agent whose mental state is being described. MSDRT will play its proper part only in Section 3 and most of what we need to know about it in this paper will be said there. But the discussion of indexicality belongs here. So we will have to say a little about MSDRT now, in anticipation of the fuller explanation that is to come, in order to be able to explain what is special about *i*.

To explain the special role of '*i*' in MSDRT it is necessary to say something about the role that it assigns to *Entity Representations*. According to MSDRT mental states can be described as sets consisting of (i) Propositional Attitudes and (ii) Entity Representations. A Propositional Attitude is a pair <MOD,*K*>, where MOD is an 'attitudinal mode indicator', such as BEL, to indicate that the attitude is a belief, DES to indicate that it is a desire and so on, and *K* is a DRS that represents the propositional content of the attitude. Entity Representations are triples of the form <[ENT,α],*K*,$\mathcal{K}$>. Here 'ENT' serves to mark the Entity Representation as representation of some

[31] The point is an old one, but it bears repeating here. For essentially the same discussion see (Kamp (1990)).

entity (and thus that it is not a Propositional Attitude), and α is a discourse referent, the so-called *distinguished discourse referent* of the Entity Representation. The distinguished discourse referents of Entity Representations in a mental state can occur in the content representations of its Propositional Attitudes, thereby introducing the referents of their Entity Representations into the propositions determined by those content representations. (On the semantics for such mental state representations these propositions are singular propositions with respect to those referents.)[32]

Argument positions in the content representations of propositional attitudes can also be filled by the discourse referent i. Here the result is also a content representation that determines a singular proposition, one that is singular with respect to the agent to whose mental state the content representation belongs. But in this case the singularity is one of self-attribution, in that irreducible sense of *de se* attribution that has been a prominent theme in the work of for instance Perry and Lewis.[33]

There is an obvious connection between i and the indexical discourse referent *sp*: Both can be used to represent, as part of the semantic representation of an utterance, the contribution made to the content of that utterance by occurrences of *I*. But there is also a crucial difference between i and *sp*. i can *only* occur as constituent of mental state descriptions. This is not so for *sp*. In fact, the one occurrence of *sp* in a semantic representation we have so far encountered, the one in (13), is an illustration of this: (13) is not intended as part of a mental representation. It simply represents the proposition that the speaker of the represented utterance is hungry at the utterance time in a

[32]Entity Representations are much like the 'files' that have become common coinage in current Philosophy of Language and Mind. Further explanations of the structure of Entity Representations will follow in Section 3. For more about Entity Representations and files see also footnote 46 in Section 3.2.1.

[33]See in particular Perry (1980), Perry (2001), Lewis (1979a). Lewis captures the distinction between *de se* attributions that MSDRT represents with the help of i by means of his notion of centered worlds. In this set-up all attitudes are self-attributions, although in many the role of the self is a tacit one, of which there is no overt trace in the way they are most naturally expressed in natural language. This is so in particular for attitudes that in the philosophical literature are usually described as *de re*. For Lewis these are attitudes that are about individuals to which the owner of the attitude stands in some special kind of acquaintance-like relation. In MSDRT *de re* attitudes are those whose content representations contain distinguished discourse referents of Entity Representations.

form that is neutral between, and independent of, the distinct perspectives of speaker and hearer. (Whether (13) could be part of a mental representation in the sense of MSDRT at all is a matter to which we will return in Section 3.)

The symbol 'n' can also occur in the content representations of Propositional Attitudes. An example would be the Propositional Attitude that the speaker of (12) expresses by uttering this sentence. Let us make the plausible assumption that this attitude is a belief. That is, at the time when the speaker utters (12), her mental state contains a component of the form <BEL,K> and it is this belief that her utterance expresses. The DRS K adopted in MSDRT to represent the content of this belief is that given in (14).

(14)

s
n ⊆ s s: hungry(i)

The role of n in this DRS is to indicate the agent's 'psychological now' – it represents the present from the internal perspective of the speaker of (12) at the time when (12) is uttered. So the occurrence of the symbol n in (14), has a status comparable to that of i. In particular, occurrences of n in content representations of thoughts (such as (14)) guide the choice of tenses and temporal adverbs like *now*, *last week* etc. when the one who has the represented thought wants to put that thought into words. For instance, someone who wants to express the thought represented in (14) will choose the present tense (and perhaps add *now* for emphasis).
Since occurrences of n like that in (14) differ from the occurrences of n we have encountered so far in this paper, just as occurrences of i differ from occurrences of *sp*, it would arguably be better to use two distinct symbols rather than the single 'n'. However I have decided nevertheless to stick to the notational overload involved in using 'n' for both purposes, following a practice that, perhaps unfortunately, has by now become rather firmly entrenched. A comparable complacency vis-a-vis the term 'indexical discourse referent' should not be tolerated, however. From now on we will use this term only to refer to discourse referents like *sp* and the earlier occurrences of n. i and n as it occurs in (14) will be referred to as *self-reflexive discourse referents*.

In the remainder of this section we focus on the relationship between i and

the pronouns *I* and *you*. The relation between *i* and *I* was already touched upon above: It is in the verbal realization of occurrences of *i* in the thought that she is putting into words that a speaker can use the pronoun *I*, and (I now add) the pronoun *I* can be used *only* to realize occurrences of *i* in the speaker's mental representation of the thought that she is expressing. The interpretation of *I*, as part of an utterance received, is a different matter. Evidently the recipient cannot use the self-reflexive *i* to interpret *I*. For in his content representations *i* can only play a role which entails that it must refer to him, instead of referring to the speaker, to whom it ought to refer. The only option that the recipient has is to interpret *I* via his Entity Representation for the speaker. (So the representation he will construct will have occurrences of the distinguished discourse referent of this Entity Representation in argument positions corresponding to those occupied by *I* in the utterance he interprets. The recipient's representation of the content the speaker has expressed will thus be structurally different from the speaker's own representation. Note well, however, that in spite of this structural difference the two representations will determine the same propositional content – a proposition that is singular with respect to the speaker.

There is also a systematic relation between *i* and the second person pronoun *you*. When a speaker A uses *you* when speaking to some other person B, then the constituent of the thought she is putting into words won't be *i*, but rather the distinguished discourse referent of the Entity Representation that she has of her addressee, *as* the one she is now addressing). But her use of *you* is an invitation to the addressee to 'self-attribute' what she is saying to him, i.e. to use *i* in all those positions occupied by *you* in the utterance he is interpreting. Here too, then, there will be a structural difference between the speaker's representation of what she says and the representation that her interlocutor will construct from the words she utters. Once again, however, the two representations will determine the same propositional content (a proposition that is singular with respect to the addressee).

The double-sided relation in which *i* stands to the first and second person pronouns is illustrated by even the most trivial exchanges between two speakers in which one uses *I* and the other *you* to refer to same one of them. Here is one arbitrarily chosen example.

(15) a. A: Did **you** put the garbage out?

b. B: I did.

In this mini-dialogue B at first plays the part of interpreter, of the question that A puts to him. In this capacity he represents A's use of *you*, as part of his representation of the question that A has put to him, by means of i. Then, in his answer to the question B – now in his role of utterance producer – asserts the content he has extracted from A's question, as a way of confirming what A has asked, and in his answer he now realizes the i of his representation of the question by the word *I*. In processing B's answer A will interpret his use of *I* with the help of her Entity Representation for him. In this way she will (applying familiar strategies for resolving the VP deletion in B's answer) obtain a content representation isomorphic to her original representation of the proposition that her question was meant to verify.

The comparison of indexical discourse referents and self-reflexive discourse referents helps to see more clearly what the indexical discourse referents *sp*, *ad*, *n* are not. They are *not* constituents of mental states. They occur as parts of user-neutral representations, which capture the contents of utterances, texts and conversations and which can be constructed on the basis of the grammar of the given language and the overt language-relevant properties of the context.

It is important to keep this distinction, between the environments of indexical discourse referents and those of self-reflexive discourse referents, firmly in mind in connection with the central aim of Section 2 – the unification of utterance context and discourse context. Such a unification can be achieved, we have seen, by augmenting DRS languages with indexical discourse referents, which import the contributions that are made by utterance contexts to utterance contents into the DRSs that represent those contents. I stress once more that such DRSs, in which we find occurrences of indexical discourse referents, do not challenge the user-neutral view of meaning that is part and parcel of Formal Semantics as we have known it since the days of Montague and Kaplan. These DRSs can still be seen as abstract structures postulated by a theory which treats semantic properties as properties of autonomous linguistic systems. The only difference with the tradition that was established in the early days of Formal Semantics is that between the logical form approach adopted by DRT and an un-mediated account of semantic values, in which the syntactic structures of the natural language itself are the

only interface that mediates between those values and the natural language expressions themselves.[34]

3. Articulated Contexts

Between them utterance context and discourse context account for a good many aspects of definite noun phrase interpretation; but as we noted in Section 1 by no means for all. Section 3 is devoted to what is missing. We start with another overview of the section, with a little more detail than the one given in Section 1.

Section 3.1 discusses in greater detail the need for a richer notion of context and gives a first rough impression of the information that is missing from the

[34]Only at a very late stage – too late for proper integration into this paper and, I admit with contriteness, solely through my own negligence, – did I become acquainted with two other approaches which present themselves as alternatives to Kaplans three level account of indexicals and come to see their direct relevance to the incorporation of Utterance Context into Discourse Context proposed in the present section. One of these is by Julie Hunter (Hunter and Asher (2012), Hunter (2012)) and the other by Emar Maier (Maier (1998), Maier (2009)). Hunter treats the indexical pronouns *I* and *you* as triggers of identification presuppositions, a treatment she proposes for all definite noun phrases. (A similar unified view of the semantics and pragmatics of definite noun phrases also forms the foundation of the general approach to nominal reference of which some parts are presented in Section 3 of the present paper; the treatment of indexicals that has been sketched in this section can also be cast into that mold.) Hunter secures the correct interpretation of the first and second person pronouns (including their rigidity, in virtue of which the sentences containing such pronouns express singular propositions), via a special constraint imposed on the resolution of the presuppositions they trigger. (In the approach presented here these restrictions would take the form of insisting that occurrences of *I* be resolved through identification with the indexical dref *sp* and occurrences of *you* through identification with the indexical dref *ad*.) Maier accounts for the fact that indexicals make their semantic contributions 'at the level of character by adding a new level to the semantic representations of DRT (its DRSs). In Maier's treatment indexicals are represented at the 'top level of the enlarged DRSs, and this makes it possible for them to behave as terms with variable characters that determine constant contents when applied to the relevant contextual information. Both Hunter's and Maier's proposals are like that of the present paper in that they put greater emphasis on what indexicals have in common with other definite noun phrases and not only bring into prominence what sets indexicals apart. As far as I see, it is not easy to decide which of the three proposals does a better job at keeping a proper balance between what makes indexicals like other definites on the one hand and what makes them different on the other; if a choice has to be made at all, then it will have to be made on the basis of other considerations than those pertaining just to their treatments of indexical noun phrases.

combined discourse-cum-utterance contexts discussed in Section 2.

In Section 3.2 we look more closely at the form of Entity Representations[35], distinguishing between unanchored and anchored ERs and, within the category of anchored ERs, between different forms that anchoring can take.

Section 3.3 is devoted to the methodological status of Articulated Contexts and to their dynamics. It is argued that Articulated Contexts must be seen as components of the mental states of speakers and interpreters, which are directly involved in the choice and interpretation of noun phrases. We then turn to context dynamics. The discourse context of an Articulated Context evolves while text or discourse interpretation progresses; in this regard these discourse contexts are 'dynamic', like the discourse contexts of classical DRT. But there is one important difference: Some of the new elements imported in the course of interpretation into the discourse context of an Articulated Context are transfers from other components of that AC, rather than novel elements which hadn't thus far been part of the context in any way. (As we will see later, the other components of Articulated Contexts are subject to their own kinds of dynamics. But the changes in these components tend to be less frequent and the result of different mechanisms.)

Our main focus throughout will be on language interpretation, just as it has been in earlier versions of DRT. But we will nevertheless also be led, more or less inevitably, to questions concerning language production. The choices speakers make of the noun phrases they use to refer to the things they want to talk about depend on what they know or guess their audiences are able to process. That is, in order to make such a choice a speaker must make assumptions about the interpretational resources that are available to the addressee.

Questions about context dynamics and speaker-hearer coordination arise in connection with all types of definites and also with specifically used indefinites. But different NP types and different uses of those types contribute to context dynamics in different ways. The examples that we will have a closer look at in this paper are for the most part deictic uses of demonstrative

[35]I will often abbreviate 'Entity Representation' to 'ER'. The expression 'ER' (a predicate true of all and only Entity Representations) is to be distinguished from the italicized '*ER*', which we will use as a singular term referring to particular ERs.

NPs.[36] This focus on demonstrative NPs is motivated by one of the central concerns of this paper: a reassessment of Kaplan's account of demonstratives from the communication-theoretic perspective that we will be led to adopt. From this perspective indexicals and demonstrative NPs don't seem to have much in common. (Both give rise to singular propositional content, but that is a property they share with certain uses of other NPs as well.) According to the analysis of deictic uses of demonstrative NPs that will be offered in Section 3.3.1 the reference mechanisms to which they are subject and that are responsible for singular content are very different from those that govern indexicals.

In view of the emphasis that is put in Section 3.1 on the challenge posed by discourse-new definite descriptions and proper names the absence of examples of this kind from the present paper may seem incongruous. But on this point considerations of length were decisive. But in their case considerations about excessive length prevailed.[37]

3.1 Additional Context Components

In Section 2 we saw how utterance context and discourse context can be integrated into a single structure. But whether integrated or kept apart, discourse context and utterance context fail to cover all the contextual information that recipients may need in order to interpret the full range of definite noun phrases in accordance with the intentions of those who use them. The next sections explore what information is missing and propose a way of integrating it into a more comprehensive type of context, which includes the discourse-cum-utterance context as well as other components.

3.1.1 Shared Assumptions on the Basis of a Common Culture or a Common Past

[36] Recall: we are using the term 'demonstrative NP' in a morphological sense, viz. as applying to those noun phrases which begin with *this* or *that* (or that consist of one of those words and nothing more, as in 'That is the one I want'. I am using the term 'demonstrative' (without 'NP') in Kaplan's sense, viz. as denoting the class that includes both demonstrative NPs and indexical NPs.

[37] For an account of proper names within the present framework see Kamp (2015). A detailed treatment of definite descriptions is in preparation.

A notorious problem for the treatment of definite noun phrases in early versions of Dynamic Semantics (including the older versions of DRT) is that the discourse context can't account for the 'discourse-new' occurrences of definite NPs that are ubiquitous in actual conversations and texts. Exactly how this problem should be phrased depends on what further assumptions we are prepared to make. The assumption I am making here has been prominent in the treatment of noun phrases since Heim's development of File Change Semantics in Heim (1982,1988) and is now widely accepted: All definite NP occurrences come with 'identification presuppositions' — presuppositions to the effect that the interpreter must be able to identify, in some appropriate manner, the NP's referent or semantic value. Like other presuppositions, identification presuppositions act as constraints on context: when a presupposition trigger occurs as part of an utterance made in a certain context, then this context ought to be one which enables the interpreter to resolve the presupposition it triggers. It is in this way that definite noun phrases impose constraints on the context in which they are used. And when a definite noun phrase is discourse-new (and not indexical) it is neither the discourse context nor the utterance context that can satisfy the constraint which its presupposition imposes. The occurrence of discourse-new definite descriptions was documented extensively in the wake of the early work on dynamic semantics and was seen as a challenge to those early dynamic accounts.[38] As a matter of fact, proper names pose a similar problem: by definition the first occurrence of a name in a discourse or text is discourse-new, but neither discourse context nor utterance context have anything to offer towards the resolution of its identification presupposition. Linguists seem to have been far less exercised by this fact about proper names than by the high frequency of discourse-new definite descriptions. But discourse-new definite descriptions and discourse-new proper names point in the same direction: Often context-based interpretation requires more than is provided by either utterance context or discourse context.[39]

[38]See footnote 7. The original versions of DRT (Kamp (1981b), Kamp and Reyle (1993)) didn't commit themselves explicitly to the familiarity principle and had nothing specific to say about definite descriptions (apart from the tentative Ch. 3 of Kamp and Reyle (1993)); and the treatment of proper names in those early versions (to the extent that there was any) was little more than a stop gap. But the most straightforward way to extend the coverage of those versions would have followed the lines of File Change Se mantis, and discourse-new definite descriptions would have presented the same problems for those extensions.

[39]Why discourse-new occurrences of proper names should not have been seen as a chal-

For someone unencumbered by linguistic theory there is nothing mysterious about discourse-new uses of definite noun phrases. Such descriptions pose no problems for the interpreter so long as they refer, in a manner he can recognize, to things that are known to him, and that are known to him not because those entities were mentioned previously in the very discourse or text in which they are mentioned now, but because they were known to him already before his interpretation of the discourse or text got under way. We all carry with us large libraries of entities — of people, animals, artefacts, places, past events — and there is considerable overlap between the entity libraries of different speakers.[40] What is more, we are often aware of those overlaps: aware that there are many entities the knowledge of which we share with others belonging to the same social communities or cultural circles that we belong to ourselves. And on a smaller and more personal scale, sharing knowledge of certain entities is also an important part of sharing a life, or a common past, with those conversation partners that are close to us or with whom we are interacting on a regular basis. Thus many of the persons, ar-

lenge in the way one came to see the discourse-new occurrences of definite descriptions may seem a bit of a mystery. I presume that it may have had to do with the circumstance that proper names are subject to principles of reference and interpretation that differ importantly from those that govern definite descriptions, pronouns and demonstratives. There may have been a tacit assumption that identification presuppositions do not play a part in the interpretation of proper names in the way they do for other definite NPs. (The view of proper names outlined in Kripke's 'Naming and Necessity' has no doubt played an important role here.) However, there has been a growing consensus in recent years that proper names too are a species of definite noun phrase, which generate their own type of identification presuppositions. When proper names are classified as definites in this sense, then the discourse-new occurrences of names are no less of a challenge to a theory in which utterance and discourse contexts are the only sources of contextual information than the discourse-new occurrences of definite descriptions.

[40]In Computational Linguistics considerable efforts are made to establish large inventories of entities that authors treat as generally known to the readerships for whom they write. (There is an actual branch of CL devoted to this task, known as 'Named Entity Recognition'. For a survey see Nadeau and Sekine (2007).) Such inventories are indispensable in automatic text processing for purposes of machine translation and other CL applications. They reflect an aspect of our ability to handle language which transcends the "pure" linguistic competence that most linguists take to be the object of linguistic theory. (That these inventories do reflect something other than general linguistic competence is indicated for one thing by a point alluded to earlier: two people who both count as fully competent in a given language (e.g. English) may nevertheless diverge widely with regard to what names are familiar to them.)

tifacts, places etc. we want to introduce into our conversations and mention in the texts we write, are entities of which we can assume that they are known to our audience. And when that is so, using a definite description that contains enough descriptive content to enable the recipient to select the intended referent from those for which he has representations in his entity library is often a good strategy for getting your referential intentions across to him. In fact, it often is the best strategy, the only one that holds a solid promise of success. In such cases using a definite noun phrase to introduce the entity we want to talk about into the text or conversation isn't just an option; it is *the right* way to introduce the entity.[41]

The information that is missing from the contexts of Section 2 isn't only information about the existence of individual entities. Interpretation of definite NPs also often relies on a general understanding of 'how the world is ticking'. Among the NP uses for which such information plays a role the ones most often discussed are so-called 'bridging' uses of definite descriptions. A definite description is interpreted as a *bridging description* if it is taken to present its referent as related in a certain way to some other entity that is independently known or identifiable. A notorious example is *the chandelier* (which for some mysterious reason has made it into the pop charts of the bridging literature): a definite description that can be used to refer to the chandelier in some room that is salient in the context. Or, for another example, the definite description *the cover* can be used to refer to the cover of

[41] In particular, the right way to introduce an entity in such situations is to use a *definite* noun phrase rather than an *indefinite* one. Indefinite noun phrases come with a *novelty* condition, a signal to the recipient that he is *not* assumed to be familiar with the entity to which the speaker or author is referring through her use of the indefinite (see Heim (1982,1988)).
A clarification should perhaps be added: There are (at least) two different kinds of use that speakers can make of indefinites, *specific* and *non-specific* uses. A speaker uses an indefinite *specifically* when she wants to talk about a particular entity she has in mind, but chooses an indefinite in order to convey to her audience that she does not assume this entity to be familiar to them. *Non-sepcific* uses of indefinites are those where the speaker does not have a particular entity in mind, but wants to express that something satisfying the descriptive content of the indefinite fits the predicate of which the indefinite is an argument. The non-specific uses of indefinites are the ones on which logicians and logically minded semanticists have traditionally focused, a preoccupation that has led them to the view that indefinites are the existential quantifiers of natural language.
For a discussion of specific indefinites in a framework close to the one described in this paper see Kamp and Bende-Farkas (2018).

some salient book. And so forth. The task of an interpreter who understands a given description α as a case of bridging is to recover the salient entity to which the referent is supposed to be related, so that he can identify the referent of α as the entity that stands in the given relation to the one he has retrieved. But in order to be able to retrieve the latter entity and be confident that the referent of α *is* the entity related to it in the intended way he has to know that entities like the one retrieved are always, or regularly or at least sometimes, coming in the company of entities of the kind described by α, so that he can infer with reasonable plausibility that the two entities are related in this way. Thus the interpreter of *the chandelier* will, in the normal case, be able to identify some contextually salient room, hallway, auditorium or whatever — in a text this will typically be the location of the current topic of the discourse or at least it will be a room, hallway etc. that has been mentioned in one of the preceding sentences — and then identify the referent of *the chandelier* as the chandelier 'of' that salient space. Knowing that rooms (and hallways, auditoriums and so on) sometimes have chandeliers will make it possible for the interpreter to recover the related entity and relate the referent in the intended way to the entity recovered. It is this kind of general knowledge about what kinds of things can and do occur in the company of which other kinds that the interpreters of bridging descriptions must have ready for use when they are confronted with bridging descriptions.

I propose to keep information of this generic sort — of the actual and possible connections that structure our world, hold it together and make it halfway predictable — separate from information about individual entities. The context component that contains such generic information about the world will be called the *generic component* and denoted as 'K_{gen}'. The component containing information about particular entities will be called the *encyclopedic component* and denoted as 'K_{enc}'.[42])

'K_{gen}' and 'K_{enc}' are two of the additional components that will be part

[42]These names may not be optimal, but they are the best I have been able to come up with. The distinction between 'K_{gen}' and 'K_{enc}' corresponds roughly to that between what in AI has been referred to as the "T-box" and the "A-box": The T-box contains information about the way the world functions — this information takes the form of strict or defeasible conditionals that describe causal and other systematic relations between events, states, individuals and so forth; the A-box consists of "entities", with properties and relations to other entities attached to them. See for instance Nilsson (2010).

of Articulated Contexts. But there is one further component we will need. That component isn't required for all forms of verbal communication, but it becomes essential in connection with NPs that speakers use in face-to-face communication to refer to what they and their interlocutors can see, hear or smell with their own eyes, ears or noses.

3.1.2 The Immediate Environment

Among the uses of definite noun phrases whose referents cannot be recovered from either the utterance context or the discourse context are the so-called *deictic* uses of demonstrative NPs. A *deictic* use of an NP, as I will be using the term, is one in which a demonstrative or other definite NP is employed to refer to some entity from the environment in which the utterance takes place.[43] It is common for the speaker who makes a deictic use of a demonstrative NP to accompany her utterance of the NP with some kind of gesturing towards the referent — e.g. by pointing at it or staring in its direction. (For instance, A says to B: "Do you see that bird on the roof?" while pointing at the bird which she intends as the referent of the NP *that bird on the roof*.) This practice of accompanying deictic uses of demonstrative NPs with some kind of gesture (a "demonstration", in Kaplanian terminology) is indicative of the double function of such uses: on the one hand they are meant to draw the addressee's attention to the referent and on the other they are part of expressing some propositional content about that referent. Kaplan (1989) treats entities from the environment, to which a speaker can refer by using demonstrative phrases in this way, as elements of the utterance context, aiming for a uniform account in which both indexicals and deictic demonstratives are directly referential phrases that get their referents from the context. For the present project, in which the central aim is to develop comprehensive notions of context and context representation in which all contextual information relevant to the interpretation of noun phrases has its place, it appears more natural to treat information about the environment as separate from that which is needed for the interpretation of indexicals. One reason for this

[43] Pronouns and definite descriptions have deictic uses as well. The principles governing those uses are in essence the same as the ones that govern the deictic uses of demonstrative NPs. In this essay I am setting deictic uses of pronouns and definite descriptions aside. (Some incidental remarks on pronouns and definite descriptions can be found in Section 3.3.1 ff.)

is connected with a worry that can be found in Kaplan's own writings: the way in which reference is fixed for demonstrative phrases (and in particular the role that demonstrations play in this) seems quite different from the reference fixing for indexicals like *I* or *now*. We will have more to say about the mechanisms of reference fixing for demonstratives in Section 3.3.1, and there we will find that the environment-related component of the context is subject to a certain kind of dynamics that is unique to it and that sets this component clearly apart from the utterance context, which lacks this kind of dynamics.

The component of the context which contains information about the environment in which communication takes place is called the *environment component* and denoted as 'K_{env}'.

3.2 Articulated Contexts and their constituents: some of the formal details

At this point we have distinguished five different kinds of contextual information: the utterance context, the discourse context, the encyclopedic context, the generic context and the environment context. These are all the components of *Articulated Contexts*, the contexts that we will work with from now on in this paper.[44] Formally, we define Articulated Contexts as tuples of context components. We stick to the unification of utterance context and discourse context defined in Section 2; so utterance context and discourse context form a single component of an Articulated Context. But the other components are kept separate. This means that Articulated Contexts take the form of 4-tuples, consisting of a discourse-cum-utterance context K_{dis} –

[44]Articulated Contexts have been designed specifically for the purpose of NP interpretation. Outside this domain we find other kinds of context dependence (e.g. in connection with vagueness) that Articulated Contexts as they are defined here do not cater for. Theories that deal with other phenomena besides definite noun phrase interpretation (as eventually any semantic theory will have to) can be expected to need more extensive and/or more finely structured contexts than ACs. The term 'Articulated Context' has been chosen to reflect the idea that each component of an Articulated Context has its own part to play in NP interpretations and is subject to its own regime of dynamic development; but in spite of their relative autonomy there are nevertheless interactions between the different components, to the effect that changes in one component can lead to changes in another. Thus, the 'articulated' of 'Articulated Context' is to be understood in roughly the sense in which it is found in the expression 'articulated lorry'.

we will refer to K_{dis} simply as the 'discourse context' –, an encyclopedic context K_{enc}, a generic context K_{gen} and an environment context K_{env}.[45] The official definition is given in (16).

(16) (Def. of Articulated Contexts)

An *Articulated Context* is a 4-tuple <K_{dis}, K_{enc}, K_{gen}, K_{env}>, where

(i) K_{dis} is the representation of the discourse context (with possible occurrences of indexical discourse referents to capture the contributions of the utterance context);
(ii) K_{enc} is a set of representations of "known entities";
(iii) K_{gen} is a set of representations of items of "(generic) world knowledge";
(iv) K_{env} is a set of representations of elements from the immediate environment.

3.2.1 Contents of the new Components

How should the information that goes into the components K_{gen}, K_{enc} and K_{env} be represented? Given the commitments we have already made about the form of discourse contexts it is natural to want to adopt a similar representation format for these new kinds of information, if only to allow for a simpler and more transparent discussion of the information transfers to K_{dis} from other AC components. In this paper I will only make concrete proposals for the format of the Entity Representations that populate K_{enc} and K_{env}. I will have next to nothing to say about content and form of K_{gen}. The problems presented by K_{gen} are very different from those connected with K_{enc} and K_{env}. In the case of K_{gen} it is the content itself that I am uncertain about, and not just the form in which it should be available so that it can support language interpretation effectively. What kind of information should go into K_{gen} depends on questions about the nature of human reasoning that

[45] We already noted that the environment component is relevant only for certain types of discourse. One way to do justice to this would be to define Articulated Contexts either as 4- or as 3-tuples, depending on the communicational setting for which they are needed. Here I opt for a formally more uniform variant, in which Articulated Contexts are always 4-tuples. In those cases where K_{env} is irrelevant, we can set it equal to $\varnothing$.

are outside the scope of this paper and of which my understanding is limited. A proper account of K_{gen} will have to wait.

We assume that there are two kinds of Entity Representations, *anchored* and *unanchored*. Unanchored ERs have the form <[ENT,x],$K, \varnothing$>, anchored ERs are of the form <[ENT,x],$K, \mathcal{K}$>, where $\mathcal{K}$ is some non-empty set of *internal anchors*. In either case x is a discourse referent, the *distinguished discourse referent of* the ER. The distinguished discourse referent of an ER serves as the representative of the ER's referent within the representation to which the ER belongs – exactly what that comes to cannot be explained at this point but will become clear as we go along. K is a DRS which expresses properties of the represented referent. Every unanchored Entity Representation comes with the presupposition that there is a unique satisfier of K. (That is, that there is a unique entity $\mathbf{d}$ such that there is (i) a verifying embedding f of K with $f(x) = \mathbf{d}$ and (ii) no verifying embedding f of K such that $f(x) \neq \mathbf{d}$.) If there is such a $\mathbf{d}$, then $\mathbf{d}$ is the represented entity; if there is no such $\mathbf{d}$, then the representation is defective and does't have a referent.[46]

[46] I already noted that the Entity Representations of MSDRT have much in common with the files that have gained prominent place within the philosophy of language and mind, and those with some familiarity with that literature may wonder why a new term is being used here to refer to a variant of a notion that is so widely known under the name of 'file'. Originally the reason for this terminological deviation had to do with the way in which MSDRT is related to DRT and DRT to Heim's File Change Semantics. The original 1981 version of DRT and FCS make the same predictions (on those data that are covered by both); and they do that in closely similar ways. But one important difference between the two theories – terminological and arguably also conceptual – is that where FCS makes use of 'file cards', DRT makes use of discourse referents. This difference must be distinguished from another difference which can be found between DRT and MSDRT: MSDRT has Entity Representations (as well as discourse referents), whereas DRT has discourse referents only. And while both discourse referents and Entity Representations play the part of representing entities, they are otherwise importantly different, in form as well as in psychological status. By extension MSDRT's Entity Representations are importantly different form FCS's file cards (and even more so from files, which are to be thought of as 'collections' of file cards). Because of all this it seemed misleading to me to refer to Entity Representations as 'files', as that might have suggested a closer similarity to either the file cards or the files of FCS than I think there is; I was concerned to emphasize the differences more than the similarities, to the extent that those exist.

There are, however, many resemblances between the Entity Representations of MSDRT and the file cards that can be found in many philosophical discussions. (Resemblances can be observed at least at an intuitive level; many philosophical characterizations of files stop short of a formal definition of what 'file cards' exactly are like and of the parts that they

Anchored Entity Representations are entity representations which represent by virtue of one or more causal relations in which they stand to the entities that are represented by them. The internal anchors of such ERs are witnesses to these relations. For the time being we limit attention to *singly anchored* ERs, for which the set $\mathcal{K}$ is a singleton. Multiply anchored ERs, for which $|\mathcal{K}| > 1$, will make their appearance in due time.

Like unanchored ERs, singly anchored ERs come with a reference presupposition. But for them the presupposition is not unique satisfaction of K but rather the existence of a unique entity to which the ER stands in the causal relation witnessed by its internal anchor. If there is such an entity **d**, then that is the entity that the ER represents. (In this case **d** is also called the *external anchor* of the ER.) If there is no such entity, then, once again, the ER fails to have a referent.[47]

---

play in mental representations or in the semantics of some natural language or natural language fragment, but use the term as a useful metaphor.) As the present paper has been taking shape, I have felt the need to compare Entity Representations with what various philosophers have said about their respective file notions with a growing urgency. But in tandem with this sense of urgency my awareness has grown of how formidable a task that is. I hope to be able to go some way to meet this challenge in subsequent work. But to meet it within the confines of the present paper seemed practically infeasible, for reasons of time (mine) as well as reasons of space (the paper's).

How daunting the challenge is has impressed itself on me with particular force through Recanati's admirable monograph *Mental Files* Recanati (2012). *Mental Files* contains extensive discussions of how Recanati's own views and assumptions about files relate to those of many other contributors to the file literature. Indeed, a modest start could be made with situating ERs within the file landscape by adding references to various parts of *Mental Files*in the present text, in those places where more will be said about the forms of ERs or about their functions and applications. But even that is a task for which the time is too short at the point in time at which I am adding this footnote. The upshot is that as it stands, the text is without the many references to the philosophical literature that it should have contained. I cannot close this footnote, however, without joining Recanati in acknowledging the central importance of the work of John Perry.

[47]In other work (for instance Kamp (2003), Kamp et al. (2011), Kamp (2005), Kamp and Bende-Farkas (2018)) the form adopted for anchored representations is a slightly different one, viz. '<[ANCH,x],K''>'. In this earlier format the DRS K'' combines the information that in a singly anchored ER (as this notion is defined here) is divided between the 'descriptive' DRS K occurring in second position and the single internal anchor K' that makes up the set $\mathcal{K}$ occurring in third position.

Internal anchors come in different varieties. Not all distinctions between types of internal anchors that have proved useful in various applications are relevant to our concerns here, but some of them are, and our next task is to describe those distinctions. These descriptions, however, are meaningful only against the background of an assumption that is of pivotal importance for the remainder of the paper. This is the assumption that Entity Representations are *mental* representations – that they are constituents of the mental states of cognitive agents.

Recall from Section 1 that the notion of an ER as we are using it here comes from MSDRT, a DRT-based theory of the structure of mental states and of the semantics of attitude reports. In Section 2.8 we noted that according to MSDRT mental states consist of Entity Representations and of Propositional Attitudes <MOD,K>, with MOD the mode indicator of the attitude and K a DRS which represents its propositional content. We also observed that the Entity Representations which belong to a given mental state can enter into the propositional contents of Propositional Attitudes via their distinguished discourse referents and that when the distinguished discourse referent x of an ER ER occurs in the content representation K of a Propositional Attitude <MOD,K>, it thereby renders the content represented by K a proposition that is directly about the entity that ER represents.

A further important aspect of MSDRT is that descriptions of mental states as sets of ERs and Propositional Attitudes can occur as arguments of a predicate '*Att*', which is used to describe states that an agent a is in by virtue of having a mental state that includes components of the kinds enumerated in the description. The Att-predicate provides the theory with considerable expressive force. It allows for a wider range of simple and complex attitude ascriptions and reports.

The formal details of MSDRT will not be repeated here.[48] That imposes certain limitations on the presentation of examples, which we will be able to discuss only at a semi-formal level. I see this as a shortcoming of the present

[48]But see Kamp (1990) for a preliminary formulation of some of the main ideas of MSDRT and Kamp (2003) and Kamp et al. (2011) for a statement of its syntax and model-theoretic semantics. Related ideas can be found in work by Asher that antedates these references (Asher (1986), Asher (1987)). See also the final footnote to the present paper.

paper. But the alternative would have added a good deal to its length and would have made it a good deal harder to read. So in the end I have decided to compromise on this point.

3.2.2 The Form of Internal Anchors.

The internal anchors of anchored Entity Representations should be thought of as testimonies, at the level of the agent's own internal psychology, of certain causal relations in which she takes herself to stand (or have stood) to the entities the ERs purport to represent and of the role that the ERs themselves are playing in these causal relationships. A prime example are ERs that result from direct perception of the entity they represent. But internal anchors cam also represent other causal relations than perception. For what follows here, three types of internal anchors will be central: *perceptual anchors*, *vicarious anchors* and *memory-based anchors*, which are derivative from anchors of the first to types and are connected with these through memory.

First perceptual anchors. A perceptual internal anchor carries the information that the entity represented by the ER of which the anchor is part is being currently perceived. For an example, suppose that somebody is looking at a chair in front of her and that her perception takes the form of a representation of the chair she sees. Then, according to MSDRT, the internal anchor of this ER has the form given in (17).

(17)

s
n ⊆ s
s: i see x

The DRS (17) expresses, from the agent's own perspective at the given time (represented by the self-reflexive discourse referent n) that she (represented by the self-reflexive discourse referent i) is at that time in a state of seeing the entity represented by x.

(17) is a representative example of internal anchors based on perception. All such anchors represent the referents of the ERs to which they belong as currently perceived by the agent. The perception predicates of such anchors are

from some short finite list of sensory perception predicates, which can be assumed to include 'see', but also perceptual concepts like 'hear', 'feel', 'smell'. (I am not committing myself to any particular list; it is left to whoever may want to make use of the MSDRT framework to choose a list that best fits their particular purpose. For a little more on the specifications of internal anchors see footnote 51.)

Second, vicarious anchors. Vicarious anchors arise when an agent B takes some other agent A to make reference to an entity, and forms on the strength of this an ER for that entity which represents it *as* the entity to which A has just made reference. In our discussion here we confine attention to those cases in which the presumed reference by A takes the form of her using a definite noun phrase α to refer to the entity in question. I assume that the vicarious anchor that B forms on such an occasion – as part of a newly formed ER or else as an addition to the anchor set of an ER that he has formed on some previous occasion – has a form like that shown in (18). (There have been some disputes over the form of vicarious anchors. Here we will make do with a simple proposal, according to which the vicarious anchor identifies the referent of the vicariously anchored ER (as represented by the ER's distinguished discourse referent x) as the entity referred to by the speaker's through her use of a definite NP α. We use a 4-place predicate 'refer' for this purpose.)

The predication 'e: refer(a,α,x' means that a speaker a performed an utterance e in which she used the definite NP α to refer to the entity x. This predicate enables us to represent vicarious anchors in the following simple form:

(18)

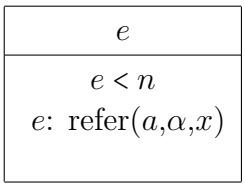

When (18) plays the part of a vicarious anchor, then x is the distinguished discourse referent of the Entity Representation of which this is a vicarious anchor. The effect of such an anchor is to stipulate that the Entity Representation it anchors represent whatever it is that the speaker has just referred

to by using α.[49] It is one of the assumptions of MSDRT that a speaker can refer through the use of an NP only when she has an Entity Representation for the entity she refers to. On this assumption a vicarious anchor can be justified only when the speaker does have an Entity Representation for the referent of her use of α. The vicarious anchor of the interpreter's Entity Representation then renders it coreferential with this Entity Representstion.

The first of the two types of annchors we have discussed was that of anchors based on current perception. An agent has an ER with such an anchor while she is observing the entity that this ER represents. For the most part, however, our perceptions are of short duration. We look at a thing, but then we turn our eyes to something else and that is the end of this particular perceptual event. But it need not be the end of the ER that has been formed on the basis of the perception, and as often as not it isn't. We still have our representation of the entity – at first almost as if we were still perceiving it – even though our immediate perceptual contact with it has ceased. The ER survives when the actual perceptual contact that gave rise to is severed.[50]

If what we have been saying about ERs that are based on current perception is taken in the strict sense in which it is intended, then an ER of this formal characterization cannot last beyond the perception that gave rise to it in the exact form in which it was first established; for the state of affairs described by the internal anchor of such an ER – that of the agent *currently* perceiving the entity – no longer holds. In other words, if we want to maintain that perception-based ERs often survive the perceptions that bring them into existence, then we must allow for the possibility that their anchors change, to reflect the shift of the perception from the psychological present to the psychological past and the fact that the referent is now accessible to the

[49](18) fails to make explicit that the speaker 'just' made her utterance, and not at some time or other in the past. At this time there still sin't a proper representation within DRT for such vague notions of temporal proximity as are expressed by gthe word *just*.

[50]Much the same applies to perceiving things by sound or smell – when we hear a thing and then the sound it is making stops, or we redirect our hearing to some other sounds; or when we perceive a thing with our noses but then the smell dissipates, or we get used to it and don't notice it any longer as something that stands out as distinct. Also, much of what can be said about perception by sight when there is light can be said about perception by touch in the dark. In all such cases we still know that we just perceived the thing, we assume that it is still where it was when we perceived it (or near to where it was); and our representation of it is much as it was when we actually were perceiving it.

agent herself only through memory. For instance, the internal anchor (17) should be minimally changed into (19).

(19)

s t
t < n t ⊆ s ¬(n ⊆ s)
s: i see x

The change from (17) to (19) is minimal in that it does no more than relocate the perceptual interaction from the psychological present represented by n to the psychological past, represented by the time t in the past of n. But often the relocation will be more specific (involving conditions, say, that indicate more narrowly where in the past the perception occurred by relating its time to those of other events that are represented in the agent's memory). This is one respect in which ERs with memory-based perceptual anchors can vary. But there are also other differences; memory-based perceptual anchors cover a range of quite different cases. At the one end there are the situations in which the agent has stopped looking at the referent because his attention has just been caught by something else, but where she takes it for granted that the entity is still where she saw it (or has moved from there over a more or less predictable distance in a more or less predictable direction) and that if she turned her eyes back to the place where it was when she perceived it, or to the place to which she expects it to have moved, her perception of it would be restored. At the other extreme there are ERs that have been part of the agent's mental state for quite some while after the time when the initial perception of the represented entity gave rise to their formation. These ERs may have changed from what they were at the time of their initial inception not only in that their anchors locate the initiating perception in a more distant past; they may also have changed with regard to the properties and relations that they attribute to the entities they represent; that is, they may also have changed in their second, descriptive component. If there are reasons to assume that the entity represented by an ER has lost some of the properties recorded in the descriptive component K of that ER, then the conditions expressing those properties, which were part of the ER's descriptive component at the time when it was formed, will have been past-shifted too.

An extreme case of how the descriptive content of an ER can change over time arises when the agent suspends her belief that the entity still exists.

Another respect in which ERs are subject to change in the course of their tenure as constituents of our mental states is that their anchoring sets may grow as time goes on, with successive encounters between agent and referent leading to additional anchors that witness those later encounters. (For more on this see the next section.)

There is considerable scope for further discriminations among Entity Representations that started out as ERs with current perception-based anchors. But this is an aspect of memory-based anchors that in the present paper we can afford to keep very simple. Besides anchors based on current perception we distinguish just one comprehensive category of anchors based on memory of past perceptions; and we assume that the content representations of such anchors can always be obtained from the content representations of anchors based on current perception through the transformation illustrated by the transition from (17) to (19), with the new anchor saying that the perceptual relation which anchored the ER initially was holding at some past time but doesn't hold any longer at the current psychological now.

Vicariously anchored ERs too are liable to change after their initial formation. But here the situation is a little different, since the events that give rise to vicarious anchors are already past at the time when these anchors are formed. (That the realization event e happened in the past is explicitly encoded in our vicarious anchor schema (18) by the condition '$e < n$'.) When the formation of a vicarious anchor has become a thing of the past, that condition remains as it was. (It is of course possible that additional information is added about the temporal location of the event, as noted above.) But there is another part of the schema in (18) which also concerns temporal location and that part does require modification when the event that gave rise to the vicarious anchor formation recedes further to a more distant past. This is the attribution, at n, to the utterer of the triggering noun phrase α that at n she is in possession of an Entity Representation for the entity that she referred to by using α. When enough time has passed after the vicariously anchored ER has been formed, the assumption that the speaker still has such an Entity Representation may no longer be supported. Put more simply: There may no longer be any reason to assume that the speaker still remembers the referent that she referred to at the time when she used α. The most dramatic form this can take (but it is one that is by no means exceptional) is when the possessor of the vicariously anchored ER has come to believe that the

person whose utterance gave rise to his formation of the ER is no longer alive.

Such changes in what the possessor of a vicariously anchored ER assumes about the person who used α on the occasion that gave rise to its vicarious anchor will lead to replacement of the condition '$n \subseteq s$' in the vicarious anchor by a pair of conditions '$t \subseteq s$' and '$t < n$' for some fresh temporal discourse referent t. But note well: taken by itself this replacement doesn't contradict the condition it replaces; the replacing conditions are compatible both with the condition '$n \subseteq s$' and with its negation. A more drastic change of the vicarious anchor is one that combines the mentioned replacement with the negation of the old condition, in analogy with what we have seen in (19).)

This is all I will say about single anchors. More could be said, but this is already more than what is strictly needed for the remainder of this paper. In the next section we look at some aspects of multiply anchored Entity Representations.

3.2.3 Recognition and Multiple Anchoring

When you encounter something you have never encountered before – a person standing on a street corner that you cross regularly on your way to work but whom you have never noticed on previous occasions, or a new tree that has just been planted somewhere along your route; the choice of examples is wildly arbitrary – then you form a new perceptually anchored ER for what you notice, and, starting as soon as you have stopped looking at the thing and walked on, the anchor of this ER turns into one based on memory (unless the ER is lost as soon as it has been created and no representation of the entity is ever imprinted in memory). But what happens when you encounter an entity for the second time? Suppose that the next morning you see the man standing on that corner again, and you recognize him as the man who was also standing there yesterday. What is the net effect of that in terms of Entity Representations?

Here is my story. Recognition of an object on a given occasion consists in retrieving an Entity Representation for this object from your entity library and linking it to the current experience you have of the object it represents. This linking takes the form of adding a new, perception-based anchor to the

Entity Representation you had.

In this way Entity Representations can grow over time, collecting new anchors with each new encounter of the objects they represent. It was the need of a way to account for recognition of known entities in terms of adding new anchors to their ERs that led to the decision to take the third components of ERs to be sets of internal anchors rather than single anchors, as had been assumed in earlier work. For the most part, when an Entity Representation is first created it will have no more than one internal anchor (although we will find some exceptions to this when we discuss deictic uses of demonstrative noun phrases in Section 3.3.1). But when the represented object is encountered again and recognized as the one represented by an Entity Representation already in one's possession, then that will have the effect of augmenting the anchor set of that ER with a new anchor, as a witness from the new encounter.

(20) is our official definition of the notion of an Entity Representation and of the distinction between unanchored, singly anchored and multiply anchored ERs. (This is nothing more than a repetition of what has already been said, but now couched in somewhat more formal terms and distinguished by its own label.) The definition is followed by a few observations about two ways in which things can go wrong with an ER and the adverse consequences that can have.

(20) (Definition of *Entity Representation* and of Entity Representations that are *unanchored, simply anchored and multiply anchored*)

(i) An *Entity Representation* is a triple of the form

$$<[\text{ENT},x],K_{descr},\mathcal{K}_{anch}>,$$

where x is a discourse referent, K_{descr} is a DRS and $\mathcal{K}_{anch}$ is a set of anchor-DRSs.[51]

[51] By an *anchor-DRS* is understood a DRS whose vocabulary is restricted to a special, limited DRS language, the 'internal anchor formalism', which is tailor-made to the task of expressing the various kinds of information that can go into internal anchors. No definitive proposal for this DRS language exists at the present time, though in applications we have

(ii) An Entity Representation ER is *unanchored* iff $\mathcal{K}_{anch} = \varnothing$; otherwise ER is *anchored*. ER is *singly anchored* if $|\mathcal{K}_{anch}| = 1$ and *multiply anchored* if $|\mathcal{K}_{anch}| > 1$.

Not all Entity Representations are proper. Every Entity Representation comes with a 'groundedness presupposition'. Only when this presupposition is satisfied can an Entity Representation perform its proper function, that of representing some particular entity (its 'referent'), on the basis of which it can make the contributions it is intended to make to content representations belonging to the same mental state. In case of failure there is no entity that is represented by the ER and all other representations in which the distinguished discourse referent of the ER occurs will be vacuous or meaningless. For details see e.g. Kamp (2003) or Kamp et al. (2011).

Groundedness presuppositions differ depending on whether the Entity Representation is unanchored, singly anchored or multiply anchored. For unanchored Entity Representations the presupposition is that their descriptive component identifies a unique satisfier.[52] For singly anchored Entity Repre-

made use of certain causal predicates, guided by an informal concern for inter-application consistency. For the purposes of this paper, pinning down the details of a special DRS-language for anchor contents is not essential. See also the remark on the use of predicates 'see', 'hear', 'smell' etc in the Legenda to the vicarious anchor shown in (18).

[52]One question that might be raised in connection with unanchored Entity Representations is whether it is the totality of all descriptive information associated with such a representation that should have a unique satisfier or only some particular part of it: If that part has a unique satisfier, then the Entity Representation might be regarded as properly grounded even if some of the remaining information isn't true of its satisfier. On this interpretation, the Entity Representation does properly represent the 'unique satisfier' but at the same time carries some false information about it. I haven't explored the possibility of such divisions within the descriptive information of unanchored Entity Representations, with one part serving as unique satisfaction condition for the groundedness presupposition and the other as 'contingent' information about the referent. When such divisions are made once and for all – when they are an intrinsic part of the structure of the ER description that allows for such divisions – then it would be natural to assume such ERs to have a separate component, their 'descriptive anchor', and to retain the second components of ERs as defined in (20) for the descriptive 'extras' that have no influence on whether or not the ER refers and on what it refers to. But it is also conceivable that the division is flexible and fleeting, with a capacity for *ad hoc* adjustment, reminiscent of the Cluster Theory of names, as first proposed in Searle (1958). In fact, a version of the

sentations groundedness is satisfied iff there is an entity to which the Entity Representation and its possessor stand in the causal relation specified by its internal anchor. Groundedness failure irises when here is something fundamentally amiss with the conditions presupposed by the anchor. Typical cases of failed perceptual anchors are those that result from visual illusions: the possessor of the ER may think he is seeing something in a certain location, but there is nothing there, his impression is caused by a 'trick of the light' and not by the object that he takes himself to be seeing. There are other sources of failure besides these, but for present needs this much will do.[53]

Multiply anchored Entity Representations can be defective for the same reason as singly anchored ERs: one or more of the internal anchors of a multiply anchored Representation can fail in one of the ways just discussed. But a multiply anchored ER can also be defective in another way. It may be that none of its internal anchors are defective as such, but that different internal anchors link the ER to *different entities*. This is a violation of a different presupposition that an Entity Representation ought to satisfy: all its anchors should point to the same referent. The most common situations that give rise to this form of defectiveness are cases of misrecognition: you take an individual or object you encounter to be the one for which you already have an ER but that ER has a referent **d** which is distinct from the entity **d'** that is actually encountered. In such situations, I am assuming, recognition takes the form of adding to your ER a new anchor, based on your current perception of **d'**. Since the new anchor links the ER that now contains it in its anchor set to **d'**, the modified ER is defective even when the original ER wasn't. But there are other cases as well. You may take a speaker to refer to the entity represented by an Entity Representation ER of yours, though in fact she is referring to some other entity. What you do in such a situation is the same that you would have done in cases where there aren't such

cluster theory might be developed on the basis of descriptive ERs that are subject to a regime of ad hoc divisions of their descriptive information while at the same time being 'labeled ERs' in the sense of Kamp (2015). As it is I do not have any compelling examples that involve shifting divisions of descriptive content and I have no idea what the principles might be that govern such shifts, in case shifting does occur.

[53]When an illusory perceptual experience is over, the defective ER to which it gave rise will be transformed into a memory-based ER, in the manner described above. The transformed ER will then of course be just as defective as the original one. In fact, most defective ERs are like this: they are the memory-based transformations of ERs that were defective at the time they were formed.

referential conflicts: you add a vicarious anchor to your Entity Representation ER as a witness to the (misconstrued) referential act of the speaker. This anchor will then point to an entity that is different from the one determined by the anchor or anchors that were part of ER's anchor set already.

Multiply anchored Entity Representations that are defective because of conflicting anchors are called *incoherent* (and those that are free from this defect *coherent*). The formal specifications of these notions are given in (21).

(21) (Soundness and Coherence of Entity Representations)

Let $\mathbb{M}$ be a model, w a world from $W_{\mathbb{M}}$ and t a time from $\mathbb{M}$ and letER = <[ENT,x],K_{descr},$\mathcal{K}_{anch}$> be an Entity Representation of an agent a existing in w at t. (That is, a belongs to the Universe of $\underline{\mathrm{M}}_w$ and exists in $\underline{\mathrm{M}}_w$ at t.) Let us assume that the causal relations between a and his environment (as captured by $\underline{\mathrm{M}}_w$ at times leading up to t) determine an external anchoring relation EXT between some subset $\mathcal{K}'$ of $\mathcal{K}_{anch}$ and the Universe of $\underline{\mathrm{M}}_w$. Then
(i) If ER is unanchored (i.e. $\mathcal{K}_{anch} = \varnothing$), then ER is *sound with respect to* $\mathbb{M}$, w and t iff there is a unique entity **d** in the universe of $\underline{\mathrm{M}}_w$ which satisfies the DRS K_{descr} in w at t in $\mathbb{M}$.[54]

(ii) An anchored Entity Representation ER = <[ENT,x],K_{descr},$\mathcal{K}_{anch}$> is *sound with respect to* $\mathbb{M}$, w and t iff every anchor A from $\mathcal{K}_{anch}$ belongs to $\mathcal{K}'$ and links ER via EXT to some entity **d** in $\underline{\mathrm{M}}_w$. Otherwise ER is *unsound with respect to* $\mathbb{M}$, w and t.

(iii) A multiply anchored Entity Representation ER = <[ENT,x],K_{descr},$\mathcal{K}_{anch}$> is *coherent with respect to* $\mathbb{M}$, w and t iff there is a particular entity **d** in $\underline{\mathrm{M}}_w$ such that every non-defective anchor A from $\mathcal{K}_{anch}$ that belongs to $\mathcal{K}'$ links ER to **d**. Otherwise ER is *incoherent with respect to* $\mathbb{M}$, w and t.

Incoherence of multiply anchored Entity Representations raises some difficult questions that I have not yet been able to explore in sufficient depth. The

[54] **d** satisfies K_{descr} in w at t in $\mathbb{M}$ if there is a verifying embedding f of K_{descr} in $\underline{\mathrm{M}}_w$ at t such that $f(x)$ = **d.** and $f(n)$ = t

central question has to do with incidental corruptions of such ERs. Suppose there is someone – let's call her Mary – who you know very well, but haven't interacted with in recent times. You have a multiply anchored Entity Representation of Mary, based on the many occasions the two of you have met in the past. Yesterday you saw her in the street, but she turned the corner before you could say hallo, and you didn't want to call after her. Or rather, this is what you think. In reality it wasn't Mary whom you saw, but some person that from that distance looked very much like Mary. In the light of what I have been saying about Entity Representations this situation is to be described as follows. When you saw the person, you added to your existing Entity Representation for Mary a new (perception-based) anchor, thus linking that ER to the person you saw, and thereby rendering the ER incoherent. But this may well be a case of incoherence that won't matter very much in practice. You still have your solidly anchored Entity Representation for Mary, almost as it was, with the old anchors and all the descriptive information about Mary that you acquired during the previous encounters of which those anchors are the witnesses. In fact, there is a good chance that you will soon forget about yesterday's incident and that your Entity Representation for Mary will return to its previous coherent condition. Or, alternatively, you may realize your mistake, thereby removing the new anchor from your ER for Mary. And even while the corrupting anchor is part of the ER, no harm is likely to come of that. The new information that you associate with your representation of Mary on the strength of your illusory encounter with her today won't amount to much – you hardly saw the person you took to be Mary – and so your image of Mary and your knowledge about her will hardly have been affected. It is only when the misrecognitions repeat themselves and accumulate that real trouble is brewing. Suppose you run into the person whom you already mistook for Mary on a previous occasion once more. If you take her to be Mary again – perhaps on the basis of the previous encounter on which you mistook her for Mary, that will make the corruption of your ER worse, and harder to disentangle in case you get wise to your mistake. And when you keep doing this, then eventually your ER will be corrupted 'beyond recognition'. Such an ER, with an anchor set that is more or less evenly divided between anchors that point in opposite directions, is truly incoherent and dysfunctional: it cannot function as the representation of a single referent.[55] It is also possible for

[55]It has been pointed out to me (Casey Woolwine, p.c.) that this is an overstatement.

the process of adding anchors pointing to the new entity to go on for so long that the new anchors (and the information associated with them) become dominant and eventually obliterate the old anchors and the information associated with those. In such cases it may be reasonable to see the ER as having transmuted into a representation of the new entity.[56]

Even when an ER has an incoherent anchor set $\mathcal{K}_{anch}$ that is evenly divided between a subset $\mathcal{K}_1$ of internal anchors that link the ER to an external anchor $\mathbf{d}_1$ and another subset $\mathcal{K}_2$ of internal anchors that link the ER to a different external anchor $\mathbf{d}_2$, it may still be possible for the ER to function more or less properly in the way it is supposed to. Here is one example. It might be that the possessor A of the Entity Representation ER always encounters the individuals $\mathbf{d}_1$ and $\mathbf{d}_2$ in distinct types of situations – for instance, $\mathbf{d}_1$ always as a famous guitarist, whom A often hears perform on the classical music station she regularly listens to and $\mathbf{d}_2$ as the composer of film scores who is among the credits of some of the films she sees and whose name appears in film reviews she reads. Somehow A has come to believe that $\mathbf{d}_1$ and $\mathbf{d}_2$ are one and the same person (for one thing because they actually happen to have the same name). So A has erroneously formed an Entity Representation ER which conflates $\mathbf{d}_1$ and $\mathbf{d}_2$ into a single fictitious person $\mathbf{d}$, the presumed bearer of, on the one hand, properties that have been acquired in situations which left behind an anchor belonging to $\mathcal{K}_1$ and that are true of $\mathbf{d}_1$ and on the other hand of properties acquired in situations yielding anchors belonging to $\mathcal{K}_2$ and true of $\mathbf{d}_2$. Let us assume, moreover, that the de facto separation of the distinct settings in which A encounters $\mathbf{d}_1$ and $\mathbf{d}_2$, via A's favorite classical music station and her viewing of and information about films, respectively, continues and that it is only through encounters of the first type that A will become aware of new properties of the fictitious referent $\mathbf{d}$ that are in fact true of $\mathbf{d}_1$ and through encounters of the second type that she becomes aware of properties that are in fact true of $\mathbf{d}_2$. Under these conditions it could be said that ER, while embodying A's belief that there is just a single individual $\mathbf{d}$ who possesses all the properties in $\mathcal{P}_1$ and all those in $\mathcal{P}_2$, nevertheless keeps what is true of $\mathbf{d}_1$ neatly separate from what is true of $\mathbf{d}_2$. In a case like this it is arguable that when A has another encounter with $\mathbf{d}_1$ (which by assumption will be an encounter of the first, classical music station, type), the role that ER plays in this event is to link A's current exposure to $\mathbf{d}_1$ with her existing information about $\mathbf{d}_1$, and likewise for encounters (of the second, film-related, type) with $\mathbf{d}_2$. In other words, ER tracks both A's encounters with $\mathbf{d}_1$ and her encounters with $\mathbf{d}_2$, much as would have been done by distinct ERs for $\mathbf{d}_1$ and $\mathbf{d}_2$. (To fully capture the way in which an Entity Representation like this one can include a record of the connections between parts of the descriptive component K of the Entity Representation and those internal anchors in its Anchor Set that witness the encounters when these parts were introduced into K more machinery will be needed.)
This is just one example of the countless forms that ERs and ER incoherence can take. (Another example is the often discussed case of the philosopher Aston-Martin discussed in Donnellan's 'Proper Names and Identifying Descriptions', see Donnellan (1970).) When incoherence interferes with the intended functions of ERs, and to what extent, is a topic in its own right, which I am leaving for further exploration.

[56]Is it possible to draw a well-motivated line between intuitively harmless contaminations

On the account I have outlined multiply anchored ERs have anchor sets in which each anchor witnesses a distinct event in which the agent is confronted with the ER's real or presumed referent and in which each such confrontation has left its trace by contributing its own external anchor. For entities with which we interact on a regular basis this would entail that their anchor sets would become very large and keep growing; and that seems quite unrealistic. The next section discusses this issue, without arriving at a solution.

3.2.4 Anchored Representations without Internal Anchors?

What, speaking intuitively, are our representations like for the entities with which we are most closely familiar and about which the strongest case can be made that we should be able to entertain singular thoughts about them — our next of kin, the objects we see and handle every day, the house or apartment we live in, the places we pass every day on our way to and from work. What kinds of representations do we have of these people and things? Continuing the train of thought we have been following so far, we might be led to the conclusion that our representations for such entities must be extremely complex, and increasingly so as a new anchor added for each new encounter. But how plausible is that? You and I know each other extremely well, we have met on countless occasions, we have been through all sorts of experiences together. When we see and talk to each other again, in what sense will that lead to my adding yet another anchor to the representation I have of you, or for you to add a new anchor to your representation of me?

When the man I saw standing on the corner of Elm Street and Main Street on my way to work yesterday is there again today, and I recognize him as the man I saw yesterday, then the assumption that this involves my adding a new perception-based anchor to the ER that I formed when I first saw him seems plausible enough. But what if the experience repeats itself day after

of anchor sets and anchor sets that are corrupted to the point of reference failure? To make any headway with this question one would have to look more deeply into the internal structure of the anchor sets of multiply anchored ERs, and perhaps also into their history. Note that the discussion above raises much the same questions that are discussed in considerable detail in Evans (1985), in relation to the cases of *Madagascar* and the 'old geyser'. I would like to think that multi-anchored ERs give us a new handle on such questions. But working out the details is a project for some other time.

day? Do I add a new anchor to my ER for each and every time I see him? That somehow feels much less plausible. Memory just doesn't keep count that way. Somehow, one can't help thinking, all those different encounters get increasingly blurred into one, making the ER into the representation of 'the man who is standing on the corner of Elm Street and Emerson Boulevard when I pass by there on my way to work.'. (But note that this isn't an unanchored ER, which refers by virtue of unique satisfaction of the quoted description. Rather, the description is meant as a shorthand for a kind of 'super anchor' which bundles together an unspecified number of links, all of them based on perceptual causation, and all pointing to the same external anchor.)

In relation to this particular example a case could perhaps be made for an anchored representation with an internal anchor that quantifies over perceptual experiences — an anchor which describes the represented entity as the man with whom I have entered into a perceptual relation each time I passed the corner of Elm Street and Emerson Boulevard and saw him. But with many of the people and things that are closest to us no such readily articulable quantifications will be available. They are encountered in all sorts of situations, and in nearly all of those we recognize them without even paying attention to the fact that we do — they are just there, as we know them. Presumably there is no single concept that subsumes all the different ways in which we apprehend them, or one that subsumes all the different occasions on which that happens. Even in such cases it may be true that each new encounter adds some new anchoring support, linking the Entity Representation I have had for some time ever more firmly to its referent; but there won't be any general formula that captures what all these contributions have in common.

If these speculations are on the right track, then many of the ERs in our entity libraries are neither singly nor multiply anchored in the sense defined above, but have anchoring components of a more 'generic' form. This will be true in particular of many of the ERs we make use of when interpreting definite NPs. As things stand, I have no viable proposal for what such anchors are like. So in what follows we will have to make do with the types of ERs we have defined explicitly in (20). But the matter of 'generic' anchors has to be addressed, and sooner better than later.

3.2.5 Entity Representations as constituents of Articulated Contexts: Some methodological implications

When we introduced the notion of an Articulated Context we motivated this richer concept of contextual information by referring to the problem of discourse-new definites in texts. This was part of the motivation for two of the new components of articulated contexts, K_{enc} and K_{gen}. And then there was the last component, K_{env}, which plays a part only in face-to-face communications about things in the mutually accessible environment. So far not much has been said about the contents of these components of Articulated Contexts. Now that we have said more about the structure and varieties of ERs we are also in a position to say more about the context components that have ERs as constituents, viz. K_{enc} and K_{env}.[57]

Before we go into more detail about K_{enc} and K_{env} there is a general issue that needs addressing first. Entity Representations are constituents of mental states. But if that is so, then K_{enc} and K_{env} must also be regarded as mental categories. And once that has been admitted, we have to face a further question: What are we to make of Articulated Contexts as wholes, with all their different components, given that two of those components consist of mental entities? The answer I will propose – it will be stated more explicitly in the next section – is that Articulated Contexts should be seen as properties of mental states in their entirety: all components of Articulated Contexts are to be treated as features of the mental states that the relevant agents are in at the relevant times.

The present section explores the methodological implications of a mentalistic interpretation of Articulated Contexts for a theory of linguistic meaning. Note that we have reached a crucial choice point. The choice we are facing is that between (i) treating linguistic contexts as aspects of the minds of language interpreters (and producers) and (ii) treating them as characterizations of the objective, user-independent situations in which interpretations of linguistic utterances and texts take place. Both these conceptions go back a long way, at the very least to the time when Montague laid the foundations

[57]Once again, the contents of K_{gen} are of a different sort. They are propositional representations of some sort, not ERs. The explorations of the last few sections throw no new light on this AC component. But K_{gen} has already been declared off limits in this paper.

of formal semantics as we know and practice it today. Montague unambiguously opted for the second view, and the conception of utterance context that can be found in his work became the received view of what kind of context is needed in the denotational natural language semantics he bequeathed upon us. But it was Kaplan's analysis of the implications of such a concept of context that was decisive for the wide currency that this concept has had, both within semantics and in the philosophy of language.

A challenge from within the formal semantics community came with the advent of File Change Semantics and DRT. As it was originally presented, DRT allowed for two different ways in which it could be interpreted, which correspond to the two conceptions of context I just mentioned. Since the choice that confronts us here is closely related to those two possible interpretations of DRT, let me give a brief summary of the choice as it presented itself at the time when DRT received its first explicit formulation.

Early versions of DRT could be taken either as offering a method for describing the semantics of natural languages that treats languages as independent systems, with their inherent syntactic and semantic properties, or as attempts to relate the semantics of natural languages more directly to the psychology of their users. From the first perspective the aims and methods of DRT are much the same as those of Montague Grammar, the only major difference being that DRT is a 'logical form' theory, which describes form-meaning relations by articulating a method for assigning logical forms (DRSs) to expressions of the given natural language and where it is these logical forms that endow the expressions to which they are assigned with semantic values, determined via the model-theoretic semantics for the representation language.[58] On the second view DRT is a theory of *interpretation*, a theory of how language users process the utterances that reach them and the texts that they read. On this view DRSs are models for the mental representations that constitute utterance and text understanding, and DRS construction is a model for the procedures that interpreters follow when they build those representations.

[58]There is of course also the dynamic dimension of DRT that sets it apart from Montague Grammar: DRSs play the double role of content representations (of the sentence or sentences from which they have been constructed) and of discourse contexts (for the interpretation of what comes next in the discourse or text). But important as this difference may be, it is not essential to the issues that confront us right now.

Note well: these two views need not be incompatible. Suppose the combination of a DRS construction algorithm for a given language or language fragment and a model-theoretic definition of how DRSs determine semantic values is seen as a description of the language as a user-independent system; and suppose it would then be concluded (on the basis of evidence relating to human language processing, say) that form and construction of DRSs provide useful models of what goes on in the minds of human interpreters. There is no obvious reason why that conclusion should be regarded as ground for abandoning the first view, according to which the theory captures properties of the language as autonomous system.

However, once we have committed ourselves to the position that the linguistic contexts on which utterance and text interpretation depend can only be understood as parts of mental states, such a eclectic position is no longer available to us. Such a commitment forces us to view any account of meaning that makes use of such contexts as indissolubly tied to the mental interpretation processes that make use of these contexts.

This is the first major conclusion to which the adoption of ERs as resources for the interpretation of discourse-new definites leads us. But does this mean that all ties with a user-independent semantics have been severed? No. Recall how we got to the point at which we find ourselves right now: Our turn towards a psychological account of meaning was brought about by our acknowledgement that discourse-new definites are interpreted with the help of ERs, which on our account are mental constituents. But the role that ERs play in our account of the contributions that definite NPs make to text and discourse meaning is that of aiding the reference resolution for those NPs. And reference resolution of definites is a special case of something more general, viz. the resolution or justification of presuppositions of any sort. In Kamp (2001a), Kamp (2001b) it is argued that presupposition justification belongs to a processing level of linguistic input that comes after the construction of a preliminary logical form (a so-called 'preliminary DRS' when DRT is assumed as general semantic framework).[59] The construction of pre-

[59]The observations of Kamp (2001a) and Kamp (2001b) build on the work of Van Der Sandt and Geurts (Van Der Sandt (1992), Van Der Sandt and Geurts (1991), Geurts (1999)), which develops a explicit account of presupposition within DRT. Kamp (2001a)

liminary DRSs can be seen as a strictly linguistic process which is guided by principles that belong to the language as an independent system (i.e. principles that are part of the 'grammar' of the language' in a narrow sense of the term). The preliminary DRSs that are the outputs of this process contain explicit representations of the presuppositions that are contributed by the various presupposition triggers contained in the linguistic inputs from which they have been constructed (including representations of the identification presuppositions generated by definite noun phrases). But the represented presuppositions are still awaiting resolution.[60] What happens up to this point – the computation of a syntactic representation for the input and the construction of e preliminary DRS from this syntactic representation – can still be regarded driven by the grammar of the language as an autonomous system.

Presupposition resolution is a process that operates on preliminary DRSs, and takes place at a second processing level, at which various sources of contextual information can be accessed; and when successful, it will convert the preliminary DRS into a DRS from which the presupposition representations are gone and that can serve as input to pragmatic processes such as, for instance, various forms of Gricean reasoning. On this (essentially Van-Der$_S$andtian) view of presupposition presupposition processing involves two successive stages, which are markedly different. This makes it possible to come up with a novel answer to the question whether 'presupposition is part of semantics or part of pragmatics'. The answer is 'both and': The

and Kamp (2001b) differ from this work in that they put a strong emphasis on the question how the logical forms of presuppositions are constructed from the syntactic structures of the sentences that contain their triggers, and not only on the problems of presupposition resolution and accommodation that have been the central concerns in most discussions of presupposition in the formal semantics literature. (For an authoritative somewhat older survey of the formal presupposition literature see Beaver (1997) and Beaver's own contributions to the topic of presupposition in Beaver (2001).)

[60]The terms 'presupposition justification' and 'presupposition resolution' relate in the first instance to, respectively, presuppositions with a purely propositional content (e.g. the presupposition that Mary was sick on some earlier occasion or occasions triggered by *again* in the sentence 'Mary was sick again') and 'anaphoric' presuppositions (e.g. those triggered by third person pronouns and other definite NPs), which require finding an antecedent in the discourse context for some particular variable or discourse referent. But the terms are often used interchangeably. In this paper I will from now on only use the term 'resolution'.

first half of presupposition processing as we have just described it, in which preliminary DRSs are derived from syntactic analyses, is part of semantics, the second half, in which the presupposition are resolved and the preliminary DRS is converted into a completed ('non-preliminary') DRS, can be regarded as part of pragmatics.

The conclusion that presupposition resolution must be part of pragmatics isn't literally forced upon us. But it is forced upon more or less. We have already conceded that many cases of presupposition resolution require contextual information that is essentially mental in nature: in the way in which we have set things up Entity Representations are essentially mental in nature. And we have also been claiming that Entity Representations (and thus with them the mental states of which they are part) are essential ingredients in the resolution of certain presuppositions. In other words, mental states, and therewith the agents whose mental states they are, enter into this stage of presupposition processing. So the language user – the agent who does the interpreting of a sentence or sentence sequence – is essentially involved in the account of presupposition resolution to which we have been led. So, if we accept the widely held principle that those linguistic phenomena are pragmatic in which the language user plays an essential role (see e.g. Morris (1946)), then for us presupposition resolution must be part of pragmatics.

Drawing the line between semantics and pragmatics in this way – as dividing two aspects of presupposition interpretation which in other presupposition accounts are often not explicitly distinguished, may be unusual, but I contend that it makes perfectly good sense and that as far as presupposition is concerned it is probably the right place to draw to draw it. But equally important, and more directly relevant to the methodological reflections we have been pursuing, is the distinction between those parts of the description of the phenomena that can be thought of a descriptions of an autonomous linguistic system and this that cannot. From all we have been saying so far, treating the computation of the logical forms of presuppositions as an aspect of autonomous grammar remains a possibility. But treating presupposition resolution in this way is not. For us presupposition resolution is an unequivocally mind-related process.[61]

[61]The picture that is emerging from our outline of definite NP interpretation is closer to the traditional conception of the semantics-pragmatics divide than DRT in its original

For the remainder of this essay this distinction between abstract, user-independent linguistic processes and those that have an essential mental component and thus involve the processing language user, will not be important. We will take the stance that all linguistic processing is processing by discourse participant, even where a more abstract description of a process, in which no reference is made to an inter peter or producer would be possible in principle.

So much for the dividing line between semantics and pragmatics, and we return to the question of user-independent accounts of linguistic meaning and approaches that account for linguistic meaning in terms of interpretation as a mental process.There is a further aspect to the user-neutrality of meaning and interpretation that I have not yet mentioned. The DRSs that result from preliminary DRSs through the resolution of identification presuppositions are by and large like the DRSs from earlier DRT versions, but they may differ in that they contain occurrences of the distinguished discourse referents of ERs

form – or at least according to early presentations DRT by some of its advocates (myself included). One of the preeminent claims that were made in those presentations was that the line between semantics and pragmatics needed to be redrawn. Discourse anaphora, a phenomenon that had traditionally been classified as belonging to pragmatics because it transcends the bounds of the single sentence, should be reallocated to semantics, the argument went, since it could not be separated from the sentence-internal manifestations of anaphora, which had generally been seen as part of semantics. The argument relied on the treatment that DRT offered of certain anaphoric phenomena, which accounts for sentence-internal and trans-sentential anaphora in the same way. (The alternative would have been to treat some or all sentence-internal anaphoric relations as pragmatic too. To my knowledge that alternative option was never seriously explored. But plausible or not, that too might have required some adjustment of the semantics-pragmatics border.) Within the over-all architecture of the present proposal this argument loses its force. Discourse anaphora now involves, like all other instances of definite noun phrase anaphora, two stages: (i) that of the construction of a preliminary DRS with its explicit representations of anaphoric presuppositions and (ii) that at which the presuppositions of that preliminary DRS are resolved. One half of the interpretation process, the construction of explicit representations for the identification presuppositions triggered by anaphoric NPs, belongs to semantics; the other half, the resolution of these presuppositions, is part of pragmatics. Note well: One consequence of this is that sentence-internal resolutions of identification presuppositions will now also qualify as pragmatic. But that is not unreasonable, for even such resolutions will often require plausibility reasoning, to arrive at the identification of the intended antecedent for the pronoun from among a number of logically possible antecedents; and often such reasoning makes use of extra-linguistic assumptions (of the sort that in our set-up belongs to K_{gen}).

(those ERs that were used in the resolution of these presuppositions). The content that is expressed by such a DRS K is a proposition about the 'real referents' of those ERs (i.e. about their external anchors): a proposition that is singular with respect to each of the occurrences of those distinguished discourse referents within K. In an intensional model $\mathbb{M}$ in which the universe of the actual world w_0 contains the referents of the ERs whose distinguished discourse referents occur in K this proposition can be defined in terms of verifying embeddings into those models $\underline{\mathrm{M}}_w$ whose universes also contain these referents, viz. by restricting the possible embeddings to those which map each distinguished discourse referent in K to the referent of its ER.[62] This gives us a user-neutral characterization of the content that the interpretation assigns to the linguistic input, even though the process that leads to this interpretation can only be described in psychological terms.[63]

3.2.6 Articulated Contexts as Parts of Mental States

Here is where we have got. I have argued that Entity Representations are often essential to the interpretation of definite NPs. I have also mentioned (though not yet argued, that's still coming) that the interpretation of definite NPs – the resolution of their identification presuppositions – often involves the interaction of different components of Articulated Contexts. And Entity Representations, as we have defined them, are mental entities. That means that part of the story we want to tell about the interpretation of definite NPs that part in which Entity Representations play an essential part, will have to bring the mental a states of which those Entity Representations into play. But since the story will involve interactions between ERs and other AC constituents, these other AC constituents will need to be given a psychological interpretation as well. Otherwise the story cannot be told.

[62] Typically, of course, this definition will yield a partial proposition, since there will be worlds in $W_{\mathbb{M}}$ in which not all of these referents exist.

[63] For some people even the user-dependent account of presupposition resolution that is part of the over-all architecture I have sketched may be more than they are willing to swallow. But what alternative could there be? One possibility might be to develop a non-psychological notion of entity representation, so that all components of Articulated Contexts can be treated as non-mental. This would be a very different enterprise from the one we are engaged in here, with its heavy reliance on MSDRT. I do not want to exclude this option categorically, but I have no idea how it might be realized. If there is a way of proceeding along such lines, I will be grateful for any tip as to how this might be done.

This means that we need a psychological interpretation of Articulated Contexts in toto: What are Articulated Contexts as parts or aspects of mental states? Perhaps not everyone would want to answer this question in the same way. But for us, who have been led to the assumption that Articulated Contexts are parts of mental states because the many of their constituents are ERs and ERs are components of mental a states as described in MSDRT, the only natural way to proceed with the question is this: How can we describe the different components of Articulated Contexts in MSDRT terms?
We start with K_{enc} and K_{env}. We have already made the commitment that each of these consists entirely of Entity Representations. So, since mental states are composed of ERs and Propositional Attitudes, the natural conclusion must be that both the K_{enc} component and the K_{env} component of the Articulated Context of a given agent a at a given time t are subsets of the set of ERs that make up the entity representing part of a's mental state at t.

However, apart from this common feature K_{enc} and K_{env} are quite different.

First K_{enc}. I have spoken of the elements of K_{enc} more than once as 'items from the agent's entity library'. In fact, that is all that can and needs to be said: In principle, any element of an agent's entity library – any ER belonging to the agent's mental state – can be used to resolve the identification presuppositions of definite NPs; and because that is so, any ER from the agent's mental state should be included among the members of K_{enc}. In other words, at any point in time the K_{enc} component of an agent's AC may be identified with his entire library at that time, i.e. with the set of all ERs belonging to his mental state.

Whereas any Entity Representation qualifies as member of K_{enc}, there are strong restrictions on the kinds of ERs that can be members of K_{env}. The members of K_{env} are Entity Representations for referents that are currently being perceived (or that have just been perceived) in the situation in which verbal communication is taking place. Such Entity Representations have anchors which testify to the current or recent perception of their referents. The little that I have said about the forms of anchors does not provide us quite with what is required for a formal characterization of this constraint. Needed is a formal notion of short time perceptual memory – of that kind of memory in which the perceptual experience one has just had is still lingering, as it were, and is accompanied by a tacit conviction that the experience could

be reactivated if one took the trouble of turning one's attention back to the object that one was just perceiving.

Suppose we have a form of internal anchor that captures this kind of interrupted but lingering perception (as distinct from other kinds of explicit or implicit memory of perceptual experiences). Then we could define the members of K_{env} to be those ERs that either have an internal anchor testifying to a current perception of their referent or else an anchor that testifies to this kind of lingering of a perception in memory. I won't try to determine the form of anchors of this second type here, but simply assume that it is by virtue of their form that certain memory-based ERs qualify as members of K_{env}.[64]

The other two AC components, K_{dis} and K_{gen}, raise different issues. K_{gen} has already been declared *terra incognita* and it would be inconsistent with this earier resolve to try and visit this territory now. As a compromise, the following footnote provides a brief aerial exploration.[65]

[64]One step towards determining the form of internal anchors that capture the sense of a lingering perception could be this: (i) The descriptive content of anchors that witness current perceptions is assumed to not only mention the perceptual event (as involving the perceptual relation between agent and referent) but also information about where the agent locates the referent in relation to himself through his perception of it. When such an anchor is then transformed into one that testifies to the lingering memory of the perception event, the description of the perception can be past-shifted along the lines shown in the transition from (17) to (19), while the information about the (stable or shifting) location of the referent will be retained as holding currently. Surely this is no more than a first approximation, but it goes some way towards what is special about the kind of memory-based Entity Representations that can occur in K_{env}.

[65]K_{gen} has been described as a collection of propositions that express general connections between things, states and events within our world. If K_{gen}-constituents are elements of mental states as these are defined in MSDRT, then (since obviously these constituents are not entity representations) they should presumably take the form of some sort of Propositional Attitudes – pairs <MOD, K> with some Mode Indicator MOD and a content-representing DRS K. I already hinted at one aspect of the content representations of items in K_{gen}: the contents that need to be represented are defeasible conditionals and statements of compossibility. The representation of such contents is intimately bound up with the use that can be made of their representations in defeasible reasoning. In the absence of a substantive account of how defeasible conditions behave in defeasible reasoning there is no point in trying to be precise about their form and so we won't try. The mode indicator of the attitudes belonging to K_{gen} must be some kind of doxastic operator, presumably a special kind of belief operator which captures the kind of non-

K_{dis} raises different questions. Its content is propositional in nature. (In this respect it resembles K_{gen}.) So if K_{dis} is to find its place among the constituents of a mental state in the sense of MSDRT, it too should be among the Propositional Attitudes and thus of the form <MOD,K>. About the content part K nothing needs to be added to what has been implicit in the things we have been saying in relation to the role and content of discourse contexts: K_{dis} is the discourse context. In the DRT-based setting we have adopted the discourse context is constructed from the input discourse or text. It has the form of a DRS (and its content is the content this DRS has according to the model-theoretic semantics for the given DRS language to which it belongs). Although we are assuming that this construction proceeds differently from what was assumed in earlier versions of DRT, with each successive sentence interpretation now going through the two phases of preliminary DRS construction and presupposition resolution, the final result will always be a DRS. So that is the general form of a K_{dis}.

But what could be the mode indicator for K_{dis}? In this case choosing some kind of doxastic operator once and for all would not be right. Even if we restrict attention to utterances with the status and force of assertions, whose purpose it is to communicate propositional information to the interlocutor, such communications do not always achieve the supposed effect of getting the recipient to accept this propositional content as true and storing it among his beliefs. The recipient may understand perfectly well what he is being told, but may have reason to doubt the truth of it or even judge it to be false – doubt and rejection presuppose understanding no less than acceptance.; all such judgments presuppose that the content of what the speaker is saying has been correctly identified. It is this attitude, which consists in having identified the propositional content of a verbal input but without any further commitment of acceptance, rejection of vacillation, that the mode indicator of K_{dis} should stand for. Let us introduce a label for a mode indicator that does precisely this and refer to it as 'CCD' (for Content of Current Discourse), and add it to MSDRT's mode indicator repertoire. This enables us to assume that the K_{dis} component of the AC of agent a at time t will appear as one of the Propositional Attitudes of the mental state of a at t in the form <CCD,K_{dis}>.

contingency that we attach to information of this sort.

Pairs of the form <CCD,K_{dis}> are propositional attitudes more in name than in substance. A genuine propositional attitude with K_{dis} is established only when the interpreter forms a judgment about the truth of the proposition expressed by K_{dis}, e.g. by accepting or rejecting it or by remaining neutral on the matter. At that point a 'genuine' Propositional Attitude will be formed, with BEL as mode indicator in case of acceptance and some other Mode Indicators in cases of doubt or rejection.[66] These genuine attitudes will typically outlive the interpretation process that yields the construction of their content representation. But for the 'Propositional Attitude' <CCD,K_{dis}>. this is not so. In fact, we assume that ACs do not last beyond the episodes of discourse and text processing for which they are needed. When such a process starts it activates an AC, which recruits most of its information from the mental state that the interpreter is in at that moment. This AC will have an empty K_{dis}, which will be gradually transformed into the final representation of the discourse has been constructed. At that point the AC ceases to exist as a unit, but its components will continue as parts of the interpreter's mental state – as part of his entity library, as some of his Propositional Attitudes (those that make up his non-episodic world knowledge) and as some Propositional Attitude with the final K_{dis} for its content representation. As the interpretation of a multi-sentence discourse proceeds, the discourse context DRS K_{dis} changes with each sentence. Note well, however, that the intermediate K_{dis}s may become the content representations of genuine Propositional Attitudes as soon as they have been constructed. For instance, it is possible

[66]Rejection of the proposition expressed by K_{dis} could be expressed as belief in the negation of that proposition. But such a reduction is not possible for all Mode Indicators that are relevant here. An agent can be in a state of perfect uncertainty vis-à-vis a proposition, assigning equal subjective probabilities to it and to its negation. (Such an attitude may be accompanied by an urge to find out whether or not the proposition is true, but it need no be.) We could use a new mode Indicator PUN for this attitudinal mode (PUN = 'Perfect UNcertainty'). In earlier work I have used WON as name for an Indictator representing a mode close to this one– WON being short for 'Wonder (whether)'. But the verb *to wonder* has a strong connotation of wanting to find out the truth, so if WON does justice to the choice of its name, it should be distinguished from PUN.
The attitudinal mode expressed by the verb *to doubt* has a good deal in common with both PUN and WON. It too expresses a state of suspension somewhere between acceptance and rejection. On the whole it tends more towards denial than to confirmation. Because of these two features it would not be right to identify it either with the mere absence of belief nor with outright rejection. So doubt is another attitude that cannot be defined in terms of belief.

for an interpreter to accept the K_{dis} of the first part of a discourse as one of his beliefs, but then baulk at the remainder of the discourse and form a negative attitude (one of 'disbelief') towards the content represented by the K_{dis} for the complete discourse.[67]

3.3 Dynamics of Articulated Contexts

Now that we have settled on some of the structural properties of Articulated Contexts we are in a position to explore their dynamics: How do ACs change over time, and how are these changes brought about by the interpretation processes which they support and that define their life spans?

Change can affect each of the components of an Articulated Context. The most obvious and prominent changes are those in the discourse context K_{dis}. The K_{dis} component invariably develops as the interpretation of an utterance, discourse or text proceeds. In this regard K_{dis} is like the DRSs of earlier versions of DRT (such as that of Kamp and Reyle (1993)). But there is nevertheless a difference with these older versions. In the older versions all there is to the dynamics of discourse contexts is that they grow as a function of progressing discourse processing. The discourse context components K_{dis} of Articulated Contexts also grow, and in much the same way, but as they grow they may take on board information that was already present elsewhere

[67] It has been suggested to me that it might also be possible, and preferable from a more traditional conception of formal semantics of natural language, to go in the opposite direction: abstract away from the psychological dimension of Entity Representations and turn them into constituents of contexts that guide the semantic evaluation of expressions of the object language as an autonomous formal system. In such a theory Entity Representations would still be playing a decisive part in the evaluation of definite noun phrases, but now as constituents of such abstract contexts. In this way the communication-theoretic principles for noun phrase interpretation – some of which we have discussed already and more of which will follow below – would become principles that can be considered part of the grammar of this autonomous formal system.
I have made some efforts to work out the details of such an alternative, but not with much success. Too much of what makes the communication-theoretic approach adopted in this paper of any interest has to do with the potential and actual differences between the Entity Representations available to the speaker and those available to the hearer. It is hard to see how this aspect of the approach could be plausibly preserved in a 'user-neutral' theory, which treats syntax and semantics of the language as determined independently of any use that might be made of it. So far my efforts at working out the details have reinforced my suspicion that this is not a useful way to go.

in the Articulated Context – some of the bits that get integrated into the discourse context are *transferred from other components.* [68]

It will be useful to make this a little more concrete. Let us assume that when a definite NP is interpreted with the help of some Entity Representation ER belonging to K_{enc} or K_{env}. That is, the NP's identification presupposition is resolved by identifying the NP's referent with the entity represented by ER. This will take the following form: (i) A discourse referent β is introduced to represent the referent of the NP (β is called the *referential argument* of the NP). (ii) β is inserted into the argument slot that corresponds to the one occupied by the NP in its sentence or clause, and is also added to the universe of the DRS whose Condition Set contains the predication involving that slot. (iii) The condition '$\beta = \alpha$' is added to this Condition Set, where α is the distinguished discourse referent of ER.

[68] This feature of Articulated Contexts – that their dynamics can take the form of transferring information from one component to some other component – was the central initial motive for wanting to develop the concept of an Articulated Context (in a talk contributed to a workshop which took place as part of the XVII-th International Congress of Linguists, Prague, 2003). The idea was that certain expressions, in particular 3rd person singular pronouns, have access only to the discourse context – it is only discourse referents made available by K_{dis} that are available as antecedents for such pronouns. But, the consideration was, this doesn't mean that entities that one cannot be referred to by means of pronouns at a given stage of the discourse must be *unknown* to the interpreter. He may know them perfectly well, but even so, before a pronoun may be used to refer to them, they must first be turned into 'referential objects of the discourse context'. That is, a representation of such an entity must first be added to the discourse context and that is something which can be effected by first referring to it using some other type of definite noun phrase, such as a proper name or a definite description.

In the opening parts of Section 3 of the present paper I motivated the need for a comprehensive context concept like that of an Articulated Context by pointing out that utterance contexts and discourse contexts do not suffice for the interpretation of discourse-new occurrences of definite descriptions and proper names. This is of course a different phenomenon from pronoun interpretation and the restrictions to which it is subject. But in fact ACs provide us with the means to account for both phenomena and for how they fit together. At the start of a discourse the K_{dis} component of the articulated context is empty, but the other components need not be and typically they are not. That is why the use of a pronoun in the opening sentence of a discourse comes across as marked and requires some kind of accommodation. But names and definite descriptions can be used without sounding marked or infelicitous; and that is because they can be interpreted by exploiting some other AC component, which is not empty at the start of discourse processing. (Most often this is K_{enc}, but as we will see in Section 3.3.1, it can also be K_{env}.)

The upshot of this is that whenever an ER from K_{enc} or K_{env} is used in the referent identification of an NP, then its distinguished discourse referent α will have a 'free' occurrence in K_{dis} (viz. its occurrence in the equation '$\beta = \alpha$'); this is the form in which NP interpretation transfers information from K_{enc} or K_{env} to K_{dis}.[69] For an example consider an utterance of the sentence 'Mary slept' that is interpreted by someone whose mental state (and thus the K_{enc} of his AC) contains an Entity Representation <[ENT,x], K,$\mathcal{K}$> for the person that the speaker is referring to as Mary and that the interpreter is using this ER to interpret this use of 'Mary'. The construction of the DRS for this sentential utterance requires choosing a referential argument for the NP *Mary* (i.e. a discourse referent to represent the referent of the NP) inserting this referential argument into the argument slot of the DRT-predicate 'sleep' that represents the concept expressed by the verb *sleep* and adding it to the Universe of the DRS. Let us assume that the chosen discourse referent is y. Using the mentioned Entity Representation in the interpretation of the NP *Mary* takes the form of setting its referential argument y equal to the distinguished discourse referent x of the ER, by adding the condition '$y = x$' to the condition set of the DRS. The result is shown in (22)

(22)

e t y
t < n e ⊆ t e: sleep'(y)
y = x

(In this example the discourse referents x and y play the part of, respectively, α and β in the schematic description given above.)

K_{dis} is not the only AC component that is subject to change. The other components can change as well. The causes and modes of change in K_{gen} are a story in their own right and since K_{gen} plays little part in our considerations anyway, I will make no attempt to tell it. But the ways of and reasons for change in K_{enc} and K_{env} are directly relevant to the role they play in noun phrase processing and thus deserve our attention. The next section and much of what comes after it will be devoted to the use of K_{env}, an important part

[69] This is not the only way in which we can implement the use of ERs from K_{enc} or K_{env} for the resolution of definite NP presuppositions. But any implementation must produce the effect that the distinguished discourse referent of the ER has one or more free occurrences in K_{dis}.

of which is its dynamics. In the remainder of this section we will look at the dynamics of K_{enc}.

K_{enc}, we noted, can be regarded as the agent's entity library. Libraries change as new items are added to them and old ones are lost or removed. So too it is with our mental entity libraries. When we encounter entities that are unfamiliar, new ERs are added. On the other hand old ERs may disappear from them (or be moved to some remote part of the stacks where they are no longer accessible). Among the situations in which we are led to add new ERs to our libraries are those in which we encounter unfamiliar entities through the words of others. Let us once more restrict attention to cases where some other agent refers to some entity through the use of an NP, which the given agent then represents by forming a new ER for it. The 'canonical' instances of this interpretation strategy are those in which the NP is an indefinite of which the speaker makes a specific use and that the recipient recognizes as having been used specifically. As argued in Kamp and Bende-Farkas (2018), this leads to the formation of a vicariously anchored ER for the referent that the recipient takes the speaker to talk about. This ER is a new addition to the recipient's entity library and thus to his K_{enc}.

The formation of new ERs in response to indefinite NPs is something to be expected in the light of the *Novelty Condition*. The Novelty Condition is a signal that accompanies the use of all indefinite NPs and that indicates to the recipient that the utterer of the indefinite does not expect her audience to be able to identify what she is talking about on the basis of the descriptive content of the NP in combination with information he already has. In this respect indefinite NPs are the opposite of definite NPs, which – as we have been assuming all along – come with a signal to the audience that it is expected to be in a position to identify the referent. With specific uses of indefinite NPs (as opposed to non-specific ones) the Novelty Condition takes on a special significance: the producer does not expect her addressee to be able to identify the particular entity she is talking about in using the NP – in our terms: that he has no ER for this entity, for which she herself does have an ER, which was the basis for her use of the NP. For the addressee this is an invitation to introduce an ER for the entity in response to the speaker's use of the NP, with a vicarious anchor that secures the entity as referent for the ER. In Kamp and Bende-Farkas (2018) we assume that the interpretation of an indefinite that the interpreter takes to have been used specifically always

leads to the formation of a vicariously anchored ER.

As said, definite NPs are subject to the Familiarity Condition: The use of a definite NP signals that the speaker assumes the addressee to have the means for identifying the referent; that is – again, in our terms –it signals that the addressee is assumed to have an ER that he can recognize as the representation of the NP's referent and that he can therefore use to resolve the identification presupposition which the NP contributes to the preliminary DRS for the sentence containing it. If the addressee is in a position to resolve the identification presupposition, then there is neither need nor room for the formation of a new ER. But of course, the producer can be wrong in assuming that her addressee has an ER that he can use to resolve the presupposition. If he doesn't, then all can do is accommodate the presupposition. And accommodation will take the form of introducing a vicariously anchored ER for the presumed referent.

For many the term 'accommodation' carries the connotation of a last resort operation. But in the case of certain definite NP types this connotation is hardly present if at all. One might have thought that if definite noun phrases come with identification presuppositions, producers would avoid them whenever they have reason to think that what they want to refer to is unfamiliar to their interlocutors. But in fact, it is quite common for a speaker orb author to use a definite without being particularly concerned whether the audience is familiar with its referent; what is more, producers will use definite NPs also in situations where they know that the referent is unknown to their audience. In such situations the audience will usually play along by constructing the ER they are lacking. In this way they acquire new information through the back door, as the matter is sometimes put. And the speaker has made a conscious use of the back door to get the information across that way. We will encounter such cases, in which the speaker or author forces her audience to accommodate an ER and may do so on purpose, in Sections 3.3.4 - 3.3.6.

In all cases we have considered the newly created ER will have a vicarious anchor. Usually this will be its only anchor, but that isn't so invariably. In the next section we will see that it often isn't so for the ERs which are created in response to deictic uses of demonstrative NPs.

3.3.1 The Use and Dynamics of K_{env}

Our next topic is the dynamics of K_{env}. What I want to say about this is best conveyed by looking at a couple of examples. We will be looking only at a handful, the choice of which is motivated by the question how deictic uses of demonstrative NPs differ from indexicals. Most of our examples will be deictically used demonstratives.

Interpreters can accommodate unfamiliar definites by creating vicariously anchored Entity Representations for their referents. That applies in particular when the definite is a proper name (see Kamp (2015) for details) and when it is a definite description. It also applies to deictic uses of demonstrative NPs. But there is an aspect to the deictic uses of NPs that appears to be unique to them. Our first example illustrates what I have in mind.

Suppose that a speaker S utters (23), using the demonstrative *that bird on the roof over there* to pick out an entity from the environment shared between her and her addressee H and claiming of this bird that it is a blackbird.

(23) Look at that bird on the roof over there! It is a blackbird.

(23) can be felicitous irrespective of whether the addressee already has an Entity Representation of the referent for the demonstrative phrase that (23) contains and irrespective of whether S assumes that H has such a representation. Let us first suppose that H doesn't have such a representation yet. Then the likely impact of S's utterance, perhaps accompanied by a gesture in the direction of the bird she is referring to, can be described as follows. The effect of the first, imperative sentence will be that H looks in the right direction, sees the bird, forms a perceptually anchored Entity Representation of it, adds this representation to the K_{env} component of his Articulated Context and also forms the beginnings of a K_{dis} by introducing a discourse referent y for the bird he has just registered in its Universe. This K_{dis} then enables him to construct a representation for the second sentence of (23), using y to interpret the pronoun *it*. H also adds a vicarious anchor to his ER for the bird, as a testimony to the reference to the bird that S has made through her use of the phrase *that bird on the roof over there*.

The scenario for (23) in which H's attention is drawn to the bird through the S's words and the one in which he had noticed the bird beforehand have

much in common. In either case H will end up with an ER that is doubly anchored, via a perceptual anchor to the bird he sees and via a vicarious anchor that reflects the reference that S has just made to it. The only difference is that in the second case a perceptually anchored ER for the bird was already in place before the utterance of (23) and in the first case such an ER is created *ab ovo*, in response to S's utterance.

Let me try to make this a little more concrete by describing in some formal detail how the different components of H's AC change as he interprets the first and second sentence of (23). We only consider the first scenario, in which H has not noticed the bird on the roof until S' utterance of (23) draws his attention to it.

Suppose that before S has said anything H's AC is as in (24)

(24) <∅, $\mathrm{K}_{enc} \cup \{ER_a\}$, K_{gen}, $\mathrm{K}_{env} \cup \{ER_a\}$>

Here ER_a is H's ER for the speaker S. It is assumed (i) that a is the distinguished discourse referent of this ER and (ii) that this ER belongs both to H's K_{enc} and to his K_{env}. (Nothing much hangs on this second assumption.)

We assume that the effect on S's uttering the first sentence of (23) has the following effects:

(i) H forms a perceptually anchored ER ER_b for the bird he now sees:

(25) ER_b: $\left\langle [ENT, b], \boxed{\begin{array}{|c|}\hline \\ \hline \mathrm{bird}(b) \\ \hline \end{array}}, \left\{ \boxed{\begin{array}{|c|}\hline s \\ \hline n \subseteq s \\ s\text{: } i \text{ see } b \\ \hline \end{array}} \right\} \right\rangle$

Next, H adds a vicarious anchor to the anchor set of this ER and also forms a rudimentary K_{dis} in which there is a discourse referent for the bird represented by this ER, as indicated by its being set equal to its distinguished discourse referent b.

(26) ER'_b: $\left\langle [ENT, b], \boxed{\begin{array}{|c|}\hline \\ \hline \mathrm{bird}(b) \\ \hline \end{array}}, \left\{ \boxed{\begin{array}{|c|}\hline s \\ \hline n \subseteq s \\ s\text{: } i \text{ see } b \\ \hline \end{array}}, \boxed{\begin{array}{|c|}\hline e \\ \hline e < n \\ e\text{: refer}(a, \text{`Bird'}, b) \\ \hline \end{array}} \right\} \right\rangle$

where 'Bird' is short for 'that bird on the roof over there'.

At this point H's AC looks like this:

(27) $\left\langle \begin{array}{|c|}\hline y \\ \hline y = b \\ \hline \end{array}, K_{enc} \cup \{ER_a\}, K_{gen}, K_{env} \cup \{ER_a, ER'_b\}\right\rangle$

Note well: I have not presented a preliminary representation for this first sentence, with an identification presupposition for the demonstrative NP *that bird on the roof over there*. But the various moves described – the formation of ER_b and its modification to ER'_b and the formation of K_{dis} are all part of the resolution of this presupposition. (This gives a flavor of how much will be involved in spelling out the resolution principles for the identification presuppositions of certain types of definite NPs.)

For the second sentence of (23), however, it will be useful to start with a preliminary representation, in which the identification presupposition for the pronoun *it* is explicitly represented.

(28) $\left\langle \left\{ \begin{array}{|c|}\hline z? \\ \hline \text{non-human}(z) \\ \hline \end{array}_{3d.p.pr} \right\} , \begin{array}{|c|}\hline \\ \hline \text{blackbird}(z) \\ \hline \end{array} \right\rangle$

Resolution of the presupposition in (28) must take the form of finding an antecedent for the queried discourse referent z? and this antecedent must be found in the Universe of the discourse context K_{dis}. This is possible, since K_{dis} contains the discourse referent y and that is the antecedent that we want.

Resolving z to y has the formal effect that z is everywhere replaced by y in the non-presuppositional DRS of (28). After this substitution has been carried out, the presupposition of (28) is removed and the non-presuppositional DRS is merged with K_{dis}. The new K_{dis} is as in (29); and after H has interpreted the two sentences of (23), his AC is as in (30).

(29)
$$\begin{array}{|c|}\hline y \\ \hline y = b \\ \text{blackbird}(y) \\ \hline \end{array}$$

(30) $\left\langle \begin{array}{|c|}\hline y \\ \hline y = b \\ \text{blackbird}(y) \\ \hline \end{array}, K_{enc} \cup \{ER_a\}, K_{gen}, K_{env} \cup \{ER_a, ER'_b\}\right\rangle$

I opted for the phrasing of example (23) because it has a natural ring to it. However, in this particular form the example has the drawback that its first sentence is an imperative. There still is no generally accepted way of dealing with non-indicative utterances in DRT. And so the treatment of H's interpretation of this first sentence, which leads him to adopt ER'_b and a K_{dis} consisting only of the discourse referent y with its link to the distinguished discourse referent b of this ER, has been plainly ad hoc. However, (23) can be rephrased in ways that are not much less natural and which do not raise the problem of non-assertoric speech acts. One such rephrasing is (31).

(31) That bird on the roof over there is a blackbird.

Interpretation of the demonstrative NP in (31) raises the same issues as the one in (23). In particular, the demonstrative of (31) shares with that of (23) the power to draw attention to its referent in case that referent had thus far gone unnoticed. So the effect of (31) on H is once again the construction or reuse of an ER for the bird together with a DRS for the sentence that is now filling the slot of K_{dis} in his AC.

There is one more move that H could have made (voluntarily or involuntarily) in the course of his interpretation of (23): He could have added his Entity Representation ER'_b to the second component K_{enc} of his AC (assuming that it wasn't there already). Addition of an ER to your K_{enc} amounts to recording your knowledge of the represented entity, as a new item in your entity library. This happens on some occasions when we encounter new things, but mostly it doesn't. Most of us are blessed with a sound capacity for not burdening ourselves with information that doesn't recommend itself as something that might prove to be of later use.

This might well be the outcome of the episode around S's utterance of (23): H notices the bird and makes sense of what S is saying to him. But his interest in the case is no more than fleeting and momentary and shortly after his ER ER'_b will have been wide from his mind, as if it had never been there. (But of course, things could also go otherwise. Perhaps this was a

memorable event. Perhaps blackbirds are rare in the place where the event takes place and H knows that; or there is something remarkable about the blackbird's behavior, as when it drops dead while they are watching, or by suddenly breaking out in the song of a nightingale.)

The next example is one in which the use of a demonstrative is more likely to have a lasting effect upon the interpreter's K_{enc} and it also raises some other issues. Here is the story: A gentleman with money and some education has hired a private guide to give him a tour of the centre of Florence. At one point the guide says (32) to him.

(32) Cellini's Perseus over there has recently been cleaned.

Suppose that the guide's use of the noun phrase *Cellini's Perseus over there*[70], combined with the way she is looking or pointing, directs his attention to the statue she is referring to and causes him to form a perceptually anchored representation of it. What else happens, as part of his interpretation

[70]This example also illustrates a further point. According to the simple-minded classification of definite noun phrases in terms of their morphology, on which we are relying tacitly throughout this paper, the phrase *Cellini's Perseus over there* doesn't classify as a demonstrative NP, since it doesn't begin with one of the demonstrative determiners *this* and *that*. But from a functional perspective it behaves as a deictic demonstrative no less than the canonically demonstrative phrase of the pervious example. A similar consideration applies to the definite description *the roof over there* that is part of the demonstrative NP *that bird on the roof over there*. When talking about (32) above I ignored the compositional structure of the demonstrative NP so as not to overload the discussion. Skipping over the description contained in the demonstrative was part of that strategy. But it is clear that a more careful analysis of that example, which also pays attention to the way in which the descriptive content of the demonstrative can be computed from the semantic representations of its parts, will have to deal with the definite description as yet another definite NP whose identification presupposition must somehow be resolved or accommodated.
It seems intuitively obvious that in the context at hand the interpretation of *the roof over there* proceeds along the same lines as the demonstrative of which it is a constituent: if the interpreter has an ER for the roof in question he can use that; if not, the speaker's words will draw his attention to the roof and he will construct a perceptually anchored ER which he can then use in the interpretation of the NP *the roof over there*. From this interpretation-related point of view, *the roof over there* is thus a 'deictic demonstrative', just like the demonstrative NP *that bird on the roof over there*. In other words, there are 'deictic definite descriptions' just as there are 'deictic demonstrative NPs'. (There is a further question about the constituent *over there*. I take this expression to be a spatial demonstrative, but won't go into details of how exactly it makes its contribution to the semantics of the phrase *that bird on the roof over there*.)

of what his guide is saying, will depend on his antecedent state of information. If he already knew about Cellini's Perseus – say, from his preparation for this trip – and has formed an entity representation for it on the basis of what he has read or heard about it (or perhaps from some pictures he has seen), then he will now link or merge his new perceptually anchored Entity Representation for Cellini's Perseus with the one that he acquired back home. Once more I assume that this takes the form of adding a current perceptual anchor to the existing anchor set of the ER he already had.[71]

A second possibility is that our tourist has never heard of either Cellini or his Perseus. In that case he can be expected to represent the new information – that the referent of his perceptually anchored ER is by Cellini and goes by the name 'Perseus' – as part of the new ER for the statue that he adds to his K_{env}. And presumably he will also form an ER for Cellini himself. Furthermore, it is plausible in this case that both these ERs will become more permanent items in the tourist's entity library K_{enc}. (We are assuming that he is eager to learn.)

Our Perseus scenario allows for more variations. Here is a third variant. Our tourist didn't know about Cellini's Perseus but did know something about Cellini himself – as a flamboyant 16-th century Italian artist, who wrote a famous autobiography that Goethe found interesting enough to translate into German (but which our tourist hasn't read). In that case he may be expected to introduce into his K_{enc} a new Entity Representation for the statue and establish some suitable link between it and the representation of Cellini that was already part of his K_{enc}.[72] The converse case, that where the tourist

[71] I note for the record that this need not be the only possible reaction. Perhaps our gentleman isn't altogether certain that this is the statue that he read about. In that case his mental state might now include a new, perceptually anchored ER coexisting with the old one he already had and with a propositional attitude that queries whether the two represent the same statue

[72] The link that needs to be established between the new Entity Representation for the Perseus and the existing one for Cellini is a relational one: the second entity is the one who made the first one. It seems a natural question to ask where this relational information should be encoded: as part of the new Entity Representation; as part of the old one; as part of both; or in yet some other place? This is one of a number of questions that an account which makes use of Entity Representations forces us to ask. They are reminiscent of similar questions that arose, some thirty five years ago, for Heim's File Change Semantics Heim (1982,1988): Where in a file card system of the sort envisaged in

already knew about the Perseus, but didn't know about its maker, is yet another possibility. But these cases do not seem to bring to light anything of interest about the interpretation of demonstrative NPs that we haven't yet encountered in the cases we considered.Readers may go through a few of these variants for private amusement.

3.3.2 Speaker's Intentions and Common Ground

It is one thing for an interpreter to introduce a new Entity Representation in response to a definite NP; it is another for the speaker to intend him to do so. It needs no argument that this distinction is important. It is connected with two aspects of verbal communication which are prominent in much of the theorizing that can be found in natural language pragmatics: (i) the speaker's communication-related intentions and (ii) the kind of coordination between speaker and hearer that is commonly described in terms of their 'Common Ground'.

So far we have assumed that the Articulated Contexts which figure in our account of meaning and interpretation are representations in the minds of interpreters and that the interpreter can use the information in his AC for resolving the identification presuppositions of definite NPs. A speaker who wants to attune the choice of her words to the interpretational possibilities of her addressee also needs information of this sort. (For instance, if the speaker didn't have ERs for the entities she wants to talk about she couldn't

---

FCS should propositional information be filed that involves more than one card? What seems to me the most plausible answer in the case before us is that the relationship should be encoded as part of the new Entity Representation and that this encoding should make use of the distinguished discourse referent of the old one. (This would make the new ER referentially dependent on the old one. But such dependencies are inevitable one way or another.) We should not forget, however, that, for all we know, mental states may be in perpetual flux: at pretty much any point in time information can be retrieved from the depths of memory and made salient to consciousness; also – and again: for all we know – information present in some form can be reformatted in various ways, with the possibility of shifting bits from one place in the over-all representation to some other place. In particular, information that is part of one Entity Representation may be shifted to, or copied onto, some other Entity Representation. Redrawing the borderline between reference-determining descriptive information of an unanchored Entity Representation and contingent information about the thus determined referent could be seen as a further instance of such flexibility.

have had the intention to talk about those entities.) But in addition, and crucially, she needs information about what relevant information is available to her addressee. More precisely, she needs information about the addressee's Articulated Context; and in particular, she needs information about his K_{enc}.

The speaker's choices are guided by this kind of information in the following manner. If she assumes that the addressee has an Entity Representation for the entity she wants to talk about, then she should use a definite noun phrase, and one that will enable him to recognize that what she is talking about is the referent of this Entity Representation. If she assumes that he doesn't have an Entity Representation for the entity she wants to say something about, then she may have a prima facie preference for using an indefinite, thereby signaling that she is not expecting him to be able to identify what she is talking about on the strength of the information he has.[73][74]
There is a clear asymmetry between the epistemic preconditions for definite noun phrase interpretation and those for definite noun phrase selection. Normally all the interpreter needs is information about the world that the speaker is talking about (more specifically: about what entities inhabit that world). The speaker needs such knowledge too. (Otherwise she couldn't trefer in the way we are assuming she does.) But what she needs in addition is information about the addressee's information about these things. In terms of the hierarchy of mutual knowledge – beliefs or assumptions about the other's knowledge, beliefs or assumptions about your beliefs or assumptions about the other's knowledge, beliefs or assumptions .. and so on – part of the information the speaker needs is, you might say, on the first rung of the ladder,

[73]We already noted that speakers abide by this principle only up to a point. They often use a definite noun phrase (typically a definite description or a name) even though they know that their interlocutors are not familiar with their referents. But they rely, consciously or unconsciously, on the readiness of their interlocutors to accommodate the Entity Representations that they do not have and that they are known not to have. But this doesn't alter the fact that speakers often do choose their noun phrases in the light of what they know or assume about their audience.

[74]In what form speakers represent this sort of information about their interlocutors is a matter that I won't pursue here. But let me note that the MSDRT formalism provides the means for representing such assumptions. It has the power to express iterated attitude attributions – like 'A believes that B believes that p' and so on – of arbitrary nesting depth. One application of this aspect of MSDRT is in describing the attitudes involved in forms of common knowledge and other kinds of attitude sharing, attitudes of the form 'I think that you think that I think that ...' .

whereas the information required of the hearer is at the level of the ground that the ladder stands on.[75] But while the information that speakers need to make suitable choices of noun phrases is 'one step up' from the information needed by hearers, this information falls short, and by a very long stretch, of what belongs to the 'Common Ground' of speaker and audience, as that notion is widely understood. On that understanding an item p of information is Common Ground between two or more agents a_1, .., a_n if (i) p is assumed by each of the a_i; (ii) it is assumed by each of the a_i that p is assumed by each

[75]There are cases where an interpreter will reason, as part of trying to determine what a given NP refers to, about what the speaker may have assumed about his knowledge. And speakers, from their end, will on occasion speculate about such second order speculations on the part of their addressees. Considerable ingenuity has been expended in the literature on the construction of communication situations which prompt such higher order reflections on the epistemic state of one's communication partner. I am not concerned about such cases here but only with the default cases where interpreters rely just on what they know about the world of the topic and speakers rely on their own knowledge of the topic together what they assume is known about the topic by the audience.
Only after completion of the present incarnation of this paper did I become conscious of the important and sophisticated work in probabilistic modeling of speaker-hearer interactions known as the *Rational Speech Act model* (RSA). The models developed within this framework that I am familiar with consider three levels of utterance processing: (i) that of the 'literal listener', who assesses the informational content of possible utterances on the basis of (a) a prior probability distribution of what the world can be like and (b) his knowledge of the (syntax and) semantics of the language to which the possible utterances belong; (ii) that of the speaker who, knowing about the state of the world and reasoning about what the literal listener will infer about the world on the basis of each of the possible utterances she could use, will choose the utterance that is optimal according to this evaluation, in the sense that the litterer will infer that the state of the world is the one she has observed with the highest degree of certainty; (iii) that of the 'pragmatic listener', who infers what the world is or may be like on the basis of the utterance he actually receives and of his model of how the speaker will choose her utterance given her knowledge of the world and her model of him as literal listener. Such models give us insights into the possible effects of speaker-hearer interaction in utterance selection and interpretation that (as far as I can see) could not be gained through purely qualitative considerations. For one thing, it is now possible to form an educated opinion whether further iterations of the back-and-forth between speaker and listener – with a 'second level speaker', who reasons about the pragmatic listener, as next step – could be relevant to a communication in that they would output distributions that would be significantly different from those obtained lower down in the iteration hierarchy. (Whether such further steps would be psychologically realistic is of course another question, which of course also can and also should be raised about the first three levels.) For two representative contributions to the research program of RSA see Frank and Goodman (2016), Frank and Goodman (2014), Stuhlmüller (2013).

of the a_i; (iii) each of the a_i assumes each of the items described under (ii); and so on indefinitely.[76] There are kinds of verbal communication for which this notion of Common Ground seems to offer an attractive model. Striking in this connection is the substantial body of evidence that was collected over many years by Herb Clark and his associates about noun phrase selection in situations where two agents have to carry out a shared task and where this requires them to refer back and forth to certain entities involved in the task that cannot be jointly seen by them, so that there is no way for one speaker to elucidate what she is referring to with an NP by pointing at the referent in a way that the other speaker can track (Clark (1996)). In such situations the participants usually manage to converge with remarkable speed on definite descriptions that they will then both use for the duration of their cooperation to refer successfully to the various entities that their joint task requires them to communicate about. It seems quite natural to model the state of mutual understanding that the subjects of such experiments reach about the use of these descriptions as examples of Common Ground in the strong sense just described.[77]

But even if Common Ground provides a plausible model for some conversational phenomena, that does't mean that it should be the right model for all instances in which discourse participants need to make assumptions about the knowledge, beliefs or assumptions of their conversation partners. As noted, the claim that a piece of information p is Common Ground between a speaker A and an addressee B is a highly complex one, which entails an infinite hierarchy of mutual attributions of ever greater iteration depth. In

[76]Formally an epistemic operator $\mathrm{CG}_{\{a_1,..,a_n\}}$ with this force can be defined as a 'fixed point' operator. When applied to any information item p it yields a formula $\mathrm{CG}_{\{a_1,..,a_n\}}p$ that satisfies the following axioms: (i) $\mathrm{CG}_{\{a_1,..,a_n\}}p \rightarrow p$; and $\mathrm{CG}_{\{a_1,..,a_n\}}p \rightarrow \mathrm{ASS}_{a_j}p$ for $j \leq n$. (Here 'ASS_{a_j}' is short for 'a_j assumes that'.) For the logic of such operators (which is of no direct relevance to the present essay) see Pfau et al. (2015).

[77]That is, both know what each of the descriptions they have adopted refers to, both know that the other knows this and so on. Much of this knowledge will be implicit, and premised on the assumption that the other has pushed Common ground reasoning to the point where they have reached the conclusion that you are attributing to him. But it excludes the possibility that either will disbelieve, let alone be prepared to categorically deny, any of the statements that are part of the infinite hierarchy involved in the definition of Common Ground given above. For instance, if the partners are A and B and it is part of their Common Ground that p, then it can be inferred that A doesn't believe that B B doesn't believe that p.

comparison the attributions that speakers must make to their addressees in order to make effective choices of referring noun phrases are usually very simple – only the first of the infinitely many iterations is involved. On the other hand there is an aspect to the attributions that are at issue here which presents a challenge of a different sort. (This is a challenge that to my knowledge hasn't been previously discussed, if only because it is directly connected with the concept of an Entity Representation, which as far as I know has not previously been put to the purposes to which it is being put in this paper and its companions.) At an intuitive level the problem is easily explained. Suppose that A wants to talk about an entity that is represented in her mind by the Entity Representation ER_A and that she has to choose between a definite and an indefinite NP for the purpose of referring to that entity. Her choice should be guided by what she assumes about the presence or absence of ERs that represent the same entity in the mind of her interlocutor B that is represented in her own mind as ER_A. Let us formally represent these assumptions that A makes about what ERs B has or doesn't have as representations of the referent of ER_A as '$f_b(ER_A)$'.[78] Here f_b is a function that is defined for a certain subset of A's entity library, and ER_A is among the ERs that belong to f_b's Domain. But what should we take the values of f_b to be – such as, for instance, $f_b(ER_A)$? This is my proposal: Either A assumes that B has no ER for the referent of ER_A, in which case $f_b(ER_A)$ is the empty set. Or, alternatively, A assumes that B has one or more ERs for the referent of ER_A, in which case $f_b(ER_A)$ is a set of *ER types* – form-related properties that ERs will have or fail to have in virtue of their structure. If the set $f_b(ER_A)$ is not empty, then this means that for each type $\mathcal{ER}$ in $f_b(ER_A)$ A assumes that B has an ER of this type which is coreferential with ER_A, and furthermore that these instances of the types in $f_b(ER_A)$ are all the ERs in B's entity library that are coreferential with ER_A.[79]

When $f_b(ER_A) = \varnothing$, then the Familiarity Principle would predict that A should select an indefinite NP (presumably with a descriptive content that will give some useful clues about what kind of entity she has in mind). But as we already noted, in practice this constraint isn't binding and speakers

[78] Think of 'b' as the distinguished discourse referent of the Entity Representation that A has for B.

[79] I will not in this paper make a proposal for how an agent like A will *represent* the ER types that make up the values of a function f_b that she may entertain in relation to a discourse partner B.

often ignore it. If $f_b(ER_A)$ is a singleton $\{\mathcal{ER}_b\}$, then A ought to choose a definite NP which provides the information that B will need. (The idea here is that the property $\mathcal{ER}_b$ will give A the information she needs to choose such an NP.) Finally, if $f_b(ER_A)$ has more than one member, then S must choose one of the ER types belonging to it and then choose an NP that B will be able to recognize as referring to the entity represented by his ER of that type that is coreferential with (ER_A). [80]

[80]The cases in which if $f_b(ER_A)$ contains more than one element are somewhat unusual and perhaps they should be ignored in what follows. (Ignoring them probably makes it easier to understand the role that f_b is meant to play in the more central cases, for which $|f_b(ER_A)| \leq 1$.) $f_b(ER_A)$ will contain more than one element only when A has reason to assume that B has two or more ERs for the same referent (the referent of her own ER_A), but without realizing that these ERs are coreferential. Agentd with such multiple representations of the same entity have been a topic of avid discussion in the philosophy of language sand mind. One prominent example is the fictitious character Pierre, introduced into the world of analytic philosophy in Kripke's paper 'A puzzle about Belief' (Kripke (1979)), who has two ERs that are both anchored to London (the capital of the UK) but who thinks they represent different cities. With one of these ERs Pierre associates the name *London* and with the other the name *Londres*. (Pierre was born and raised in France and then at some point in life landed in one of London's slums, where from the on he has carried on a lingering existence. His 'London'-labeled ER represents to him London as he currently experiences it. His 'Londres'-labeled ER dates back form his childhood in France.) The discussions of this and similar cases in the literature have focused on the question how the propositional attitudes of agents suffering form this kind of 'double vision' predicament are best described – for instance, what the right way of identifying the content of the beliefs that Pierre associates with the names 'London' and 'Londres' respectively? – and with the semantics of the attitude attributions that other agents may make to him or her. (For instance, what can we say about the semantics of the belief report 'Pierre believes that London is ugly'?) These are non-trivial issues. But the relevance of such agents for our present discussion is a different one: How is a speaker to choose her NPs when she is speaking *to* an agent with double vision? For instance, if a speaker S wants to make statement about London that is addressed to Pierre, what NP should she choose to refer to London as part of what she is going to say? Asking this question in the way I have just done may suggest there is more of a problem there really is. especially when the situation is one in which ER_A is A's own representation for London and $f_b(ER_A)$ consists of two ER-types that differ from each other in that one is of an ER that includes the information 'is called London' and the other the information 'is called Londres'. What noun phrase is a speaker with this information to choose whern she wants to address Pierre and tell him something about London? In general she should have no problem in guiding Pierre's attention to the ER that she wants him to make use of in his interpretation of her words, for instance by using a complex NP like 'London, the place where you live' when she wants to guide him to his 'London'-labeled ER. But of course she could also just explain to Pierre, using whatever words that would take, that

In order that A's function f_b gives her proper guidance in her choice of NPs, the assumptions it encodes must of course be correct. The assumptions could be off the mark in more than one way, either because $f_b(ER_A)$ is empty while B has an ER for the referent of ER_A, or because $f_b(ER_A)$ is not-empty although B has no such ER, or because the ER B has for the referent does not fit the type that $f_b(ER_A)$ specifies or because the ER that B has and that is of the type specified does not represent the referent of ER_A. In the first case A will assume that the entity she wants to refer to is unfamiliar to B and so she may choose to use an indefinite to talk about this entity. B may then perhaps recognize what she is talking about nevertheless and relate what she has to tell him to the ER for the entity that he has. Alternatively he may treat what A is talking about as something new to him sand perhaps create a new ER for it. This will lead him to have to ERs for the entity A has just talked about. That would not be optimal but as a rule it doesn't need lead to serious problems.

In the second case, in which A wrongly thinks that B has an ER for the entity she wants to talk about, she is likely to choose a definite BP to refer to it. All that B, who doesn't have such an ER, can do in response is to accommodate a new ER for the entity that he takes A to have referred to. What happens in the third case is hard to predict from the little that I have said about it. If A chooses her NP on the basis of her assumption that B's ER for the entity is of the type of the one member of $f_b(ER_A)$, then there is a good chance that B won't be able to recognize that she is talking about the referent of the EDR that he does have for the entity. In that case B will presumable end up with two ERs for this entity, just as that is the likely outcome in case 1. But here too it is just possible that B will recognize that A is talking about the referent of his ER for the entity she is talking about and use that ER to interpret her words. Th fourth case is the one that is most likely to lead to serious troubles. If S chooses an NP geared towards leading B to an ER of the type specified by $f_b(ER_A)$ there is a good chance he will be induced to use his ER of that type; but of course in that case he is bound to misinterpret what A is saying (viz. as about another entity than

---

his 'London'-labeled ER and his 'London'-labeled ER represent the same place.
The moral: Cases where $f_b(ER_A)$ consists of ,more than one elements are of marginal interest for the p;resent discussion and are set aside. In other words we make the assumption that $|f_b(ER_A)| \leq 1$.

the one she is talking about).

To summarize: The different ways in which $f_b(ER_A)$ can misrepresent what information B has about the referent of ER_A can lead to different kinds of problems for B's interpretation. Some of the mishaps that can occur will be fairly harmless, but some can be serious.

3.3.3 Common Ground and Discourse Context

Everyone has his own entity library. But as a rule the entity libraries of different agents show significant overlap and the better the agents know each other, or the closer the social circle or circles to which they both belong, the larger the overlap will be. Yet the overlap is rarely if ever total and even if it were, neither agent would normally be in a position to know this for sure. Since we have identified entity libraries with the K_{enc} components of ACs, it ought to be possible to state this observation also in terms of K_{enc}.[81]

[81] What is overlap? That is a non-trivial question. A simple characterization of the term is the one according to which the overlap between the ER sets of A and B is determined by the set of entities **d** such that both A and B have at least one ER that represents **d**. We can then say that an ER of A (or of B) belongs to the overlap if it represents one of the entities in this set. This characterization, however, is unsuitable for an account of how definite noun phrases function in verbal communication. A more appropriate definition should pay attention to what A and B each assume about the entity library of the other. There are several ways in which such a definition can be formulated so long as we assume, as we did in Section 3.3.2, that communication partners have representations of what kinds of ERs are available to their interlocutors. Let us assume once more that this information is given in the form of a function f_B that maps ERs of the agent to sets of ER types (where for a given ER (ER_A of the agent $f_b(ER_A)$ is either empty or a singleton, and when a singleton, then it is the type of the ER she assume he has for the entity represent by her (ER_A)). Then this ids a way of defining overlap: **d** belongs to the overlap between the sets of entities for which A and B have Entity Representations if (i) both A and B have ERs that represent **d**; (ii) for each Entity Representation ER_A that A has for **d**, $f_b(ER_A)$ is *completely correct*: and (iii) likewise, $f_a(ER_B)$ is completely correct for each Entity Representation ER_B that B has for **d**. Here 'completely correct' is defined (for the case of $f_b(ER_A)$) as: $f_b(ER_A)$ is *completely correct* iff (a) for each type $\mathcal{ER}_b$ in $f_b(ER_A)$ B has an ER that represents **d** and is of that type, and (b) every ER that B has for **d** is of one of the types in $f_b(ER_A)$. The overlapping ERs in the entity libraries of A and B can then be defined once more as all those that represent one of the entities **d** in the just defined set. Other definitions could be argued for as well, but this much must suffice for here and now.

To state the observation in terms of K_{enc} we need to make a further assumption. But I think it is an assumption that is quite plausible. Let us once more focus on a situation of face-to-face communication in which A does the talking and B the listening. Thus far we have spoken of K_{dis} construction, and of Articulated Contexts that contain the constructed K_{dis} as one of their components, only in relation to the interpreter B. But there are good reasons for taking it that such mental state constituents can also be ascribed to the producer A: as A is putting her utterance together, she herself can apply the interpretation rules of the language to the words she is producing, and in much the way that her addressee can. Doing that won't tell her anything new, of course, for it is after all her own thoughts that she is in the process of putting into words; but it is a way of keeping, from the 'listener's perspective' she is adopting, tabs on whether her words are readily understandable for her interlocutor, and also as a way of keeping track of the parts of what she wants to say, that she has already expressed and the parts that still remain to be put into words; moreover, this will also give her a better grip on what words should be found for the remaining parts, given the semantic representation of what she has said or written so far.[82]

Let us suppose, then, that as A is speaking to B they each develop their own ACs and build their own K_{dis}s as part of them. In general, the discourse contexts they construct, $K_{dis,A}$ and $K_{dis,B}$, will not be identical in structure. In that respect they are like the other components of the ACs AC_A and AC_B that A and B entertain at any point of their conversation. For the reasons pointed out above, $K_{enc,A}$ will as a rule differ from $K_{enc,B}$, and for similar reasons $K_{gen,A}$ can also be expected to differ from $K_{gen,B}$; and in situations in which there is an environmental component to A and B's ACs, $K_{env,A}$ can be expected to often differ from $K_{env,B}$.

In this regard, then, K_{dis} is like the other components of ACs. But whereas the differences between the other components are as a rule substantial, in that they represent significantly different information, in the case of K_{dis}this is typically not so. When both A and B construct their K_{dis}-components by following the rules of the language, then their respective K_{dis} representa-

[82]Whether this is exactly what people actually do when they speak (or write) doesn't matter all that much. It is enough that the speaker or writer is in a position to slip into the role of interpreter of her own words, and in this way to keep a record of what has so far been said or written.

tions will be very similar in form (if not in general fully isomorphic; recall the discussion in Section 2.3) and they will determine the same propositional content. Moreover, A and B can both feel certain that their respective K_{dis}s will be semantically equivalent – since that will be the case so long as their conversation partners are competent speakers and are making the correct use of their linguistic competence. And when the discourse participants do feel confident about all this, then they can also be certain not only that their respective K_{dis}s are equivalent, but also that their conversation partner is similarly confident of this; and confident not only of that, but also of their own confidence that the two K_{dis}s are equivalent. And so on. In other words, it seems reasonable to consider it part of the Common Ground between A and B that $K_{dis,A}$ and $K_{dis,B}$ are content-equivalent. Or, to put it in yet other terms, A and B can both be expected to consider the intended interpretation of A's words to be *public* in roughly the sense of Lewis' scoreboard (Lewis (1979b)): it is as if A, by saying what she says, puts up a semantic representation of her words in a place where both she and B can see it, so that they both know that the other can see it too and rely on that knowledge when they want to refer to some of the things that are posted there; and what is on the scoreboard is not open to misinterpretation.

With the assumption that K_{dis}-components are Common Ground it may seem that we have come full circle. Aren't we saying in essence that these K_{dis}-components are simply personal copies of the user-neutral DRSs that are the results of discourse representation according to DRT in its original form? Almost, but not quite. Often the K_{dis}-components from the Articulated Contexts of the different discourse participants contain distinguished discourse referents from ERs belonging to their respective K_{enc} components. That renders them essentially private, mental structures, with links to other parts of the mental state they are part of. But of course that doesn't prevent these representations from being content equivalent: As long as the ERs from other parts of A's AC that have lent their distinguished discourse referents to A's K_{dis} are coreferential with the ERs from B's AC that have lent their distinguished discourse referents to B's K_{dis} in the same places and so long as both K_{dis}s have been constructed from the shared linguistic input according to the rules of the language, they ought to determine the same proposition.

The picture we have arrived at is this: of the different components of the Articulated Contexts of two agents involved in face-to-face communication it is

the K_{dis} components that represent information with the status of Common Ground. That is in general not true for the other AC components; normally those show some overlap, but as a rule the overlap will only be partial; and it is usually difficult for the agents to determine how much overlap there is and what belongs to it.

By and large this picture is correct, but there is one part of it that needs retouching. In face-to-face communication the environment has a mutual transparency which approaches that of the current discourse context. The interpretation of the deictic use of a demonstrative NP requires, we argued in Section 3.3.1, detecting the referent in the environment, just as the interpretation of an anaphoric NP requires the detection of the relevant *referent representation* as a salient constituent of the discourse context.[83] In either use, deictic or anaphoric, the speaker can make interpretation easier for her audience by supplying collateral information that assists the interpreter in targeting the intended deictic referent or anaphoric antecedent. For deictic demonstratives this can take the form of a pointing (or some other non-linguistic deictic act) or of descriptive information built into the NP itself; and often it will be a combination of the two. For anaphoric NPs guidance is provided by the descriptive information that is part of the NP (though the choice between the determiners *the*, *this* and *that* is important too, as is the choice between pronouns and non-pronominal anaphoric NPs).[84] To a significant extent the role played by the descriptive content of anaphorically interpreted NPs can be seen as analogous to the one that demonstrations play in the interpretation of deictic NP uses.

Looked at in this way deictic and anaphoric interpretations have much in common: in both cases the expression is taken to pick out an element from a display that is accessible to both speaker and hearer, in a sense that Lewis' scoreboard metaphor so aptly captures. This commonality may be at the root of a remarkable cross-linguistic fact: in language after language the

[83]In the terms of our DRT-based framework this referent representation will be a discourse referent that is accessible in the technical sense of DRT, and also 'salient' within the given discourse context DRS. For some work on salience within formalisms akin to DRT see for instance Dekker (2001).

[84]For the different constraints on the choice of anaphoric antecedents for pronouns and non-pronominal anaphoric NPs see in particular Centering Theory (Grosz et al. (1983), Grosz et al. (1995)).

types of definite noun phrases that can be used anaphorically are the same as those that can be used deictically.

3.3.4 Deictic and anaphoric uses of demonstratives and pronouns

So far I have been using the terms 'deictic' and 'anaphoric' in an intuitive sense. In this I have been following what appears to be a wide-spread practice within linguistics, and I think that the use I have been making of these terms should have been reasonably clear. However, the concept of an Articulated Context we have been developing enables us to give more precise characterizations of these notions:

(i) An NP is *interpreted anaphorically* if its identification presupposition is resolved via the discourse context K_{dis}.

(ii) An NP is *interpreted deictically* if its identification presupposition is resolved via the environment context K_{env}.[85]

What it means to resolve an identification presupposition using K_{env} has already been stated explicitly: An ER is chosen from K_{env} and the referential argument of the NP is set equal to the distinguished discourse referent of this ER. The resolution of the identification presuppositions of anaphorically interpreted NPs – in one sense of 'anaphoric', see below – is a closely similar story, which differs on only one point: the choice of an anaphoric antecedent now takes the form of selecting a discourse referent from the main universe of K_{dis} (rather than choosing an ER from K_{env}). But in both cases the selection is from a readily accessible scoreboard-like display.

Note well, however: the use of K_{dis} in the resolution of presuppositions does not have to take this particular form. Many presuppositions are *propositional*; they require the truth of some particular propositional content. The

[85] A speaker or author can use a definite NP *intending* it to be resolved anaphorically or deictically. The default situation is that in which the interpretation by her audience is in agreement with her own 'speaker's intention': either she intends the NP to be interpreted anaphorically and the audience interprets it that way or she intends the NP to be interpreted deictically and that is also the way the audience interprets it. In such cases the NP occurrence can be called 'anaphoric' or 'deictic' without qualification.

way to resolve a propositional presupposition is to show that it is entailed by the available contextual information, with 'context' understood in a broad sense which includes all information provided by Articulated Contexts as we have defined them here. And among the ways in which contextual information can be used for the resolution of such presuppositions there are those in which the discourse context provides all the information that is needed, or at least the most prominent part of it. Presupposition resolutions that rely on the discourse context can be regarded as 'anaphoric' for that reason, irrespective of whether or not the presuppositions are propositional.[86]

This notion of anaphora should be distinguished from the one assumed two paragraphs back. The notion I implicitly referred to there is the one that seems to correspond more closely to the intuitions that have traditionally been associated with the term 'anaphora': that notion according to which an expression is anaphoric if it refers by virtue of being interpreted as 'anaphorically linked' to another element from the discourse – a link which renders it 'coreferential' with that other element. These two notions – of anaphora in the sense of being interpreted as coreferential with some element that can be recovered from the context and of anaphora in the sense of being resolved via the *discourse* context – are not always clearly distinguished in the literature. But it is important to see that and how they are different. In particular, expressions that trigger propositional presuppositions can be anaphoric only in one of these senses, whereas so-called 'anaphoric pronouns' qualify as 'anaphoric' in both senses:

(i) they are anaphoric in that their presuppositions must be resolved via the discourse context; and

(ii) the resolution of the presupposition of an anaphoric pronoun takes the form of finding an anaphoric antecedent for the pronoun from among the elements that the discourse context makes available.

One reason, I suspect, why the two notions have often been conflated is precisely that with third person pronouns – arguably the expression type that has been the most prominent in discussions of the relations between

[86]This conception of anaphoricity is in essence the view of 'presupposition as anaphora' put forward in Van Der Sandt (1992).

presupposition and anaphora – the two notions coincide. For instance, in the framework we are using in this paper the conditions (i) and (ii) can be naturally combined into a single one:

(iii) the identification presupposition of an anaphoric pronoun must be resolved by identifying the referential argument of the pronoun (i.e. the discourse referent introduced to represent the pronoun at the level of logical form) with a discourse referent from the Universe of the DRS K_{dis}.

I believe this is an intuitively plausible explication of what has commonly been understood by the term 'anaphoric pronoun'. But the explication exposes the term 'anaphoric pronoun' as a kind of misnomer. As a syntactic category, third person singular pronouns are open to two kinds of interpretations, anaphoric and deictic. An 'anaphoric pronoun' is a pronoun – more correctly, a pronoun occurrence – for which it has already been decided that its interpretation should be anaphoric (as opposed to deictic). What makes the notion 'anaphoric pronoun' seem natural is that for a great many pronoun occurrences, among them those that have been prominent in discussions of anaphora, anaphoric interpretation is the only option. This is true in particular for pronouns occurring in texts, whose interpretation context is (with marginal exceptions) lacking a K_{env} component. For such occurrences the interpreter knows from the start that a deictic interpretation is out of the question. Sometimes this is obvious also for pronouns in spoken utterances.[87] But even though many pronoun occurrences can be recognized as 'anaphoric' off the bat, because a deictic interpretation is obviously ruled out, this justifies the use of the term 'anaphoric pronoun' only up to a point. (To repeat, as members of a grammatical category pronouns are strictly speaking never 'anaphoric' – qua types they can get both anaphoric interpretations and deictic interpretations, and so are neither 'anaphoric' nor 'deictic'.)

[87] As far as I can tell, it is still a matter of the dispute whether definite descriptions can be anaphoric in both of the two senses distinguished here; that is, not only in the sense that the presupposition triggered by the definite description is justified on the basis of the discourse context, but also in the second sense, that the description gets its reference via linking with an element from the discourse context. This question is only one of many difficult issues that have to be settled as part of a comprehensive account of the range of possible interpretation mechanisms for definite descriptions in English (and corresponding expressions in other languages). A comprehensive description of the different uses of the different English definite NP types and of the constraints that govern those uses is a topic for another research project, and I will no further elaborate on it here.

Towards the end of Section 3.3.3 it was observed that by and large the same NP types that can be used and interpreted deictically can also be used and interpreted anaphorically and conversely; and I ventured the hypothesis that this is because the interpretational resources involved, the discourse context and the perceptually accessible environment, resemble each other in that both are 'transparent' to the discourse participants in roughly the sense of Lewis' scoreboard. But there nevertheless appears to be some difference in the accessibility (or 'transparency', or salience') of these two resources This shows up in particular when an intended referent is accessible both via K_{dis} and via K_{env}. In such cases it seems that the use of K_{dis} tends to be preferred.

One such situation is that where a first, deictic use of an NP has rendered its referent available at the level of K_{dis} and the need then arises to refer to this referent again. We encountered one example of this in (23) in Section 3.3.1, repeated here.

(23) Look at that bird on the roof over there! It is a blackbird.

For all we have said about this example, the pronoun *it* of the second sentence could either be interpreted anaphorically, by resolving it via the discourse referent for the bird that belongs to the discourse context established by the first sentence, or else interpreting it deictically by making use of the ER for the bird that was involved in the interpretation of the demonstrative *that bird on the roof over there* in the first sentence. But I have a strong hunch that it is the anaphoric option that an interpreter of (23) will actually use. For one thing, *it* in the second sentence would be naturally pronounced without the accompaniment of any demonstrative gesture, and also without a hint of the emphatic stress that is typical of deictic uses of pronouns. This suggests an anaphoric and not a deictic interpretation, as if the special effort involved in deictic reference – that of activating or creating an ER – need not and should not be made again once it has been made already: now that the referent is represented in the discourse context, further reference to it can exploit this new and easier route.

Some further evidence for the preference of anaphora over deixis in situations where both are available in principle is provided by the possible reactions to an utterance of (23) or, likewise, to an utterance of (31). Some possible ways

of reacting to (31) (repeated here as (33a)) are shown in (33b) and (33c).

(33) a. A: That bird on the roof over there is a blackbird.
b. B: Yes, and it cannot be more than a few months old.
c. B: Yes, and that bird on the roof over there/that bird on the roof/that bird over there/that bird/that there/that cannot be more than a few months old.

If the contrast between (33b) and (33c) is not conclusive outright, it certainly highly suggestive. What ti suggests is that after the bird has been introduced into the discourse context (by A's utterance of (33a),) then when you want to refer to it again in a follow-up utterance anaphoric reference is preferred over another deictic reference. Given what has been said about pronouns and demonstratives so far – that both can be used deictically as well as anaphorically – this assessment needs support from additional considerations. First, pronouns. When a pronoun is deictically it is usually stressed and there it is possible and often desirable to aid the reference one wants to make with it with a demonstrative gesture. Neither of these conditions are fulfilled by the *it* of (33b). In a natural pronunciation of (33b) *it* is unstressed, and pointing once more at the bird while uttering the pronoun seems unnecessary and might even be felt like a distraction. By these criteria the use of *it* in (33b) is anaphoric. In our terms, *it* establishes reference via an anaphoric link with the discourse referent that represents the bird in K_{dis}, as it exists after the utterance and interpretation of (33a).

Second, demonstratives. Although demonstrative phrases can sometimes be used anaphorically, their anaphoric use appears to be restricted to situations in which some kind of contrast is involved: Several entities of a certain kind are salient in the discourse context an one uses a demonstrative phrase to establish an anaphoric link with one of those. An example was given early on (1) in Section 2.1. It is repeated here.

(1) If a Texan steals the cattle of another Texan, then that Texan will get very cross.

Selection from a multiplicity of representations of entities of the same kind in the discourse context is not at issue in (33). (In the scenario we are considering there is just one bird on a roof in the visible surroundings.) So it

seems that in (33c) only deictic interpretations of the listed demonstrative phrases are a possibility.

The conclusion we may draw from (33), then, is that when an entity has become represented in the discourse context (and this representation is sufficiently salient because of its recency), then anaphoric reference is preferred over deictic reference. This finding is not an isolated case; there are many other examples to show that when anaphoric reference is a salient option, it is preferred over other options. The question is¿ why should that be?

3.3.5 Reference and Anaphora

The definition we have given of anaphoric interpretations of definite NPs – as interpretations that are based on the discourse context – sets these apart not only from deictic interpretations but also from any other definite NP interpretations that involve AC components distinct from K_{dis}. We have already drawn attention to certain differences between K_{dis} and the other AC components, viz. to the special 'scoreboard status' of K_{dis} and to its dynamics. But in addition there is a further important difference, between the uses that NP interpretations can make of K_{dis} and the uses they can make of the other AC components. This difference has to do with the logically complex structure of many K_{dis}s: A K_{dis} may contain sub-DRSs, which come with their own DRS Universes. These sub-universes also contain discourse referents that can be used as antecedents for anaphoric NPs – so long as they are accessible from the positions of those NPs in the technical sense of DRT (Kamp (1981b), Kamp and Reyle (1993)). When a discourse referent from such a subordinate DRS-universe is used as antecedent for an anaphoric NP, this as a rule imposes upon the NP a 'non-referential' interpretation.

But what is 'referential'? When we want to make this notion more precise it is useful to distinguish between two notions of referentiality, *direct referentiality* and what I will call *referentiality simpliciter*, or, more concisely, just *referentiality*. Within the formal setting of this paper the first of these notions is the one that is more easily defined. When an NP is interpreted via K_{enc} or K_{env} it makes its contribution to the proposition determined by the representation that the interpreter constructs for the sentence containing it via the ER he uses to resolve its identification presupposition; and in our framework as it has been set up this always has the effect that the proposition thus determined is singular with respect to the entity represented by the

ER. The same is true also for NPs that are interpreted anaphorically, but where the discourse referent used as anaphoric antecedent is from the main Universe of K_{dis} and is linked in K_{dis} to the distinguished discourse referent of some ER. In all these cases the NP gets a directly referential interpretation, both in an intuitive sense and in the sense of the definition we will give below.

What makes these NP interpretations directly referential in the sense in which that term is understood in the literature on direct reference is that the referential argument α of the NP – i.e. the discourse referent that is introduced to represent the referent or semantic value of the NP in the semantic representation for the sentence or clause to which the NP belongs – is *anchored to* some particular entity **d** in the technical sense that only embedding functions are admitted in the semantic evaluation of the sentence or discourse representation that map α to **d**. We have seen that this is always the case when α is linked to the distinguished discourse referent of some ER that has **d** as its referent. But as we saw in Section 2, NPs that are interpreted via ERs are not the only cases where the possible embedding functions are under such a constraint. The others are the indexical NPs, whose referential arguments are anchored to components of the utterance context.

In our set-up these are two ways in which an NP can get a directly referential interpretation, and they are also the only two ways. The following 'definition', which does little more than repeat what has already been said, makes this explicit:[88]

[88]From a conceptual perspective definition (34) is unattractive. Intuitively speaking, direct referentiality is a unified concept – that of some kind of direct link between the referring expression and its referent, which injects that referent as a constituent into the propositional content of the sentence or discourse in which the expression occurs. Against the background of this unified conception the disjunction in the definiens of (34a) seems inappropriate and one might hope that a deeper analysis would make it go away. But I have not been able to find a way to make this happen. All attempts to find a uniform definition that does justice to the facts of direct reference in English and other languages seem to run into the problem that the mechanism underlying the interpretation of indexicals is fundamentally different from the mechanism or mechanisms involved in the interpretations of demonstrative NPs even if we limit attention to their deictic uses. Traces of how difficult it is to bridge the gap between these mechanisms can be detected in Kaplan's tentative discussions of the role of demonstrations, as playing a part in the deictic interpretation of demonstrative NPs that is presumed to be somehow analogous the the role played by the utterance context (in our narrow sense of the term) in the interpretation of the indexicals *I*, *you*, *now*, ... Eliminating the disjunction from (34) couldn't be more than papering over

(34) a. An NP interpretation is *directly referential* iff

the referential argument of the NP (the discourse referent introduced to represent its referent or semantic value) is either

(i) identified with a component of the utterance context or

(ii) directly or indirectly linked to the distinguished discourse referent of some ER.

b. An NP *type* can be said to be *directly referential* if its interpretations are always, and of necessity, directly referential.[89]

According to what has been said so far, the only NPs that are directly referential as types are the indexicals. The other types of NPs we have discussed – demonstrative NPs, definite descriptions and pronouns – can get directly referential interpretations but they can also get interpretations that do not produce the effect of direct reference. As noted, these NPs also allow for anaphoric interpretations, in the sense that their identification presuppositions are resolved by the discourse context; and as will be made explicit below, among the anaphoric interpretations of these NPs some are not even referential, let alone directly referential. Nevertheless, for demonstrative NPs directly referential interpretations are extremely common (all their deictic

a difference that is better left exposed.

[89]There is a terminological tension in this definition. As stated, clause (ii) of (34a) is compatible with the possibility that the ER to whose distinguished discourse referent the referential argument of the interpreted NP is linked may be improper. When that is the case, the referential argument of the NP is of course *not* linked to an actual referent. One cannot help feeling that there is an intuitive implausibility in including such cases among the 'directly referential' interpretations of NPs; for after all these are cases where the NP doesn't refer, let alone 'directly'. We could correct for this difficulty by adding to clause (ii) of (34a) the requirement that the ER in question be proper. But when we make this addition, we get into trouble with (34b). The only NP types that will survive the addition as directly referential are the indexicals; those NPs that are meant to qualify as directly referential via clause (ii) won't do so any longer, since it is always possible that their referential arguments get linked with improper ERs.

Intuitively it is clear that the only way out of this predicament is to distinguish between two notions of direct referentiality for NP types: (i) direct referentiality *in spirit*, which is the notion defined in (34) as it stands, and (ii) direct referentiality *de facto*, which results when clause (ii) of (34a) is amplified in the way described in this note.

interpretations are directly referential and so are some of their anaphoric interpretations). To a somewhat lesser extent this is also true of definite descriptions. For pronouns non-referential interpretations are more frequent, but pronouns too get referential interpretations often enough.

The only way for a definite NP to get a non-referential interpretation is when its referential argument is linked to an unanchored discourse referent from K_{dis} (where by 'unanchored' I mean here that the discourse referent is neither linked to a component of the utterance context nor directly or indirectly to an ER). Unequivocal examples are provided by NPs that get resolved through identification with a discourse referent occurring in the universe of a subordinate DRS in a logically complex K_{dis}. As a rule such discourse referents are quantificationally bound in the manner made explicit by the model theory of DRT and their status as quantificationally bound variables will then extend automatically to the referential arguments of anaphoric NPs to which they serve as antecedents.

The classical examples of non-referential definites are pronouns. Typical examples are quantificationally bound pronouns like the *him* of (35a), but also donkey pronouns like the *he* of (35b).

(35) a. Every man admires some woman who has never heard of him.
b. If a man lives in Athens he doesn't live in Sparta.

But as we have seen, pronouns are not the only definite DPs that can get non-referential interpretations. Demonstrative NPs also have non-referential uses. That demonstratives can be used in this way is illustrated by (1) which was repeated in Section 3.3.4 and which I won't repeat yet again. (36) shows that definite descriptions can be used in this way too.

(36) If Fred has a hamster and a guinea pig, he will like the guinea pig better than the hamster.

Since pronouns, demonstratives and definite descriptions must be interpreted either deictically or anaphorically and deictic interpretation entails direct reference (in the sense of definition (34)), interpretations of these expressions that are non-referential must be anaphoric. But on the other hand anaphoricity is no guarantee for non-referentiality. Consider the example in (37):

(37) So it was decided that Bobby Kennedy should run for President. But he/that Kennedy got assassinated too.

If we assume (i) that the occurrence of *Bobby Kennedy* in the first sentence is interpreted by linking its referential argument to the distinguished discourse referent of an ER that is externally anchored to the man Robert Kennedy, and (ii) that the pronoun *he* or the demonstrative *that Kennedy* are interpreted anaphorically through use of the discourse referent introduced into K_{dis} by the interpretation of the name *Bobby Kennedy*, then the anaphoric connection between the referential argument of the pronoun or demonstrative in the second sentence and the referential argument of the name in the first sentence will ensure that the pronoun or demonstrative is also linked to the man. Thus, an interpretation of (37) that satisfies this assumption is an example of a pronoun or demonstrative that is interpreted anaphorically, but nevertheless directly referential.

It ought to be plain by now, but is worth emphasizing nonetheless, that in the framework we have developed, and given the terminology we have adopted, anaphora and direct reference are notions of very different sorts. Anaphora pertains to the way in which an occurrence of a definite noun phrase is interpreted: the interpretation is anaphoric iff the referential argument of the NP is equated with an accessible discourse referent from the discourse context. Direct reference, on the other hand, is a property of the interpretations of definite NPs that has to do not just with how the NP is interpreted but with the consequences that the interpretation has for the contribution that the NP is making to the content of its sentential context. To repeat once more: the NP makes a directly referential contribution according to the interpretation it has been given iff its referential argument is linked (in one or more steps) to a component of the utterance context or to an ER. As we have found things to be, anaphoric interpretations can lead to directly referential contributions, but they do not need to. Likewise, directly referential contributions are sometimes the result of anaphoric interpretations, but not always. This is shown by indexical NPs like *I* and *you*, which never involve anaphora – both on the account of these indexicals mapped out in Section 2 and on any other that I am aware of – and also by deictic interpretations of demonstrative phrases of the kind discussed in Section 3.1.[90]

[90]Proper names are another example of NPs that make directly referential contributions,

3.3.5 Non-direct Reference

In a couple of places I have used the word 'referential' without modification by 'directly'. Since I never defined or explained referentiality simpliciter, as distinct from direct referentiality, this could have passed for a kind of shorthand. But it wasn't. A difference between direct reference and 'reference simpliciter' is implicit in much of the recent philosophical literature on reference. And it was the second of these two notions that I had in mind when talking about reference without using the word *direct(ly)*.

Since 'reference simpliciter' means the same as 'reference that is either direct or non-direct' and a definition of direct reference is already in place, the problem of defining reference simpliciter and that of defining non-direct referent come to much the same thing. In the brief discussion that follows we will focus on the definition of non-direct reference. Furthermore, we will consider just one form of non-direct reference, that of reference by virtue of unique satisfaction. (If there are other forms, I do not know what they might be.)

Since the time when Donnellan, Kaplan and Kripke revolutionized the theory of reference by demonstrating the legitimacy, importance and ubiquitousness of direct reference, the distinction between direct reference and 'non-direct reference' (for lack of a better term) has been a central issue in the philosophy of language. But curiously, satisfactory definitions of 'non-direct reference' are not easy to come by. This is true in particular of the notion of reference that was all-dominant before the mentioned revolution, that according to which NPs refer by virtue of the unique satisfaction of predicates associated

but the interpretations of proper names are rarely anaphoric if ever. That proper names are capable of anaphoric interpretations has been argued in particular by Geurts (see Geurts (1997)). Geurts' examples indicate that anaphoric interpretations of proper names need not even be referential. More directly relevant for our considerations here is that proper names can make directly referential contributions and that result always or nearly always from non-anaphoric interpretations. (Directly referential interpretations of proper names are the focus of Kamp (2015). For those who accept Geurts' examples and his analysis of them proper names not be directly referential as types in the sense of (34b). For those who reject his examples it is possible to hold the position that names are directly referential as types. As far as I can tell this is a position that many philosophers and semanticists subscribe to, even while some of them would object to the terms in which that position is described in the terms I am using.

with them. The suggestion that definitions of satisfaction-based reference[91] are hard to come by may sound surprising; for wasn't that the notion that the reference literature before Donnellan et al was all about? And even after the revolution got under way, the advocates of direct reference themselves never expressed any doubt that non-direct reference does occur – witness Kaplan's discussions of definite descriptions like *the shortest spy* or *the first man to be born in the 22nd century*. Surely they knew what they were talking about.

In fact, the reason why the definition of satisfaction-based reference might be thought easy (and thoroughly old hat at that) is that the examples of non-directly referential expressions that most readily come to mind are precisely definite descriptions like these; and, indeed, for these – definite descriptions with explicit descriptive contents that play the part of the uniquely satisfiable predicates associated with them – a definition is obviously easy to give.

In particular, there would be no problem with giving a satisfactory definition applying to definite descriptions of this sort in the framework of this paper. Basically all we would need to do is to appeal to the traditional identification presuppositions for definite descriptions – the propositions which say that the description's descriptive content has a unique satisfier. The contribution of a definite description δ to the interpretation of the sentence or discourse σ containing δ would then be referential iff the DRS K_{id} which represents the identification presupposition for the description (which is part of the preliminary representation of σ that is constructed along the way to its final representation) has a unique satisfier. More precisely, let $\mathcal{M}$ be a model for the DRS language in question, w a world from $W_{\mathcal{M}}$ and t a time from the time structure of $\mathcal{M}$. Then the contribution of the NP to the given interpretation of σ would be *referential in* $\mathcal{M}$ *in* w *at* t iff there is a unique entity $\mathbf{d}$ in the universe of $\mathcal{M}$ such that K_{id} can be verified in $\mathcal{M}$ at w by an embedding function f such that $f(x) = \mathbf{d}$ and $f(n) = t$, where x is the referential argument of the definite description and by no embedding function f such that $f(n) = t$ and $f(x) \neq \mathbf{d}$.

But there is a well-known problem with this definition. Very often when it seems intuitively right to say of a definite description that it refers because of

[91] I use *satisfaction-based reference* of an expression δ as a shorthand for 'the description associated with δ has a unique satisfier'.

unique satisfaction the descriptive content of the definite description itself is not uniquely satisfied – in the world at large it is multiply satisfied, often by very large and open-ended sets of satisfiers. If unique satisfaction is essential to the way such a definite description refers then it is because the context in which it is used limits the 'search space' within which its satisfiers must be found and does so to such an extent that the space contains just one of the many satisfiers at large. This observation is so obvious that it must go back to when the unique satisfaction analysis of definite descriptions was first contemplated. For long it remained a truth that must have been evident to all, but without making much of a dent in the reference theories that were being discussed. In more recent years, however, considerable work has gone into spelling out how context-driven 'Domain Restriction' works. How successful these efforts have this far been is a matter for dispute. (In my own view a comprehensive account of Domain Restriction is still outstanding.) But here all that I want to note is that a good definition of satisfaction-based reference presupposes a satisfactory account of Domain Restriction.[92]

Nothing that has been said so far addresses the question what contributions non-directly referring NPs make to the content of the sentences in which they occur. In the literature on satisfaction-based reference this question has been of central importance. It was this question that over many decades has fueled the debate. For long the main opposition in this debate has been between the view that definite descriptions come with unique satisfaction presuppositions and the Russellian view that they should be treated as quantifiers which incorporate unique satisfaction into the truth-conditional content of the sentences to which the given definite descriptions belong. One of the commitments made in this paper is that all definite noun phrases come with

[92]There is a further issue that arises both for the characterization of non-direct and of direct reference. A pronoun (or other type of definite noun phrase) that is anaphoric to a referring expression deserves to be regarded as likewise referential. It inherits its referentiality so-to-speak from its anaphoric antecedent. In principle extending a definition of referentiality from a noun phrase to another that is anaphorically linked to it is problematic. It isn't that even though anaphora need not be just a one step affair. The interpretation of a text or longer monologue may involve the formation of anaphoric chains, which can in principle be of arbitrary length. But formulated in the right way the definition will include among the referential NPs all those that are linked to a referential antecedent through being anaphoric all connected to it by some such chain. (Within the framework assumed in this paper, chains can be immediately read off the semantic representations that interpretations build for sentences, discourses and texts.

'identification presuppositions', including definite descriptions that are interpreted as referring on the strength of unique satisfaction. This puts the present account squarely in the first of these two camps. But the treatment of the identification presuppositions of definite descriptions that is being assumed here makes the difference between a presuppositional and a Russellian treatment fairly small. The reason is that satisfaction of the presupposition of a definite description, either through verification by the context or through accommodation, has the effect that the unique satisfaction condition is verified by the discourse representation of which the interpretation of the sentence that contains the description is a part. Whether unique satisfaction is imported into the discourse representation in Russellian fashion by the sentence itself or enters into of the discourse representation through verification or accommodation of the description's presupposition makes no difference to the truth conditions of the discourse representation as a whole.

A difference between such a treatment of the presuppositions associated with definite descriptions and a Russellian account remains in those cases where presupposition verification is impossible and accommodation is blocked because the unique satisfaction presupposition is known to be false. These are the classical cases of non-satisfaction of the presupposition, as in a sentence like 'The largest prime is odd'. According to the presuppositional treated alluded to here the interpretation of this sentence aborts because the presupposition can't be verified and resists accommodation. So we do not end up with a viable interpretation and no truth conditional content is determined. A Russellian would conclude that the sentence, or the discourse of which it is part, is false.[93]

3.3.6 Underspecified Speaker's Intentions and Audience Diversity

In the final three sections of the paper we return to more specifically communication-theoretic aspects of definite NP use. In Section 3.3.1 we noted that

[93] A Russellian twist can be given to the presuppositional treatment I am assuming by modifying it in the following way: Every presupposition gets accommodated if it cannot be verified in the context – even those that are known to be false. In that case the unique satisfaction always gets incorporated into the discourse context and thus causes it to be false rather than causing its abortion. I do not know on what basis a choice can be justified between this and the unmodified version

deictic uses of demonstrative NPs can be felicitous irrespective of whether the recipient has already noticed the referent or not. The situation on which we focused in our discussion of (23) was one in which the interpreter doesn't yet have a perceptually anchored Entity Representation for the bird the speaker is talking about – a representation which he would have formed on his own if he had noticed the bird before – and where the speaker correctly assumes that he has no such representation. But (23), we noted, would also have worked if the addressee had already noted the bird. This is even clearer in the case of (31), which too is just as felicitous in a situation where the bird had already been noted as one at which it wasn't. The only difference between the two situations would be that in the first the addressee can make use of an existing ER, to which he now only needs to add a vicarious anchor that links it to the reference the speaker has just made, whereas in the second situation he has to come up with a new ER 'ab initio'.

Since it makes no appreciable difference to the final result whether or not the addressee already has an ER for the bird when the utterance is produced, it doesn't really matter either whether or not the speaker assumes the addressee to have such an ER. And since the speaker's assumptions on this point do not really matter to the success of the communication, there is no need for her to make any such assumptions; this is not one of the things she needs to worry about.

Similar observations apply to the second example discussed in Section 3.3.1, in which the guide draws the tourist's attention to Benvenuto Cellini's Perseus (example (32)). She too may have had thoughts about whether her addressee has already seen the statue, or about whether the statue is already familiar to him in some other way. In this case too, the details of the addressee's (i.e. the tourist's) information about Cellini or the Perseus may affect the details of the interpretation process, but not the content of the outcome. So for the appropriateness of (32) it isn't crucial either what the guide expects the tourist to know, or whether she has any expectations on this point at all.

There are, then, speech situations in which it doesn't matter exactly what the speaker thinks about the addressee's knowledge of the entities she is referring to, or whether she has any such thoughts. There are also situations in which the speaker may have such thoughts, but where it is less obvious how her NP choices should be guided by these. These are situations in which the

speaker is addressing a group consisting of several addressees, or a crowd (the audience of a lecture, say, or of a public address), or when an author writes (as one does more often than not) for a larger audience. Often a speaker or author has good reasons to believe that some members of the audience are familiar with a certain entity but that other members are not. With what kind of NP should the speaker/author then refer to this entity? Suppose for instance that the speaker of (23) is part of a larger group of people and that she wants to communicate the content of (23) to all the others. She may suspect that some of them have already noticed the bird on the roof but that there are also some who haven't. (She may have concluded this because she has noticed where everyone is sitting or standing on the balcony where they are all gathered and which directions the different people are facing.) Under these conditions (23) still seems a perfectly good way of conveying what she wants to say, and of conveying it to everyone of the people present. It will get those who hadn't noticed the bird so far to form Entity Representations for it now, and get those who had noticed it before to put the Entity Representations on alert that they formed when they noticed the bird. And then everyone will be able to use their Entity Representations to construct representations of the propositional content of (23) (the proposition that the represented entity is a blackbird).[94]

The Perseus example (32) of section 3.3.1 can be 'collectivized' along the same lines as the bird example involving (23). Suppose that the guide isn't just addressing a single tourist but a group. She too may be facing an audience consisting of people that vary as regards their information about the artifact she is referring to; and in this case there may be even more diversity than in the bird case, for the people she is addressing may vary not only in that some have already noticed the statue while others have not, but also with regard to their antecedent information about the artist and the statue. In spite of this additional dimension of variation there is no significant difference between this case and the one of the people on the balcony: The members of the tourist group can be expected to process their guide's utterance of (32) in ways that may differ in their details but that nevertheless all capture the speaker's communicative intention. So in this case there is no need either for the speaker to make assumptions about what the various

[94]I owe the point of this and the next section to Arndt Riester. See in particular Riester (2008).

people in her group may know about the statue she is referring to, or may have noticed by looking around.

And a good thing that is too. For, first, speakers often don't know much about the information states of the different members in their audiences; and, second, even when they do know more about the information states of the individual members of a larger audience than we normally do, it would be a considerable nuisance to have to put this knowledge to use by talking to different members of the audience in different, tailor-made, ways instead of using a single formulation with which to address them en bloc. If custom-made communication was always necessary, then there would be little room for public address and – a far more serious impediment – for the use of language in print.

3.3.7 Choosing between definites and indefinites when addressing a crowd

The difference between definite and indefinite noun phrases isn't a principal part of the agenda of this paper, in which I have made an effort to limit attention to the behavior of definite NPs. But Articulated Contexts can also be used to explain an important aspect of the definite-indefinite distinction: Why is it that in some situations the choice between using a definite or an indefinite NP doesn't seem to matter very much, whereas in others it is crucial? In this and the next section I will try to give an idea of how the concept of an Articulated Context can throw light on this question.

Consider the following variation of our Perseus example, to be compared with our earlier (38).

(38) That there is a statue known as the 'Perseus', by a 16-th century Florentine artist called Benvenuto Cellini. It has recently been cleaned.

In the members of our guide's group (38) will produce the same sorts of information updates as (32), and that irrespective of the different information states from which each of them starts. This point deserves some emphasis, given that the linguistic principles involved in the interpretation of (32) and (38) are quite different. The crucial difference is between the interpretation rules for definite NPs like *Cellini's Perseus* in (32) and the demonstrative

phrase *that there* in (38) on the one hand and indefinite NPs like *a statue known as 'Perseus'* and *a 16-the century Florentine artist called Benvenuto Cellini* in (38) on the other. Recall once more the opposition between familiarity and novelty. Definite NPs come with a familiarity constraint: By using a definite NP the speaker conveys to the addressee that she takes him to be familiar with its referent, or to have the means needed for its independent identification. Indefinites come with the opposite constraint (the so-called Novelty Constraint): the speaker who uses an indefinite (and uses it specifically) signals to the addressee that she does *not* take him to be in a position to identify the entity she is talking about independently of what she is saying about it. These different signals define different tasks for the interpreter. The familiarity constraints (or identification presuppositions, in the story as I am telling it) that accompany the definites in (32) and (38) require him to come up with ERs from his Articulated Context to interpret these NPs. In contrast, the novelty constraint that accompanies the indefinites of (38) tells him that he is not supposed to be in possession of ERs that identify what those indefinites are being used to talk about. Rather, if he takes the speaker to be using the indefinites specifically – an assumption that in the case of (38) is almost inevitable – he will assume that she assumes that he doesn't have ERs for the referents of these indefinites.[95]

It would take a fair amount of time and space to work in detail through the different interpretations that the definites and indefinites in (32) and (38) can be expected to trigger throughout a 'mixed audience', whose members vary in what they know about the statue and its maker. Since I do not think that much can be learned from such an exercise that would be of interest for us here, I will make no attempt at this. (In the light of all that has been said about NP interpretation in context, readers shouldn't have any trouble in reconstructing how certain group members will process the different NPs.) What we should retain from this discussion is that in the scenarios described by our Perseus example and its variants the choice between the indefinites of (38) and the corresponding definites of (38) isn't all that important – either choice will do. Perhaps those in the tour group who know that the statue the guide is pointing at is the Perseus of Cellini may have a tendency to react to the indefinites in (38) with a 'Yes, yes, I know what that thing is.' And perhaps some of those who have never heard of Cellini and his work may

[95]See Kamp and Bende-Farkas (2018) for more discussion.

have a slight reaction along the lines of: 'Oh, should I have known about this already?' when the guide uses the words in (32). But that is probably going to be the limit of cognitive disruptions that (32) or (38) could produce in recipients with varying degrees of preparation. And no doubt the information that both (32) or (38) are meant to convey will be unproblematically accessible to anyone in the group.

3.3.8 When definites are fine and indefinites really bad

There is one aspect of the choice between definite and indefinite noun phrases, however, that deserves an explicit comment. It has to do with two kinds of tension there can be between the noun phrase that is chosen by the speaker/author and the information state that the recipient brings to the task of interpreting it: (i) the NP is definite but the recipient has no ER for the referent: (ii) the NP is an indefinite, but the recipient does have an ER for the referent. The first kind of tension has been discussed at length: the interpreter of a definite NP who cannot resolve its identification presupposition by using an ER he already has can be expected to accommodate a (vicariously anchored) ER for the referent. The other tension hasn't been discussed as yet. But in the light of what has been said about ER management in NP interpretation it would seem clear what the interpreter should do in case (ii): He should use the ER that he recognizes as representing the entity that he takes the speaker to be talking about, by identifying the referential argument of the NP with the distinguished discourse referent of the ER.

The remarks of this last paragraph suggest that neither tension ought to be much of an obstacle to successful communication: In either case the interpreter has an easy strategy for straightening out the wrinkle that the apparent mismatch between the chosen NP and his knowledge or ignorance of its referent might produce. But as a matter of fact that isn't so. Mismatches of the second kind can be serious obstacles to proper understanding. This might strike one as surprising, for why should having an ER for the entity that is being talked about be a problem, even when the speaker signals that she doesn't expect you to have one? Indeed, in many cases the mismatch is unproblematic. This is so in particular for the indefinites of (38). Someone in the group who already knows about Cellini and Cellini's Perseus won't have any difficulty in relating his ERs for the person and the artifact to the

indefinites the guide is using. After all she might just underestimate the group's preparation for the trip; or she might be right in thinking that some of the members are not as well prepared as he is and that it was to them that she is adjusting the way she is expressing herself.

But there are also cases where the unwarranted use of an indefinite can be truly confusing and unsettling to the recipient. The most striking examples of this that I am familiar with are those where the speaker introduces an individual into K_{dis} and then goes on to talk about this individual using an indefinite instead of an anaphoric definite. Here is an example.

(39) a. I just bought a new car

b. The dealer is going to bring it around tomorrow morning.

c. The dealer is going to bring a car around tomorrow morning.

d. The dealer is going to bring a new car around tomorrow morning.

Suppose that the speaker wants to say what is naturally expressed by the combination of (39a) and (39b) by using the sentence pair ((39a), (39c)), with the indefinite *a car* in (39c) replacing *it* in (39b), or by using the pair ((39a), (39d)), in which *it* has been replaced by *a new car*. These alternatives simply do not work. You just cannot use either *a car* or *a new car* and expect that your audience will readily understand you as talking about the car that you just introduced into the discourse through your use of the indefinite *a new car* in your first sentence. In such situations the novelty constraint comes down hard and irrevocably. What it militates against can be stated as follows:

(40) By using an indefinite a speaker or author signals that she is talking about something that is new in the sense that it is *not salient in the discourse context*.

That this principle is almost impossible to override is another indication of the transparency which characterizes the discourse context – something that we captured earlier by saying that the discourse context has the status of Common Ground: with regard to entities in the Common Ground there is no room for the supposition that the speaker might think that the entity she is talking about is unfamiliar to you. And so there is no room for the addressee to suppose that the speaker might be mistaken about this and for

him to leave the matter at that.

It is this that makes the difference between some of the mismatches of the second kind (the violations of the novelty constraint) and all of the mismatches of the first kind (the violations of the familiarity constraint). Whenever a speaker makes use of a definite NP to refer to something unknown to her addressee, thereby conveying her presumption that he is familiar with the referent, it is in principle always possible for him to assume that she is mistaken and to make up for the defect by accommodating. But when the discourse context makes clear that an entity is mutually known, you cannot construe the speaker as assuming erroneously that you aren't familiar with the entity; and because of that the option of correcting for this error through accommodation never comes to life.

What these considerations do not explain is that our attitudes towards the use of definites are even more tolerant than is entailed by them. As we noted earlier, speakers can use definites in situations where they have good reason to assume that the referent is unknown to their audience and in which they moreover assume that their audience must know this (i.e. where they are aware that the audience is not familiar with the referent and the audience for their part has good reasons to assume that the speaker knows that they don't know). Even in such situations recipients will normally be quite happy to play along, accommodating without demur the ER that they would have had if the familiarity constraint had been satisfied.

I do not quite see my way through to a satisfactory explanation of this remarkable degree of tolerance that we have towards the use of definites. Arguably this is just a special case of the more general phenomenon that speakers often knowingly violate the satisfaction constraints imposed by presuppositions: The speaker has reason to assume her interlocutor to be ignorant that a certain presupposition is satisfied, but uses an expression that triggers this presupposition nonetheless, with the intention that the addressee accommodate the presupposition, thereby taking its content on board as well as the asserted content. But subsuming a phenomenon under some other, more general phenomenon isn't much of a solution so long as no account is in sight for this second phenomenon. The observation that by insinuating presuppositions into our discourse that we take to be unfamiliar to our audiences, as an efficient way of supplying them with additional information

'through the back door', may have a certain intuitive appeal. But it seems to fall well short of what should pass as a proper explanation in either linguistics or philosophy. Alas there is no more I can offer on this point, and on this sobering note I drop the curtain.

4. Summary

The central goal of this paper has been to develop a notion of context which includes all the different kinds of information that may be needed in the interpretation of definite noun phrases. This project has been subdivided into two parts, a fairly traditional and a more revolutionary one. The more traditional part, in Section 2, deals with the problem how a unification can be achieved of the utterance contexts that have been part of Formal Semantics since its beginnings and the discourse contexts of Dynamic Semantics. The solution proposed for this problem is to enrich the vocabulary of the relevant DRS languages of DRT with indexical discourse referents' such as *sp*, *ad* and n, whose values are determined by (versions of the) character-to-content rules in the sense of Kaplan. These indexical discourse referents can then serve to represent the semantic contributions of the indexical expressions of English and other natural languages.

The starting point for the second part of the paper – all of Section 3 – is concerned with the information that interpreters use to make sense of discourse-new occurrences of definite descriptions and proper names. Here the proposal is that interpretations of discourse-new definites make use of Entity Representations. The notion of an Entity Representation is taken from MSDRT, a theory of the structure of mental states according to which important parts of such states are composed of Entity Representations and Propositional Attitudes.

The assumption that Entity Representations are the contextual ingredients used in the interpretation of discourse-new and other occurrences of definites (and also of specifically used indefinites) comes with a heavy methodological commitment. Since Entity Representations are constituents of mental states, a similar assumption forces itself upon us with regard to the comprehensive contexts (the so-called Articulated Contexts) that contain Entity Representations among their constituents; and once Articulated Contexts are

defined as parts of mental states, interpretations in which Articulated Contexts are used must be described in psychological terms as well. However, in the two-stage architecture of the framework that is assumed here the impact of this 'psychological turn' is mitigated in that it only affects the second of two interpretation stages, that at which the presuppositions of preliminary semantic representations are resolved. The first stage, at which those preliminary semantic representations are constructed from syntactically parsed input sentences, can still be described in user-independent terms.

One of the central aims of the paper is to rethink Kaplan's Theory of Demonstratives against the background of the framework presented here. The upshot of this exercise has been that the story about deictic uses of demonstrative NPs comes out as very different from the essentially Kaplanian story that is told about indexicals in Section 2. What, according to these two stories, indexicals and deictically used demonstrative NPs have in common is that both give rise to singular content; but that is something they share with many other NP uses. What sets indexicals apart from the deictically used demonstrative NPs, is that interpretation of the former involves mutually accessible elements from the utterance context, whereas interpretation of the latter is necessarily mediated by ERs whose presence in the interpreter's mind is a matter of contingency. (Because of their independence from this kind of contingency indexical NPs are unique among the NPs whose interpretation can produce directly referential contributions and lead to singular contents. In this way the indexicals regain an aura of uniqueness that in Kaplan's work they do not have.

The final parts of the paper (from Section 3.3.2 onward) explore, from the perspective of the communication-theoretic framework that has been outlined, some aspects of notions prominent within semantics and the philosophy of language: speaker intentions, speaker-hearer coordination and Common Ground, the relations between deixis and anaphora and those between anaphoricity and referentiality; and, finally, the differences between definite and indefinite NPs, as governed by the respective principles of Familiarity and Novelty.

Throughout I have tried to keep formal matters at a low level. The first part of the project in Section 2 is carried out with a fair amount of formal precision, in keeping with the level at which discussions about utterance

context and discourse context have been conducted in the literature. But the second part has been kept largely free of formalization (apart from the formal definitions of Entity Representations and Articulated Contexts and the schematic specifications of the forms of perception-based and vicarious anchors and the one explicit application of them in Section 3.3.1). I hope that keeping formalization at arm's length will have made the paper less inaccessible than it would otherwise have been to all but a select few familiar with DRT and MSDRT. But this should not be misinterpreted as a signal that I consider formalization unimportant. What genuine merit there may be to the framework that has been laid out here in largely informal terms depends ultimately on the possibility of explicit and detailed applications to particular utterances, conversations and texts; and such applications presuppose the construction algorithms of DRT and MSDRT.

A practical problem is that while substantial construction algorithms have been in place for some time (and have been applied in recent years, mostly in work carried out within the research project SFB 732 at the University if Stuttgart), a comprehensive description of them,which enables the interested reader to find out how she might put them to her own use, is still missing. All I can do at this point is to end with a plea that this state of affairs will be improved before long.[96]

References

S. Abney. *The English Noun Phrase in its Sentential Aspect.* PhD thesis, MIT, 1987.

[96]As far as DRT is concerned, a new introduction – (Kamp and Roßdeutscher (2019)) – is in preparation which presents a compositional treatment of meaning, starting at the sub-lexical level of roots and extending all the way up to full sentences, texts and conversations. This book will be a compendium of the rules and principles of a bottom-up construction algorithm which covers this entire span. Construction principles that are specific to MSDRT can be inferred from sample applications found in publications and unpublished work that has been cited earlier in this essay. (Kamp et al. (2011), Kamp (2005), Kamp (2011)). Here too a document is needed in which the syntax, model-theoretic semantics and construction rules of MSDRT are brought together. A preliminary version, consisting of a number of examples worked out in formal detail, with commentaries on the construction principles used, is being made accessible on the author's website [**add web address**]

N. Asher. Belief in discourse representation theory. *Journal of Philosophical Logic*, 15:127–189, 1986.

N. Asher. A typology for attitude verbs and their anaphoric properties. *Linguistics and Philosophy*, 10:125–198, 1987.

D. Beaver. *Presupposition and Assertion in Dynamic Semantics*. Folli, CSLI Publications, Stanford Unicersity, 2001.

D. Beaver, B. Geurts, and E. Maier. Discourse representation theory. *Stanford Encyclopedia of Philosophy*, 2007, 2015.

D.I. Beaver. Presupposition. In Johan van Benthem and Alice ter Meulen, editors, *Handbook of Logic and Language*, pages 939–1008. Elsevier, Amsterdam, 1997. URL `citeseer.nj.nec.com/beaver96presupposition.html`.

H. H. Clark. *Using Language*. Cambridge University Press, 1996.

M. J. Cresswell. *Logics and Languages*. Methuen, London, 1973.

D. Dekker. "she"'s character. In M Bras and L. Vieux, editors, *Semantic and Pragmatic Issues in Discourse and Dialogue*. Elsevier (CRiSPI Volume 9), 2001.

K. Donnellan. Proper names and identifying descriptions. *Synthese*, 21: 335–358, 1970.

G. Evans. *Collected Papers*. Oxford, 1985.

B. Frank and N. Goodman. Inferring word meanings by assuming that speakers are informative. *Cognitive Psychology*, 75:8096, 2014.

M. Frank and N. Goodman. Pragmatic language interpretation as probabilistic inference. ms., Stanford University, 2016.

K. Fraurud. Definiteness and the processing of noun phrases in natural discourse. *Journal of Semantics 7(4)*, 7:395–433, 1990.

B. Geurts. Good news about the description theory of names. *Journal of Semantics*, 14:319–348, 1997.

Bart Geurts. *Presuppositions and Pronouns*. Elsevier Science, 1999.

B. Grosz, A. K. Joshi, and S. Weinstein. Providing a unified account of definite noun phrases in discoure. In *ACL 83 Proceedings*, 1983.

B. Grosz, A. K. Joshi, and S. Weinstein. Centering: A framework for modeling the local coherence of discourse. *Computational Linguistics*, 21:203–226, 1995.

J. Gundel, N. Hedberg, and R. Zacharski. Cognitive status and the form of referring expressions in discourse. *Language*, 69:274–307, 1993.

I. Heim. *The Semantics of Definite and Indefinite Noun Phrases*. Garland Press, 1982,1988.

I. Heim. On the projection problem for presuppositions. *Proceedings of the West Coast Conference on Formal Linguistics*, 2:223–248, 1983.

I. Heim. E-type pronouns and donkey anaphora. *Linguistics and Philosophy*, 13(2):137–177, 1990.

J. Hunter. Presuppositional indexicals. *Journal of Semantics*, 30:381–421, 2012.

J. Hunter and N. Asher. *Now*: A presuppositional account of indexicals. In P. Dekker and M. Franke, editors, *Proceedings of the Fifteenth Amsterdam Colloquium*, pages 119–124, 2012.

H. Kamp. Evénements, représentations discursives et référence temporelle. *Langages*, 64:39–64, 1981a. For the original English ms. from which the paper in *Langages* is the French translation, see *Semantics and Pragmatics*, Vol 10 (2017.

H. Kamp. A theory of truth and semantic representation. In J. Groenendijk et al., editors, *Formal Methods in the Study of Language*. Mathematisch Centrum, Amsterdam, 1981b. Reprinted in: K. von Heusinger and A. Ter Meulen, 'Meaning and the Dynamics of Interpretation. Selected Papers of Hans Kamp'. Brill, 2013.

H. Kamp. Prolegomena to a structural account of belief and other attitudes. In C. A. Anderson and J. Owens, editors, *Propositional Attitudes—The*

Role of Content in Logic, Language, and Mind, chapter 2, pages 27–90. University of Chicago Press and CSLI, Stanford, 1990. Reprinted in: K. von Heusinger and A. Ter Meulen, 'Meaning and the Dynamics of Interpretation. Selected Papers of Hans Kamp'. Brill, 2013.

H. Kamp. The importance of presupposition. In Christian Rohrer, Antje Rossdeutscher, and Hans Kamp, editors, *Linguistic Form and its Computation*. CSLI, Standord, 2001a. Reprinted in: K. von Heusinger and A. Ter Meulen, 'Meaning and the Dynamics of Interpretation. Selected Papers of Hans Kamp'. Brill, 2013.

H. Kamp. Presupposition computation and presupposition justification: One aspect of the interpretation of multi-sentence discourse. In Myriam Bras and Laure Vieu, editors, *Pragmatic and Semantic Issues in Discourse and Dialogue*. Elsevier, 2001b. Reprinted in: K. von Heusinger and A. Ter Meulen, 'Meaning and the Dynamics of Interpretation. Selected Papers of Hans Kamp'. Brill, 2013.

H. Kamp. Einstellungszustände und einstellungszuschreibungen in der diskursrepräsentationtheorie. In Ulrike Haas-Spohn, editor, *Intentionalität zwischen Subjektivit"at und Weltbezug*, pages 209–289. mentis, Paderborn, 2003.

H. Kamp. Temporal reference inside and outside propositional attitudes. In K. von Heusinger and K. Turner, editors, *Where Semantics Meets Pragmatics*. Elsevier, 2005. Reprinted in: K. von Heusinger and A. Ter Meulen, 'Meaning and the Dynamics of Interpretation. Selected Papers of Hans Kamp'. Brill, 2013.

H. Kamp. Representing *de se* thoughts and their reports. unpublished, 2011.

H. Kamp. Using proper names as intermediaries between labelled entity representations. *Erkenntnis*, 80:263–312, 2015.

H. Kamp and A. Bende-Farkas. Epistemic specificity from a communication-theoretic perspective. to appear in Journal of Semantics, 2018.

H. Kamp and U. Reyle. *From Discourse to Logic*. Kluwer, 1993.

H. Kamp and U. Reyle. Discourse representation theory. In K. von Heusinger, C. Maienborn, and P. Portner, editors, *Handbook of Semantics*. De Gruyter, 2011.

H. Kamp and C. Rohrer. Tense in texts. In Bäuerle, Schwarze, and Von Stechow, editors, *Meaning, Use and Interpretation of Language*, pages 250–269. De Gruyter, Berlin, 1983.

H. Kamp and A. Roßdeutscher. From roots to riches. ms. university of Stuttgart, 2019.

H. Kamp and M. Schiehlen. *How to say WHEN it happens*, chapter Temporal Location in Natural Language. Niemeyer, 2002.

H. Kamp, J. van Genabith, and U. Reyle. Discourse representation theory: An updated survey. In Dov Gabbay, editor, *Handbook of Philosophical Logic, Vol. XV*, pages 125 – 394. Elsevier, 2011.

D. Kaplan. Demonstratives. In H. Wettstein J. Almog, J Perry, editor, *Themes from Kaplan*. Oxford University Press, 1989.

S. Kripke. A puzzle about belief. In A. Margalit, editor, *Meaning and Use*. Reidel, Dordrecht, 1979.

D. Lewis. General semantics. *Synthese*, 22:18–67, 1970.

D. Lewis. Attitudes de dicto and de se. *The Philosophical Review*, 88(4): 513–543, 1979a.

D. Lewis. Score keeping in a language game. *Journal of Philosophical Logic*, 8:339–359, 1979b.

E. Maier. *Belief in Context: Towards a unified semantics of* de re *and* de se *attitude reports*. PhD thesis, Radboud University, Nijmegen, 1998.

E. Maier. Proper names and indexicals trigger rigid presuppositions. *Journal of Semantics*, 23:253–315, 2009.

R. Montague. Universal grammar. *Theoria*, 36:373–398, 1970.

Ch. W. Morris. *Signs, Language and Behavior*. Oxford University Press, 1946.

D Nadeau and S Sekine. A survey of named entity recognition an classification. *Linguisticae Investigationes*, 30(1):3–26, 2007.

N.J. Nilsson. *Principles of Artificial Intelligence*. Canbridge University Press, 2010.

G. Nunberg. Indexicality and deixis. *Linguistics and Philosophy*, 16:1–43, 1993.

G. Nunberg. Indexical descriptions and descriptive indexicals. In *Proceedings from Semantics and Linguistic Theory IX (SALT9)*. CLC Publications, Cornell University, 1999.

J. Perry. A problem about continued belief. *Pacific Philosophical Quarterly*, 61:317–332, 1980.

J. Perry. *Reference and Reflexivity*. CSLI, Stanford, 2001.

J. Pfau, T. Miller, L. Sonenberg, and Y. Kashimac. Logics of common ground. 2015.

M Poesio and R Vieira. A corpus-based investigation of definite descriptions. *Computational Linguistics, 24(2)*, pages 183–216, 1998.

Fr. Recanati. *Mental Files*. Oxford, 2012.

H. Reichenbach. *Elements of Symbolic Logic*. Macmillan, London, 1947.

U. Reyle, A. Rossdeutscher, and H. Kamp. Ups and downs in the theory of temporal reference. *Linguistics and Philosophy*, 30(5):565–635, 2007.

A. Riester. A semantic explication of 'information status' and the underspecification of the recipients' knowledge. In *Proceedings of Sinn and Bedeutung 12*, pages 508–522. Oslo, 2008.

C. Roberts. Information structure: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics*, 5(6):1–69, 2012.

Ph. Schlenker. A plea for monsters. *Linguistics and Philosophy*, 26:29–120, 2003.

J. Searle. Proper names. *Mind*, 67:166–173, 1958.

J Spenader. *Presuppositions in Spoken Discourse*. PhD thesis, Stockholm University, 2002.

Andreas Stuhlmüller. *Modeling Cognition with Probabilistic Programs: Representations and Algorithms*. PhD thesis, MIT, 2013.

Richmond Thomason and Anil Gupta. A theory of conditionals in the context of branching time. *The Philosophical Review*, LXXXIX(1):65–90, 1980.

J. Tonhauser. A dynamic semantic account of the temporal interpretation of noun phrases. In B. Jackson, editor, *Proceedings of Semantics and Linguistic Theory (SALT) XII*, pages 286–305. Cornell Linguistics Circle Publications, 2002.

R.A. Van Der Sandt. Presupposition projection as anaphora resolution. *Journal of Semantics*, 9:333–377, 1992. Special Issue: Presupposition, Part 2.

R.A. Van Der Sandt and B. Geurts. Presupposition, anaphora, and lexical content. In O. Herzog and C.-R. Rollinger, editors, *Text Understanding in LILOG*. Springer, 1991.

B. van Fraassen. Presuppositions, supervaluations and free logic. In K. Lambert, editor, *The Logical Way of Doing Things*. Clarendon Press, Oxford, 1975.

L. van Leusen and R. Muskens. Construction by description in discourse representation. In J. Peregrin, editor, *Meaning: The Dynamic Turn*, pages 33–65. Elsevier, 2003.