How we say WHEN it happens

Contributions to the theory of temporal reference in natural language

Edited by
Hans Kamp and Uwe Reyle

Max Niemeyer Verlag
Tübingen 2002

Die Deutsche Bibliothek – CIP-Einheitsaufnahme

How we say when it happens : contributions to the theory of temporal reference in natural language / ed. by Hans Kamp and Uwe Reyle. – Tübingen : Niemeyer, 2002
(Linguistische Arbeiten ; 455)

ISBN 3-484-30455-3 ISSN 0344-6727

Gedruckt auf alterungsbeständigem Papier.
Druck: Weihert-Druck GmbH, Darmstadt
Einband: Industriebuchbinderei Nädele, Nehren

Hans Kamp and Michael Schiehlen

Temporal Location in Natural Languages

Summary

This paper deals with the semantics of various types of noun phrases which refer to time, as well as some phrases of related categories, such as common noun phrases and prepositional phrases. A prominent place play "calendar expressions", expressions such as *year*, *month*, *day*, *hour*, *February*, *Wednesday*, *Easter*, etc. The semantics we pursue in this paper is a model-theoretic one and its first part (Section 2) is devoted largely to the development of a model concept that can serve as the foundation for what is to come after. Since calendar terms take centre stage in our explorations, it is crucial that they can be given plausible interpretations in the models we use, models which reflect the typical speaker's knowledge about time in general and about the meanings of those terms in particular. This makes it necessary to look quite carefully into the models' underlying time structures, which serve as the substratum for the model-theoretic interpretation of calendar terms and other temporal expressions. Besides, the model theory of temporal concepts developed here is also intended as the semantic basis for a conceptually plausible definition of "time-logical" consequence, which incorporates general knowledge about time and about the calendar. The second part (Sections 3, 4, and 5) deals with the interpretation mechanisms of a variety of time-referring expressions, mostly at an informal or semi-formal level. A good deal of attention is paid to the context-dependent aspects of various temporal expressions. One of our aims in this part is to attain a better understanding of the different kinds of context sensitivity, and in particular of the difference between anaphora and indexicality. We have had to abandon our original plan of concluding the paper with an explicit model-theoretic treatment of a representative "fragment" of temporal NPs and PPs (generated by a set of phrase structure rules from a reasonably sized lexicon). We hope that the informal analyses in the second part will give a fairly clear idea of how such analyses work in general.

1 Introduction

The observations reported in this paper originated in the speech-to-speech translation project Verbmobil, in which German utterances were translated on-line into English. The dialogues which Verbmobil aimed at translating were, for most of the VM project, of a quite restricted kind: Their exclusive aim was to fix time and place of one or more future meetings between the dialogue participants. In these dialogues reference to (future) times was – obviously – a central issue, and not surprisingly, "calendar" expressions like *der siebte Oktober* (the seventh of October), *die dritte Maiwoche* (the third week in May), *Montag zwischen zwölf und zwei* (Monday between twelve and two), *nächster Mittwoch* (next Wednesday), *an diesem Mittwoch* (that Wednesday), etc. play a particularly prominent part. The translation of such expressions from German into English is by and large not especially problematic. For the semantic intricacies of the German phrases and their English counterparts usually match, so that a translator need not worry about them. Nevertheless there are a number of intricacies connected with these expressions of which we became aware only when Verbmobil led us to think more carefully about the systematic aspects of their interpretation.

Salient among these – in our own current state of mind and in the following pages which reflect it – are in particular two points: (i) The difference between purely semantic and epistemic factors in our use of such expressions – it is common for us to be unsure whether, say, the third of next month is a Wednesday (and, by implication, the first Wednesday of next month) or some other day of the week, although the denotations of the terms *the third of next month* and *the first Wednesday of next month* are fully determined and we know in principle how to compute them; and (ii) The special character, within the temporal domain, of context dependence, which shows a blend of anaphora and indexicality not known from the by now well-known studies of third person as opposed to first and second person pronouns.

In order to account for these and other aspects of time-referring expressions we found it necessary to provide ourselves with an explicit and detailed model-theoretic description of time as it informs human discourse. An important part of this description is concerned with the ways in which the temporal axis is partitioned by calendar predicates such as *year*, *month*, *week*, *day* and the like. We are aware that this (earlier) part of the paper makes for uncommonly tedious reading. But the model-theoretic set-up is indispensable for what follows, which we hope will not seemm tedious to quite the same degree.

One pervasive aspect of the semantics of the NPs with which we will deal is context-sensitivity. This is especially obvious, and has been extensively discussed, for expressions like *tomorrow* or *last week*. But the phenomenon is much more pervasive than that; it attaches equally to *the first of September*, *Sunday* or *half past ten*. An occurrence of *the first of September* denotes a particular day only insofar as the context fixes which year is concerned; in the case of *half past ten* the context must determine the day to which the referent belongs; and to fix the reference of an occurrence of *Sunday* the context must determine a temporal domain containing a unique Sunday. We will look at the principles which govern these kinds of contextual resolution in Section 3.

For now we only note that the context sensitivity of many temporal NPs is directly connected with the categorial ambiguity of expressions such as *Sunday* or *Christmas*. We have seen that these terms can be used as (context-sensitive) NPs. But they can also be used as common nouns, as in *next Sunday*, *every Sunday*, *the Christmas before last*, etc. One obvious connection between these two uses is that the referent of an NP-occurrence always belongs to the extension of the common noun. For instance, whatever day the NP *Sunday* denotes, it will always be a Sunday. So the semantic relationship between NP-occurrence and common noun is the same as that between occurrences of the NP *the first of September* and the complex common noun *first of*

September. This latter relationship is one that has long been seen as crucial to any compositional semantics of these and other definite descriptions: the denotation of *the first of September* should be analyzed as determined by what the singular definite article *the* does with the meaning of the common noun *first of September*. The semantics we will present in this paper endorses this principle. And it will account, in an analogous way, for the reference of 0-determiner NPs such as *Sunday* on the basis of the extension of the corresponding common noun. It is for this reason that we will begin our exploration of temporal expressions with a discussion of such common nouns and turn only after that to the problem of context-dependent NP reference. But before we can start any part of these investigations, it will be necessary to establish the general framework they presuppose.

While the impetus for writing this paper came from our involvement in Verbmobil, in which Christian Rohrer has played a central part throughout, many of the paper's actual concerns go back to issues on which he and the first author worked in the eighties (and to which an entire chapter was devoted in the still unpublished work on temporal reference in French which they then wrote). Thus the following pages would seem a suitable tribute to someone in whose earlier work tense and time played so prominent a part and who then focused increasingly on the computational concerns of which his prominent role in Verbmobil has been one of many manifestations.

2 Model-theoretic Framework

2.1 Temporal Order and Metric

Our approach follows the familiar tenets of model-theoretic semantics. Each expression of the part of natural language we will study has a denotation. This denotation depends on the one hand on the meaning of the expression and on the other on the factual properties of the world. By abstracting over the latter it is possible to focus on the former. One does this by treating the world as instantiating one from a class of possible models. Each possible model M represents one possible world (or a cluster of possible worlds between which it is unnecessary to distinguish in the context of the semantic investigation at hand). Given such a class of possible models we can identify the semantic context of an expression e with the function which maps each possible model M to the denotation of e in M.

A precise characterization of this function presupposes a precise characterization of its domain – i.e. of the class of models to be considered. It is important that this class be "just right". On the one hand it should include among its models a counterpart to each of the possible contingencies that might have their effect on what the denotation of various expressions is. But on the other it should not be too big. Models which are inconsistent with semantic properties of and relations between expressions of the language under consideration, ought to be excluded.[1]

[1] For instance, it seems reasonable that the English adjectives *dead* and *alive* are mutually

In actual practice model-theoretic treatments of natural language have usually followed the strategy of approximating the intended model classes "from above": When developing a model theory for a language or language fragment L one starts out with a model class C_o that includes more than it should and then restricts this class by adopting "meaning postulates", principles which capture certain semantic properties of and relations between expressions of L. This is more or less the way in which we will proceed. We will formulate various conditions that our models will have to satisfy. As more such "meaning postulates" are added, some of the models that were previously admitted will be eliminated.

The question what models should be included in the model class that is needed for our purposes can be divided into several parts. The first and most basic question concerns the structure of the temporal order and its metric. Here we encounter a difficulty which will confront us equally when we come to address the remaining questions: How are we to determine "the" way in which language users conceive of time? It may be expected that the conceptions that different speakers have of time do not fully coincide, so that the best we can hope for is to capture something like the "conception of the typical speaker". But who qualifies as the typical speaker? That is hard to tell. And in the absence of extensive empirical investigation as to what it is that different speakers assume about time, one is likely to identify what one takes to be the typical speaker's conception with one's own intuitions. We realize that more likely than not we are victims to such prejudice ourselves and that the proposals we will make about what should be regarded as a conceptual truth about time (and thus about what holds in all models of our class) are probably not free from some measure of arbitrariness. We do not think, however, that this is an issue that seriously affects what else we will have to say.

The first assumption we make about the conceptual properties of time is that it is a dense, linearly ordered medium. Perhaps the intuitively most satisfying way of describing this is as a structure consisting of "periods" which are linearly ordered by complete precedence and overlap and in which all periods (excepting the minimal, "punctual" elements of time, if these are to be called periods at all) are infinitely divisible. One can also think of the periods as collections from an antecedently given point set. This is arguably more in keeping with mathematical tradition. But for our present purposes the choice between these two approaches does not really matter. (See e.g. van Benthem 1984 or Kamp 1979 for discussion of how the two approaches are related.) What will be important, however, – no matter which approach is chosen – is the distinction between the non-minimal periods and the minimal ones. We will refer to the latter as *moments of time* and to the former simply as *periods.*

More precisely, we will assume, as a basis for further model-theoretic considerations, a topological time structure consisting of both moments and periods, in which the moments are arranged in a dense linear ordering $<$, there is a relation $\in$ between moments and periods ("$i \in p$" means that the moment i belongs to the period p), and in which there is for each pair $\langle i1, i2\rangle$ such that $i1 < i2$ one corresponding period

exclusive: that nothing is dead and alive at the same time is a matter of meaning, not an accident of how the world happens to be. Accordingly, the model class used in a semantics for a fragment of English including these two adjectives should contain no models in which something belongs to the extensions of both.

$(i1, i2)$. (The question whether $i1$ and/or $i2$ belong to $(i1, i2)$ is of no empirical significance and has to be settled by convention.) The temporal relations between periods of which we will make use include: (i) The relation $<_P$ of complete precedence (we will often simply write $<$ instead of $<_P$, as using the same symbol for the ordering relations between moments and between periods is usually harmless); and (ii) the relation O of overlap familiar from the cited literature (again, see Kamp 1979 or van Benthem 1984); we assume that if $p < p'$ then all moments belonging to p precede all moments belonging to p' and that any two overlapping periods have at least one moment in common. In addition to $<$ and O we also will be making use of (iii) the relation $\supset\!\subset$ of abutment and (iv) the relation $\subseteq$ of temporal inclusion. $p_1 \supset\!\subset p_2$ holds iff $p_1 < p_2$ and for no period p_3, $p_1 < p_3 < p_2$. (Thus when p_1 is the period $(i1, i2)$ and p_2 is the period $(i3, i4)$, then $p_1 \supset\!\subset p_2$ just in case $i2 = i3$.) $p_1 \subseteq p_2$ holds whenever every moment which belongs to p_1 also belongs to p_2. (Thus, if p_1 is included in p_2, then every period p which precedes p_2 also precedes p_1 and every p which is preceded by p_2 is also preceded by p_1.) In addition, we will use the notation "$p_1 + p_2$" to indicate the sum period of two abutting periods which covers all the moments belonging to p_1 as well as all those belonging to p_2. (Given the correspondence between periods and pairs of moments this sum will always exist: if p_1 is $(i1, i2)$ and p_2 is $(i2, i3)$, then $p_1 + p_2$ is the period $(i1, i3)$.)

Exactly how this type of structure is to be defined axiomatically, and in particular whether periods are defined in terms of moments, moments in terms of periods, or whether both kinds of entities are treated as primitives, we leave open.

Our conception of time is not just of an ordered, but of a metric structure: Besides the structural elements already mentioned, people take it to be part of time that certain periods stand in a relation $\equiv$ of "equal duration". $\equiv$ is an equivalence relation, which is assumed to obey a number of further assumptions. We quickly run through those assumptions which we consider essential.

The first assumption is that $\equiv$ is a congruence relation with respect to forming sums: If $p_1 \supset\!\subset p_2$, $p_3 \supset\!\subset p_4$, $p_1 \equiv p_3$ and $p_2 \equiv p_4$, then $p_1 + p_2 \equiv p_3 + p_4$.

The second assumption has to do with size comparison. The basic intuition on which it is based is one that applies to periods p_1 and p_2 whose beginnings coincide. In such a case p_2 is longer than p_1 if, intuitively, at the moment when p_1 has come to an end, additional time is needed to run through what remains of p_2. The natural way to express this intuition would seem to be this: The additional time needed to get through the remainder of p_2 is a period p_3 such that $p_2 = p_1 + p_3$. The principle of Size Comparison, in its restricted form, then says that if two periods p_1 and p_2 begin at the same time, then they are either identical or else one is longer than the other: (i) $p_1 = p_2$ or (ii) $(\exists p_3)p_2 = p_1 + p_3$ or (iii) $(\exists p_3)p_1 = p_2 + p_3$.

It seems to be part of our intuitions about $\equiv$ that size comparison applies not just to the special case of periods with coincident beginnings, but to any pair of periods whatever: If p_1 and p_2 are any two periods, then either (i) $p_1 \equiv p_2$ or (ii) $(\exists p_3, p_4)p_2 = p_3 + p_4 \wedge p_3 \equiv p_1$ or (iii) $(\exists p_3, p_4)p_1 = p_3 + p_4 \wedge p_3 \equiv p_2$. Abbreviating $(\exists p_3, p_4)p' = p_3 + p_4 \wedge p_3 \equiv p$ as pNp', we can state this generalized Size Comparison Principle in the following condensed form:

- (SCP') $p_1 \equiv p_2 \vee p_1 N p_2 \vee p_2 N p_1$.

In fact, the intuitions pertaining to size comparison go, it seems to us, further than this. They include, we believe, the idea that $\equiv$ and N form a strict pseudo-ordering: (i) N is transitive and asymmetric; (ii) $\equiv$ is a congruence with respect to N; (iii) (SCP'); and (iv) the disjunction (SCP') is exclusive.

In addition to Size Comparison we take the following three conditions to be intuitively justified:

- Repartition: Whenever $p_1 \supset\!\subset p_2$, there are p_3 and p_4 such that $p_3 \equiv p_2$, $p_4 \equiv p_1$ and $p_3 + p_4 \equiv p_1 + p_2$.
- Halving. For each p there are p_1 and p_2 such that $p_1 \equiv p_2$ and $p_1 + p_2 = p$.
- Metric Connectedness.

The next condition we consider concerns a certain form of "metric connectedness". No matter how far two periods p and p' are apart, it is always possible to reach the one by "multiplying" the other. More precisely, we can reach p' by a finite chain of pairwise abutting periods all the size of p and starting at p. Or, in other words, there are some finite number n of periods $p_i \equiv p$ which abut as in the diagram and where p_n overlaps with p':

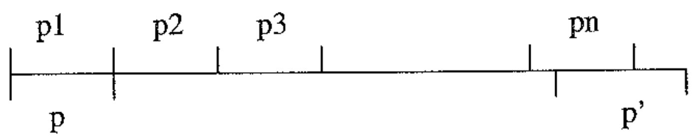

Structures $\langle T, <, \equiv \rangle$ satisfying this condition are called Archimedean.

In addition to all the assumptions mentioned so far we will also assume that time is infinite in both directions. In fact we will assume this in a strong form, according to which there exist, for any combination of periods p_1 and p_2, periods p_3 and p_4 which are each of the same duration as p_2 and which abut p_1 on the right and on the left, respectively:

- Metric Unboundedness: For any periods p_1 and p_2 there exists (i) a period p_3 such that $p_3 \equiv p_2$ and $p_3 \supset\!\subset p_1$, and (ii) a period p_4 such that $p_4 \equiv p_2$ and $p_1 \supset\!\subset p_4$.

Metric Unboundedness seems conceptually problematic in a way that, according to our perception, the other conditions are not: How plausible is it to attribute this condition to the normal user of a natural language? Not much, it would seem, especially now that it has become an item of general physical belief (or at least speculation) that time did have a beginning, and may well have an end. We have decided to include this condition nevertheless, since it makes it easier to formulate our assumptions about the extensions of calendar predicates, to which we turn in the next section.

How much room do these assumptions leave for contingency? Formally speaking there is some, though for our present purposes it does not seem to amount to much. For instance, our assumptions do not differentiate between the real and the rational numbers: they admit both structures whose moments are the reals, and whose periods are the intervals of reals, and structures where the moments are the rationals and the periods rational intervals. In fact, there are innumerable (that is, non-denumerably many) other structures not isomorphic to any of these, which are admitted as well. However, from a conceptual perspective the differences between these structures do not seem to matter, and for our purposes it appears harmless to operate on the assumption that the order and metrics of time are determined completely. Hence our advice to the reader to focus in what follows on a particular structure familiar to him, e.g. that of the reals with their familiar metric (which among other things assigns the same duration to all intervals $(n, n+1)$ for integral numbers n). At the same time we stress that for computational purposes thinking of the reals as a given time structure won't quite do. In that context what we need is an axiomatic theory which, in combination with the calendar-related axioms we will discuss towards the end of section 2, will enable the system to compute (canonical representations of) the denotations of the time denoting terms belonging to the language fragment it covers.

There is a difficulty here, connected with the Archimedean principle of Metric Connectedness. As stated, this principle is not first order (because of the unbounded existential quantification over the number n of periods needed to get from p to p'). We conjecture that for the computation of canonical representations of calendar term denotations and of temporal relations between those, Metric Connectedness is not needed and thus that an explicit axiomatization consisting of the other (first order) principles of this section together with the calendar-related postulates (CCi) of which we give some illustrative examples in Section 2.6 will suffice for that purpose.

2.2 Measure Nouns and Calendar Nouns

Very different issues are raised by the problem what assumptions should be counted as part of the typical speaker's understanding of calendar concepts. What information do people typically connect with such notions as that of a year, a month, a week, a day, an hour, etc.? This is the question which will occupy us in this and the next section.

The first point that needs to be cleared, and which is in a way preliminary to the real issue, is this. The nouns *year*, *month*, *week*, are systematically ambiguous (Arguably this is true of the words *day* and *hour* as well, but we will concentrate on the first three, for which the ambiguity we have in mind is beyond question.) In the first place, each of these three words can be used to denote any member of a certain temporal *duration*, i.e. of a certain equivalence class under the relation $\equiv$. Thus *year* can be used to denote any period of the "duration of a year", irrespective of when it begins and ends; a *year* in this sense may coincide with a calendar year – it may start on the first of January and end on the thirty first of December – but it need not. This is the way we use *year* in a sentence like *It has been exactly a year since I saw her last.* and more generally in all measure noun constructions (e.g. *two years later*) which can be uttered truly at any time, not just at the moment which separates the old calendar year from the new one. The words *month* and *week* allow for this kind of use also – just replace in the phrase just quoted *year* by *month* or *week*. We note in passing a certain ambivalence in the interpretation of *month*, which has to do with the fact that not all calendar months are of the same length. For instance, when you say on the 28-th of February that you saw Bill exactly a month ago, is this to be interpreted as the claim that you saw him on the 28-th of January, or that you saw him on the 31-st of January? To the extent that we can make out our own intuitions and those of whom we have asked, both interpretations seem possible.

Besides this first use, the nouns *year, month, week* can also be used to refer to *calendar* years, months and weeks, respectively. A calendar year is always a period running from the first of January through the 31-st of December 365 or 366 days later, the calendar months are the periods which we also denote as *January, February*, etc., and the calendar weeks are the seven day periods which start with a Sunday and end with the following Saturday (or, for others, start with a Monday and end with the following Sunday). It will be useful to have a succinct way of distinguishing the calendar concepts from the measure concepts. So we introduce for each of the three nouns two different predicates, distinguished from each other by the subscripts $_m$ and $_c$. The extension of $year_c$ will consist exclusively of the calendar years, while that of $year_m$ will be the entire equivalence class (w.r.t to the relation $\equiv$) generated by the first extension; and similarly for *month* and *week*[2]. Evidently the $_c$-predicates have extensions that are much smaller than those of the corresponding $_m$-predicates. For if p_1 is, say, a calendar year, then the first calendar year after it is the period p_2 of that length which abuts it on the right; if we assume that time is dense, then an infinity of years in the sense of $year_m$ lies between p_1 and p_2 in that they all start after p_1 and end before p_2.

With other period-denoting nouns this ambiguity is either absent or dubious. Thus *minute* and *second* are only used as $_m$-predicates. (The underlying representation of clock times, however, will have to make use of the corresponding $_c$-predicates) This seems to be the case by and large also for *hour*, although this word may still have a marginal use as a $_c$-predicate, as when we speak of the first hour of the day for instance. In the past this use seems to have been much more common. *Day*, finally, is primarily used as a $_c$-predicate; at least, to our ears it is funny to say, at three o'clock in the afternoon: *It is exactly one day since Fred left.* when what you want to say is that Fred left exactly 24 hours ago.[3]

It is the predicates $year_c$, $month_c$ and $week_c$ that contribute significantly to the "calendar structure" of time. Crucial to that contribution is that each of their extensions constitutes a *complete partition* of the time line. This follows from the more general assumption we made above that for any period there are abutting periods of the same duration on its left and on its right and that by passing from such periods to those abutting it it is possible to reach any moment of time whatever. The same is true for the extensions of $month_c$ and $week_c$. Moreover, the partition induced by $month_c$ is a refinement of the extension of $year_c$, in the sense that each member of

[2] Actually, this isn't strictly true. See the last two paragraphs of this section. Nevertheless it seems to us to capture the conceptual connection between $_c$-predicates and $_m$-predicates correctly.

[3] We can also use *day* in what seem to be measure noun constructions, such as *three days later*, but where it is nevertheless the periods from one midnight to the next that are involved. For instance, *He had the accident on Monday. But the pain started only three days later.* would normally be interpreted as meaning that the pain started on the third day after the day of the accident.

the former is wholly included in one member of the latter. But unfortunately the extension of $week_c$ is not a refinement of the extension of $month_c$, or even, for that matter, of that of $year_c$, a source of some confusion in actual life, and a point about which more below.

It is this succession of ever finer partitions of time which provides the basis for all other calendar-related concepts and terms. In particular, we have a fairly complex system of "canonical names" for the elements of these partitions. This system has its idiosyncrasies both as regards the actual forms of the names used and the concepts which determine their referential semantics. Most of the conceptual idiosyncrasies arose as man discovered the mismatches between the different cosmic phenomena in terms of which the notions *day*, *month* and *year* were originally "defined" – the change between light and dark, the phases of the moon and the change of the seasons. The result is an odd mixture of the rational and the accidental.

In addition there are the special aspects of the concept of a week. Whether the fact that one week has seven days is a reflection of a natural rhythm in the life of modern man – some kind of genetically determined preference for an alternation of work and relaxation based on a cycle of seven days – seems doubtful.[4] But whether there is some deeper reason for the seven day week or not, there is arguably no other calendar concept which so thoroughly pervades and shapes the structure of our social life. One way in which this shows is that it is more common for people not to know which day of the month it is than to be unaware of the day of the week. Also, the circumstance that our social lives are so strongly structured by the weekly cycle creates a curious kind of complexity in the ways we often refer to particular days. Rather than referring to, say, a suitable day for an appointment as the n-th day of month x one will refer to it as the Wednesday of the second week of x; for the very fact that that day is a Wednesday may be a reason for assuming that that is a day on which one will be free – since normally one is free on a Wednesday. From the information that a certain day is, say, the 17-th of the month, you may not be able to conclude this straightaway, for in order to draw that conclusion you have to know first what day of the week the 17-th of that month is. And on the whole we are pretty bad at keeping track of how days of the month (identified by number) correspond to days of the week (identified by weekday name). Thus, to find out what day of the week is the seventeenth, we often have to do a little computing.

The special role that week-related concepts play in our social life and, as a consequence, in our cognition, is thus a crucial factor in the terms which people choose to refer to times. This is an issue that demands special attention and we will postpone it to the next Section. For the moment we focus on the other calendar predicates: *year*, *month* and *day*.

What needs to be added to the notion of a model as specified in Section 2.1 (i.e. as any densely ordered point-and-period structure with a metric congruence relation $\equiv$) if it also is to model these calendar concepts? The answer is straightforward. An

---

[4] The concept of a *week* seems to have been subject to considerable variation. In other cultures one finds weeks with 3, 4, 5, 6, and 8 days. See for instance the entry for *week* in the Encyclopaedia Britannica. It is legitimate to speak of *weeks* in such cases because of their institutional role and/or their historical continuity with the week as we know and have it today.

intuitively plausible interpretation of the calendar predicates $year_c$, $month_c$ and day_c requires two decisions. First, we have to choose a "zero point of the calendar". In practice mankind has had an understandable preference for real or fictitious event times such as the birth of Christ or Mohammed's flight from Mekka, but the matter is strictly one of convention and from a formal point of view any other time would do as well. Secondly, it is necessary to fix at least one instance of one of the calendar predicates. The construction which best fits both the actual facts concerning the extensions of these predicates and common intuitions about how the calendar works, is that which takes day_c as basic. And to fix the extension of this predicate in terms of what structure our models possess already it suffices to determine just one calendar day, the one which starts at the zero point. The general constraints on the extensions of calendar concepts will generate from this one calendar day the entire extension of day_c. And the extensions of all other calendar concepts can then be determined in terms of this one by appeal to the familiar conventions about how many days there are in a calendar year (365 or 366) and how many days there are in each of the calendar months.[5] As we already noted, the effect of this is that not only the elements of the extension of $month_c$, but also those of the extension of $year_c$ often fail to stand to each other in the relation $\equiv$; this too is something that most speakers are aware of. (We take it to be part of the knowledge about time of a typical speaker of our culture that the year consists of 365 days, except for leap years, which have 366, and of how the year is subdivided into months.)

In this way we narrow down the class of intended models outlined in Section 2.2. We expand the models described there by adding specifications which give for each model, besides its densely ordered metric point-and-period structure, (i) a zero point t_0 and (ii) a second point t_1 to the right of t_0. The period (t_0, t_1) is to be understood as the first day starting at the zero point and thus is an element of day_c. This then determines, given the constraints stated in the preceding paragraph, the full extensions of the three calendar predicates discussed there.

If calendar predicates and measure predicates were actually related in the way they appear to be connected conceptually, then the extensions of the predicates day_m, $week_m$, $year_m$ and $month_m$ could be identified with the closure under $\equiv$ of the extensions of the corresponding $_c$-predicates. But, as noted, the reality of our calendar isn't that neat. For obtaining the extensions of day_m and $week_m$ the simple recipe seems correct, but for $year_m$ and $month_m$ this is clearly not so. In fact, neither the extension of $year_c$ nor that of $month_c$ is included within a single equivalence class to begin with, and this is something that most speakers know. It is also widespread knowledge that according to the accepted chronometrical conventions no member of the extension of $year_c$ belongs to the extension of $year_m$. How much about the relationship between the two extensions should be considered common knowledge we find hard to assess. But we suppose that most speakers subscribe to the principle that the two extensions are commensurable in the following sense: Even for large n, a period consisting of

[5] We ignore all complications which arise from the fact that the calendar has been reformed at specific times in history. Even with regard to the times preceding such changes we assume that the extensions of the calendar predicates are determined according to our *present* definitions, not according to the criteria applied by those living, and using the calendar concepts, at the time.

n abutting calendar years is roughly of the same duration as a period consisting of n abutting members of the extension of $year_m$. So we suggest that the extension of $year_m$ in any model M in which the extensions of the $_c$-predicates have already been fixed be chosen in such a way that this commensurability principle is satisfied.

The relationship between $month_m$ and $month_c$ is, as we saw, even more problematic. Here, it seems to us, not even the commensurability principle is valid: For all we can tell, some of the uses which speakers make of $month_m$ do not fit, nor are they assumed to fit, this principle. So, if the range of admissible models is to cover all uses, the only constraint that we think it would be reasonable to impose on the extension of $month_m$ is that the duration of its members lies somewhere between 28 and 31 days.

Once the extension of day_m has been fixed, those of the remaining measure predicates, $hour_m$, $minute_m$ and $second_m$, can be determined by subdividing the members of the extension of day_m in the familiar way. (Of course, the "order" in which these extensions are determined is irrelevant so long as they end up standing to each other in the correct numerical relations.)

We should emphasize once more that we see most of the proposals we have made in this and the previous section as open to modification or refinement. This is true especially of the last two, concerning $year_m$ and $month_m$. For the sequel, however, it is not decisive whether the reader agrees with all the proposals we have made. Many other model classes can serve just as well as a formal basis for what will be discussed in the following sections as the class that is determined by the particular constraints for which we have opted.

2.3 Implicit and Explicit Knowledge. The Limits of Model Theory

Our knowledge of how many days there are in each of the twelve months of the year enables us to identify particular days as the n-th day of month such-and-such in the m-th year before or after the zero point – as in, say, *the eighth of September 1995 (A.D.)*. Let us call such descriptions of days "canonical date descriptions". In each of the models discussed in the last paragraph of the previous section, the denotation of any such canonical date description is completely determined.

We have already drawn attention to the difficulty that speakers often have in correlating canonical date descriptions with descriptions couched in terms of the calendar predicate $week_c$ and/or the weekday names *Sunday, Monday, ...*. Not being able to tell whether, say, the third of next month is a Wednesday is, evidently, a form of ignorance. But what kind of ignorance is this?

It might perhaps seem natural to think of the apparently loose connections between the week-related notions and those figuring in the canonical date descriptions as something to be captured in terms of model variation: For instance, the model class which reflects the knowledge of someone who has no idea of what day of the week is the third of the coming September would contain models in which the seventeenth of September is a Sunday, models in which the seventeenth is a Monday, and so on for the five other days of the week. But model variation is not the right way to account for this kind of ignorance. For mostly it is ignorance that coexists with knowledge which, when properly applied, would make the ignorance disappear. I may be "ignorant" of what day of the week the third of September is in that I cannot answer you promptly when you ask me; but I may at the same time be aware that today is the tenth of August and that today is Sunday. And from those two facts and what I know about the calendar generally I can, with a little effort, derive that the third will be a Wednesday. Model theory cannot capture the difference between my knowing that the tenth of August is a Sunday and my not knowing that the third of September is a Wednesday, because any model which verifies the former fact as well as the general principles by which the calendar predicates are connected will also verify the latter fact. This is a familiar problem: model theory cannot account for the distinction between *explicit* knowledge – where what is known is directly present to consciousness in a propositional form close to that in which the knowledge attribution presents it – and *implicit* knowledge – facts of which one is not directly aware, but which can be derived from that of which one does possess explicit knowledge.

Other instances of unknown correlations, though different in some ways from the one just described, are often similar in this regard. For instance, it is not uncommon for someone not to know what day it is today. But as a rule he will have *some* information that fixes the correlation of $week_c$ and $month_c$ across the board, e.g. because he remembers that he got married on the 22-nd of April 1977 and that that was a Friday. To derive from this that the third of September 1997 is a Wednesday is more of an effort than deriving it from the fact that the tenth of August 1997 is a Sunday. But in either case it is a conclusion that can be drawn from information under his active control, and thus it qualifies as implicit knowledge.

In fact, cases where the speaker has no information at all from which such correlations could be derived are probably quite rare. The failure to correlate weekdays with canonical date descriptions is certainly not of this kind, and it is this type of failure for which we are trying to find an account. Such an account must of necessity involve the concept of a derivation or computation: it must be an account couched in terms of *proof* theory, not model theory. This places it beyond the scope of the present paper, in which we have decided to deal with matters of denotational semantics only.

A further source of ignorance has to do with the denotations of such expressions as *Christmas Eve, All Saints Day, Easter, Lent, Whit Sunday, St Patrick's Day, Tag der deutschen Einheit, Buß- und Bettag , Yom Kippur, or The Queen's Birthday.* When you ask either of us which day of the month was Easter Sunday this year, we will not be able to tell you, nor do we think we could derive the answer reliably from information that is in our possession. In this respect these cases are different from those considered above. To tell when Easter was this year one would have to know the relevant rule (perhaps one knows that this rule has something to do with the phases of the moon, but without quite knowing all the relevant details); and besides knowing the rule, one would also need to know the relevant facts about the moon this year. A person resembling either of this paper's authors won't know either of these things.

Ignorance of this sort does have model-theoretic repercussions. For instance, suppose that all you know about Easter this year is that it was some time in the second half of March. Then the model class which reflects what you know about the denotations of *Easter Sunday* will contain models in which the single member of the extension of this predicate that falls within the year 1997 is the first Sunday after the 15-th of March, models in which it is the second Sunday and models in which it is the third. The class may contain these models even if you know which day was Easter in some other year (so that all models in the class will coincide with regard to that year's date of Easter). For knowing when Easter was then does not help you to determine the time of Easter this year.

Formally, then, ignorance concerning canonical descriptions of the denotations of such terms as *Easter 1997* differ importantly from ignorance about canonical descriptions of the denotations of those NP's which contain the term *(calendar) week* or the names of the weekdays. In the latter case what is lacking is usually just explicit knowledge; in the former it is more often than not implicit as well as explicit knowledge that is missing. In actual practice, however, the difference between these two cases isn't as big as our discussion may have suggested. For even those who have not internalized the information that would enable them to calculate, say, the day in March or April on which Easter will fall in 2002, will often have easy ways of finding this out, e.g. by looking it up in a diary. In real life there isn't all that much difference between gaining explicit knowledge by deriving it from explicit knowledge one already has and gaining it by consulting an external source.

One implication of this discussion is that there can be no unique answer to the question: What is the class of models which captures our general (explicit and implicit) knowledge of the calendar? How good an answer can be given depends among other things on the calendar-related terms in which our knowledge or ignorance manifests itself. If we limit attention to terms like *day*, *week*, *month*, *year* together with the names of the months and the days of the week then the class of models characterised in section 2.3 will come close to reflecting the knowledge of the typical adult speaker of our culture. But as soon as additional terms are included, which do belong to the typical speaker's vocabulary as well, the model classes reflecting the knowledge of different speakers begin to diverge. The best one can do to capture the "common knowledge" underlying these terms is to include all models that are compatible with at least some intuitively competent speaker's knowledge of how these terms denote.

This concludes our informal discussion of temporal common nouns. In the next section we give a formal definition of the class of models which will serve as the basis for our further explorations.

2.4 Models: Formal Definitions

In the introduction we characterized models as abstract representations of possible worlds. This comparatively simple "extensional" model-theoretic conception is all we really need for the analyses of calendar terms in the following sections. However, since these analyses are meant to be interpreted in a more comprehensive semantics in which intensional notions will play their inevitable part, we have opted to follow our current practice of working with a model concept which can accommodate intensional notions as well. Such intensional models are not representations of single possible worlds, but of families of them. With respect to such a model M an expression e gets a denotation for each of the "worlds" represented in M; the function from the "worlds" of M to the corresponding denotations, the intension of e in M, can then be considered as an approximation of the meaning of e. (The approximation is the better, the more possible worlds have representations within M.)

Definition 2.1 A *model* consists of

(i) a set of "possible worlds" W.

(ii) For each $w \in W$ a time structure $\mathcal{T}_w = \langle T_w, <_w, \equiv_w \rangle$, where T_w is a non-empty set (of the moments of time in w), $<_w$ is a dense linear order of T_w, $\equiv_w$ is an equivalence relation on the set $\{ \langle t_b, t_e \rangle : t_b < t_e \in T_w \}$ of periods of $\langle T_w, <_w \rangle$; moreover, $\langle T_w, <_w, \equiv_w \rangle$ satisfies the conditions Size Comparability, Repartition, Halving, Metric Connectedness, and Metric Unboundedness defined in 2.2.

(iii) For each $w \in W$ a universe of discourse U_w, which may contain entities of various sorts. In particular, U_w may include a set of events and it will always be assumed to include T_w as well as the set of periods of the time structure (i.e. the set of all pairs $\langle t_b, t_e \rangle$ with t_b, t_e from T_w and $t_b <_w t_e$).

(iv) An interpretation function I which assigns, for each world w, an extension to the set of lexical elements of the language fragment[6].

Since the intensional dimension of our models is not strictly needed for most of our specific purposes below, we will often make do with "quasi-extensional" models, in which the set W consists of one world only. Where this is so, we suppress reference to worlds completely; thus we dispense with the component $W (= \{w\})$ when specifying a model, and we also omit all subscripts w. Moreover, even where it is assumed that W consists of more than one world, we will make the simplifying assumption that all worlds determine the same time structure. So from now on the second component of our models will always be a triple $\langle T, <, \equiv \rangle$, which describes the one, "absolute" time structure of the model.

2.5 Extensions of Calendar Predicates

The model class we have defined so far is meant to reflect what is general knowledge about time: The properties shared by the time structures of all models in this class are those to which anyone who can be said to have a concept of time is committed. But as we have seen, the typical speaker's knowledge of time will as a rule include a good deal more than that: A large part of our practical command of time is connected with our conception of the calendar and our understanding of the terms which relate to it.

In Section 2.3 it was noted that information about the calendar can be encoded with respect to a temporal structure $\langle T, <, \equiv \rangle$ by fixing the extensions of the calendar predicates day_c, $week_c$, $month_c$, $year_c$ and that for this the determination of a "zero

6 Since in this paper we do not define a particular fragment, we will not say more about the set of lexical items for which I has to provide extensions. In what follows we will be mostly concerned with nouns like those discussed already (day_c, $week_m$, *Christmas eve*, etc.).

point" and of one element in the extension of day_c suffice to fix all the rest. Stating the general principles by which the extension of the different calendar concepts are connected is, we observed, a tedious business, but it is an inescapable part of any computational account of the semantics of calendar terms that speakers have access to. A comprehensive presentation of the principles would be out of the question; but a few illustrations ought to make clear what the principles look like in general, and we will confine ourselves to those. We will state these constraints in the form of meaning postulates.

(CC1) (the extension of $year_c$ consists of abutting non-overlapping intervals)
$(\exists x \supset\!\subset year_c(x) \wedge ((\forall x \supset\!\subset year_c(x) \rightarrow (\exists u \supset\!\subset (\exists v \supset\!\subset year_c(u) \wedge year_c(v) \wedge u \supset\!\subset x \wedge ux \supset\!\subset v)$
$\wedge\ (\forall z \supset\!\subset year_c(z) \wedge z \neq x \rightarrow (z < x \vee x < z))))$

Postulates analogous to CC1 obtain for each of the other calendar concepts.

The next postulate states that each calendar week consists of seven calendar days, and that these are, in that order, a Sunday, a Monday, ..., and a Saturday; and further that each calendar is a day of this type

(CC2)
$(\forall x \supset\!\subset week_c(x) \rightarrow (\exists! u_1 \supset\!\subset \exists! u_2 \supset\!\subset \exists! u_3 \supset\!\subset \exists! u_4)(\exists! u_5 \supset\!\subset \exists! u_6 \supset\!\subset \exists! u_7)(\bigwedge_{i=1,\ldots,7}(day_c(u_i)$
$\wedge u_i \subseteq x) \wedge \bigwedge_{i=1,\ldots,6} u_i \supset\!\subset u_i + 1 \wedge Sunday(u_1) \wedge Monday(u_2) \wedge Tuesday(u_3) \wedge Wednes-$
$day(u_4) \wedge Thursday(u_5) \wedge Friday(u_6) \wedge Saturday(u_7)) \wedge (\forall y \supset\!\subset y \subseteq x \rightarrow \bigvee_{i=1,\ldots,7} u_i \bigcirc y)$

We need similar postulates saying that the calendar year consists of the twelve months January, February, ..., December, and that each of those contains its respective number of days. (We can make things easy for ourselves as regards February in that we specify that it has either 28 or 29 days; most people know more, e.g. that a February of 29 days occurs only every four years; the ultimate subtleties needed to readjust the calendar as closely as possible to the sidereal rhythm are unknown to most of us and thus should probably not be included. In any case, these are matters of detail we leave for others to sort out in whatever way they see fit.)

In addition to these postulates we need the postulates to the effect that a calendar day consists of 24 hours, an hour of 60 minutes and a minute of 60 seconds. Moreover, any two calendar weeks, calendar days, and any two hours, minutes or seconds are of the same duration. We state this last principle for the case of calendar weeks:

(CC3) $(\forall x \supset\!\subset \forall y \supset\!\subset week_c(x) \wedge week_c(y) \rightarrow x \equiv y)$

For the concepts of a calendar week and calendar day we can define the extension of the corresponding measure predicates straightforwardly. E.g. for $week_m$ we get

(CC4) $(\forall x \supset\!\subset week_m(x) \leftrightarrow (\exists y \supset\!\subset week_c(y) \wedge x \equiv y))$

(Recall, however, our remarks in 2.3 on the loose usage that one often makes of this and other measure concepts in such expressions as *a week ago*, etc.)

Less straightforward, as we noted earlier, is the characterization of the concepts of $year_m$ and $month_m$. The problem with $year_m$ is that calendar years are not all of the same length.

The proposal we made towards the end of Section 2.3 for the extension of $year_m$ has the somewhat unfortunate consequence that it cannot be captured by a single first order postulate. The best we can do is to adopt an infinite set of axioms, one for each number n, stating that a period composed of n abutting calendar years and one consisting of n abutting members of the extension of $year_m$ are of about the same length. To make this precise we can say, for instance, that the first period lies between (in the sense of N) the second period minus one day and the second period plus one day. (Other ways of expressing approximate sameness of duration might be considered as well. We leave the explicit formulation of the postulates, in the particular form just suggested or some other, to the reader.)

A much simpler theory is of course obtained by identifying the extension of $year_m$ with the equivalence class of some particular member of the extension of $year_c$. As we have argued, this theory is false, and is known to be false by most speakers. Nevertheless the computation results it yields for the denotations of time denoting expressions containing measure uses of the noun *year* may be correct often enough (and close enough to being correct in other cases) to justify its adoption as an inference basis in certain practical applications.

For $month_m$ we propose, in accordance with what was said in 2.3, a postulate expressing that its duration is between 28 and 31 days. (Again, formalization is left to the reader.)

What about predicates such as *Christmas, Easter, Whitsun* and more such terms for holidays which celebrate or commemorate events of religious or historical significance? Again, these are questions on which we do not feel we should commit ourselves. How much is known about the extensions of these concepts varies, we suggested earlier, widely with cultural background. One very weak assumption which recommends itself as capturing something that many people know about such concepts is what sort of entities belong to the extensions, viz. that they are calendar days, and perhaps also that their extensions contain exactly one instance per calendar year.

In addition to the calendar concepts already discussed there are those which play an important role in ordinary communication, but which are not part of what we see as the canonical vocabulary for designating dates and times of day. These are (i) the terms to subdivide the calendar day into morning, afternoon, evening and night and (ii) the subdivision of the year into winter, spring, summer and autumn. For both these groups of predicates there seem to be something like "official" sharp definitions. In the case of the seasons this is generally known and accepted: spring is from March 21 till June 21, summer from June 21 till September 21, autumn from September 21 till December 21 and during the remaining time it is winter. It is a curious fact, however, that people often use season terms in ways that go against these definitions. For instance, there is a colloquial use of *autumn* or *fall*, which includes all of September (or at least all of its second half), especially when in the country in question, the decisive break in the weather typically comes near the beginning of that month. This throws some doubt on the model-theoretic interpretation of these terms which would otherwise have been the obvious choice: the interpretation in $\mathcal{M}$ of *spring* is the set of all intervals (t_1, t_2) where t_1 is the beginning of the 22-nd of March of some year and t_2 is the end of the 21-st of June of that same year; and so forth. Perhaps one should distinguish here between two different model classes, the first corresponding to the "official" use of season predicates, which consists only of models in which their extensions are defined in the way just indicated, and a larger class, reflecting a more informal use, which also contains models in which the extensions deviate from these official standards. Note well, however, that even in the models in this larger class the extensions of the season terms will have to satisfy a general condition to the effect that each calendar year is partitioned into winter, spring, summer, autumn and winter, in this order.

In addition one might want to add certain further constraints, to the effect that the seasons cannot begin or end just anywhere. For instance, even if a lot of people may be unclear exactly when autumn starts, it seems reasonable to assume that it is part of general knowledge that it does not start on the 1-st of January. But where is the boundary between possible and impossible beginnings of autumn?

The terms *morning*, *afternoon*, *evening* and *night* present similar problems. While there seems agreement on the general principle that the four terms partition any 24-hours period between them, there is much less agreement over where the boundaries lie, and this, moreover, varies a good deal with custom and culture (cf. Italian *sera* with English *evening*). Within a culture like ours, there appears to be a good deal of agreement for some terms (e.g., the afternoon starts at noon and ends at six) while little for others (at what time does morning begin or evening end?). To catch the common denomination of popular assumptions one would have to admit a good deal of variation in the ways that the extension of these terms partition the calendar days. Once more vagueness is a major problem.

2.6 Days-of-the-Week and other generic Objects.

Before we conclude our discussion of the model-theoretic basis for our analysis, there is one further matter we need to address.

Consider a sentence like

(1) Monday is the worst day of the week.

For such a sentence there would seem to be a natural semantic analysis which takes the copula *is* as identity and the phrases *Monday* and *the worst day of the week* as two referring terms (a proper name and a definite description, respectively). In other words, this analysis construes the sentence as asserting that these terms denote the same object. This object is one of the seven "days-of-the-week", an entity which is not privy to any one week in particular, and thus is not a calendar day in the sense discussed so far, but which has its "manifestation" – a particular Monday – in each and every week. Perhaps it might also be possible to account for the intuitive meaning of (1) in some other way, e.g. by assuming that the sentence has a logical form which involves generic quantification over particular Mondays. But it is not at all clear to us what such an analysis should be like in detail, let alone how it could be generalized to the large variety of sentences in which week day names seem to function as names of "days-of-the-week", rather than refer to any particular day.

Within the domain of calendar-related NPs the phenomenon which (1) exemplifies is quite general. Not only are the terms we have for the days of the week often used in the way just indicated; we find the same with expressions referring to times of day, to months, to seasons, to holidays:

(2)

(i) Five o' clock in the afternoon is the time when I am least alert.

(ii) Monday at 3.45 is the time of the Institute Colloquium

(iii) It is useless to talk to him in the morning.

(iv) (The month of) May is the most beautiful month of the year.

(v) I like autumn better than spring.

(vi) Whitsun is a Christian holiday.

(vii) Easter never falls on the same day in two consecutive years.

(viii) Easter has been an official Christian holiday only since the 4-th century.

The list could be extended indefinitely. But even as it stands, it seems to us to provide very strong evidence for the thesis that a natural semantics for calendar-related expressions should acknowledge such objects as days-of-the-week, months-of-the-year, times-of-the-day, and so forth. So we want to include such objects in the universes of our models.

But how are such entities as days-of-the-week related to the objects which the universes of our models have already been stipulated to contain? Clearly there are systematic relations: as we noted, the day-of-the-week Monday has its "manifestation", a member of the extension of the predicate *Monday*, in every single week; and so it is with all other entities of the sort that is now under discussion. In fact, in the light of what we have said so far we can conclude that the word *Monday* has not two related uses, as was implied in previous sections, but three – (i) as bona fide count noun, as in *every Monday* or *on most Mondays*; (ii) as designator of some particular day, as in *I'll ring on Monday.* or *He filed the petition on Monday, the third of October, 1995*; and (iii) as designator of the day-of-the-week, as in (1).

In English it is unusual for a single word to play this syntactically and semantically multiple role; but the semantic relationships which this multiplicity brings into focus are much more general. For any common noun N we can form a corresponding definite description *the N.* The standard analysis of definite descriptions has it that on any particular occasion of use *the N* refers to the unique object in the given context which belongs to the extension of N. For the majority of uses of definite descriptions this is essentially correct.

But descriptions have other, "generic" uses as well, and these are related to (i) their "standard" uses and (ii) the extension of the common noun phrases N which constitute their descriptive contents in much the way in which the use of *Monday* in

(3) is related to those uses of *Monday* as NP in which it refers to a particular day and to the extension of *Monday* as common noun. This parallel is particularly close in those cases where the generic interpretation has a clearly visible functional character, as we find it with descriptive phrases such as the President of the US. Consider (3).

(3) The President of the US has always been an American citizen.

On the reading that concerns us here the truth conditions of this sentence involve different presidents at different times. One way to account for this is to assume that the referent of the NP *the President of the US* which gives rise to this reading is a function which yields, for each of the relevant times (those lying within the period during which the US have been in existence as a political unit), the corresponding individual who held the office of President at that time.

Often with such functional interpretations the function arguments are times, but this isn't always so. For instance, in a sentence like *In rural New England the town hall is (usually) right next to the church.* the arguments are New England towns; and examples with other kinds of arguments can be found as well. But such cases need not concern us here. On the other hand, in the functional interpretations – of *Monday* etc. – which prompted this excursus not only the function arguments are times but the function values are as well, and it is this kind of case – of functions from times to times – on which will focus here exclusively.

There is no general agreement on whether interpretations like the one of (3) we have just discussed are best analysed by assuming that they involve a special denotation for the relevant NP (in this case the President of the US), which then yields the various function values through application to the values of some variable introduced by some other aspect of the interpretation. (In (3) this would be the variable bound by the temporal quantifier *always*.) However, for the special case with which we are confronted here, viz. that of weekday names in their role as designators of days-of-the-week (as we have informally called them), this line of analysis strikes us as intuitively plausible. That is, we want to assume that days-of-the-week are genuine objects and that it is to such an entity that, e.g., Monday in (1) refers (on the reading of that sentence which we have been discussing) Similarly we want to assume similar entities corresponding to the analogous uses of the NPs in (2): seasons-of-the-year, Christmas as an annually recurring event, and so forth.

This decision implies that we have to include such entities within the universes of our models. But if we take such entities on board, exactly what kind of entities are they?

We have observed for some of them that they determine functions from restricted periods to instances within those periods, and we have informally used the suggestive term "functional entity". But are they just functions, or are they entities that determine functions, but whose identity itself is determined by other, and possibly stricter, criteria? We do not know. Arguably such notions have a conceptual dimension that transcends the strictly extensional properties of a set-theoretic function. For the purposes of this paper this conceptual dimension is of no direct relevance. But in any case our intensional model theory allows us to capture at least some of it: We can take the day-of-the-week Monday to be the corresponding *intension*, the function which for each world w and each week t in w returns the Monday of t in the world w. So we adopt this function and its six obvious counterparts as the seven days-of-the-week. Other "functional" entities that semantics may need as referents for temporal NPs will be identified with similar "intensional" functions.

Strictly speaking we should, at this point, revise our model theory once more and stipulate that the universes of our models contain the right functions. As, however, the new entities are in the set-theoretical closure of the models as we have defined them, the revision involves no more than formally adding to the universe things that were there implicitly in any case. Little would be gained by carrying this out in formal detail.

3 Semantics and Pragmatics of temporal NPs

3.1 Designators of calendar Times

Many expressions that we use to refer to times depend for their reference on the contexts in which they are used. In fact, this seems to be the case for a clear majority of those time-denoting expressions that we most commonly use in practice – we mentioned examples of such expressions already in the introduction – expressions such as: *the sixth of June, the third, at a quarter past five, on Wednesday, Sunday morning.* The frequent use we make of such expressions is in part connected with the circumstance that in daily life the times we have occasion to refer to are quite often not too distant from the time at which the referring is being done: Many belong to the same day; and if they do not belong to the same day, then they often belong to the same week, or to the same month; and when that isn't the case, they often belong to the same year. In such cases there is no need to use an "absolutely" referring expression such as *at 10.15* or *on the sixth of June 1998.* If it is clear from the context what day is in question, then *at 10.15* will do just as well and will be preferred for reasons of conciseness; and if the day isn't contextually determined, but the year is, we could and normally would say *at 10.15, on the sixth of June*, leaving out the redundant *1998.*

Expressions such as *at 10.15, the sixth of June* etc., which yield a unique denotation only when the context provides a restricted temporal domain within which the descriptive content of the expression has a unique instance, seem to behave in a kind of indexical way: The expression refers to the unique instance of its descriptive content within some limited period around the utterance time. But in other cases the periods within which the expression determines its denotation via unique instantiation do not contain the utterance time; in fact, they may be arbitrarily far removed from it. What all these cases seem to have in common, however, is that the period within which the descriptive content is instantiated uniquely is itself an element of the extension of some other calendar predicate. This predicate does not occur within the expression used, but stands to it in some systematic relation. Roughly speaking it is "one up" from the most general calendar predicate that occurs within the expression used. Thus, if the expression is something like *June* or *the sixth of June*, or *the sixth* *of June, at ten past seven*, then the predicate will be $year_c$; when the expression is *the sixth*, the relevant predicate will be $month_c$; when it is *Wednesday*, the predicate will be $week_c$; when it is *at ten past seven*, the predicate is day_c. To make this precise we need to introduce a notion that has hovered in the background of much that we have been saying so far and that now needs to be brought out into the open and made explicit. We will refer to this notion as the *Calendar Granularity Hierarchy* (CGH).

The term should give a fair idea of what we intend by it. We have stressed repeatedly in the preceding sections that an important feature of our conceptualization of time is the way in which we see it as subdivided into years, seasons, months, weeks, days, ... Each of these calendar concepts imposes on the continuum of historical time a partition of a certain *granularity* – a partition the elements of which have a certain size; and the size varies considerably from one partition to the next. These partitions make it possible for us to think of time – in situations in which we are contemplating certain processes or sequences of events – as quasi-discrete, and the multiplicity of partitions allows us to choose the size of the discrete portions so that it suits the subject of our attention. The calendar concepts supply us with a set of ready-made options for this purpose, and for most ordinary purposes the range of options it offers is adequate.

Obviously, there are systematic relations between many of these options. Some of them stand to each other in the relation of refinement, in the formal sense in which one partition can be a refinement of another. In addition there are cases of "quasi-refinement", where the portions of one partition stand in a more or less fixed proportion to those of the other – as in the case of weeks to months or months to seasons – but where in general an element of the coarser partition is not a perfect sum of a number of elements of the finer partition. For instance, though the length of a month is that of between four and four and a half weeks, it is exceptional for a month to be identical with the period covered by an integral number of calendar weeks. (In fact, this is only the case when the first of February happens to be a Sunday and the year is not a leap year.) Nevertheless, it is part of our way of thinking of time that we see its subdivision into months as "rougher" than the subdivision into weeks, since a month always lasts between four and five times as long as a week. A reflection of this fact is the frequent use of expressions like *the third week in November*.

The graph of the Calendar Granularity Hierarchy is given in (3.1):

Definition 3.1 (Calendar Granularity, Hierarchy)

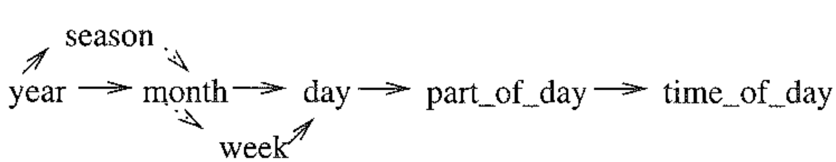

The straight edges in this diagram indicate strict refinement (with the more refined partition standing on the right). The dotted edges express the weaker relation of quasi-refinement.

We will refer to the nodes of the CGH sometimes as sorts and sometimes as (sortal) predicates; we think of the sort as the intension of the corresponding predicate. For each of the temporal predicates P with which this section is concerned (*February,* *Wednesday, half past ten, autumn, morning,* etc.) there is a sort s in the CGH which is the sort of P in the sense that the extension of P is always a subset of the extension of s. Similarly we can also speak of the sort of a temporal designator (such as *the first of May, May* (as NP), etc.). The sort of a designator is that sort of the CGH whose extension contains the possible denotations of the designator as elements.

The CGH is useful not only as a compact representation of the standard temporal granularity options, but also as an instrument in describing the reference conditions of the temporal expressions we are discussing. However, in order to make an effective use we first need to say a little more about the structure of these expressions. As the examples we have mentioned illustrate, the expressions under consideration are strings of temporal designators, and each of these designators can be assigned an element of the CGH as its sort. For instance, *the sixth of June, 1998* can be analyzed as a string of three designators, *the sixth, June* and *1998*. The first of these is of the sort *day*, the second of the sort *month* and the third of sort *year*. Similarly, *the sixth of June* involves the first two of these designators and *the sixth* only the first.

For expressions δ which consist of one or more such sorted designators we can define the *relevant predicate* as follows: Let s be the highest sort (i.e. the leftmost sort in the sense of CGH) of δ. Then the *relevant predicate* for δ is some "suitable" predecessor of s in the CGH. Exactly which predecessor is a question to which we do not wish to give a full answer right away and for the time being we will just give a few examples for the case that δ has the semantic status of a "definite description" – that is, δ is a singular NP which either begins with the definite article *the* or else is a bare NP like *July* or *Monday*. In these cases the relevant predicate is that predecessor s' of s in the CGH which stands to δ in the following relation UI(s', δ)

Definition 3.2 Let δ be a calendar time designator, s' a sort from the CGH. Then we have UI(s',δ) iff each member t in the extension of s' includes a unique time which satisfies the descriptive content of δ.

Some examples[7]: (i) NPs that consist just of the name of a day of the week, such as *Thursday*. Here the sort s of δ is day_c and the relevant predicate (i.e. the sort s') is $week_c$. (ii) NPs consisting of *the* and an ordinal, such as e.g. *the third*.[8] Again the

[7] Some of these were already discussed informally in the second paragraph of this section.

[8] With NPs consisting of *the* and an ordinal we encounter one of the idiosyncrasies of the English systemm for referring to time. In general, English admits NPs the nominal part of which consists exclusively of an ordinal, or else of a word such as *last* or *next*. There exists a general strategy for interpreting such NPs which treats them as "incomplete" and relies on the context to supply a partially ordered set, from which the adjective occurring in the NP can then select a subset as denotation of the NP's nominal part. Often such interpretations take the form of recovering from the context a common noun or common noun phrase – which may be thought of as the "missing common noun (phrase)" of the "elliptical" NP – provided the (contextually restricted) extension of this common noun or noun phrase can be seen as coming with some appropriate ordering to which the overt adjective of the NP can be applied . Thus, in *I quite enjoyed the first talk. But by the third I was solidly asleep.* the natural interpretation of *the third* is as *the third talk*, the common noun phrase *talk* having been recovered from the first sentence and the order being the temporal order in which the talks were given. (We note in passing that some

sort s is day_c, but the sort s' is now $month_c$. (iii) NPs consisting of *the*, an ordinal, the preposition *of* and the name of a month, as in *the third of July*.[9] Once more s is day_c; but this time s' is $year_c$. (iv) Compound NPs like *Wednesday, the sixth*. Once more s is day_c. But as regards s' the matter is now more involved. It is a general property of such compound NPs, which consist of two full NPs (separated by a comma or just put in sequence) that they are to be treated as pleonastic: the referent of the compound NP must be both the referent of its first component and the referent of its second component. As a consequence we now have two possible options for the relevant sortal predicate s': the first component of the NP licenses $week_c$ and the second component $month_c$. (v) *five minutes to midnight*. In this case s is the sort *time-of-day* and s' is the sort day_c.

It will be convenient to have a cover term for the expressions we have just informally described. As we are unaware of a generally agreed term in the literature on temporal reference, we will cast one of our own making. The best we have been able to come up with is *calendar time designator*. We subdivide this category into the *absolute calendar time designators*, for which the highest sort is $year_c$, and *elliptical calendar time designators*, for which the highest sort is an element of CGH that lies to the right of $year_c$.

The "indexical" uses of elliptical calendar time designators can now be characterised as follows:

Definition 3.3 Let δ be an elliptical calendar time designator, and let $U(s', \delta)$. An *indexical use* of δ is a use in which its denotation is the unique instance of the descriptive content of δ within that unique member of the extension of P which also contains the utterance time.

This definition captures the non-compound cases of elliptical calendar time designa-

languages, such as German, are more liberal with regard to such "elliptical" NPs than English, in that they allow for such elliptical NPs with arbitrary Adjective Phrases.) Besides this general use of NPs of the form "article + ordinal" there is in addition the interpretation that concerns us here. This is a conventionalized interpretation, in which the missing common noun phrase is *day of month M*; it is conventionalized in that there is no need for the context to make a predicate of this form explicitly available. It is worth observing that this conventional interpretation is distinct from the general strategy, and that the two may produce distinct results, even in cases where the referent determined by the latter strategy also is a day. For instance, in the two-sentence discourse *The meeting took five days. By the fourth everyone was totally exhausted.*, *the fourth* presumably refers to the fourth day of the meeting and this day need not have been the fourth of the month. In fact however, in a context such as this the phrase *the fourth* is genuinely ambiguous. If for instance the third day of the meeting was the fourth of March, then *the fourth* could also be taken to refer to the fourth of the month and thus to the third day of the meeting. The possibility of interpreting NPs of the form "article + ordinal" in the special sense of "day of month M" needs to be stated separately somewhere in the grammar, probably as part of the lexical information associated with ordinals.

[9] NPs of this form are subject to the same convention that was noted in the previous footnote: *the third of July* must be interpreted as *the third day of the month of July*. Here, moreover, there appears to be no room for ambiguity; the conventional interpretation is the only one such phrases allow.

tors on which we have so far concentrated. Here are a few more examples of such cases. First, suppose that you want to refer to the time 16.15 of the day on which you are speaking. Then the phrase *at 16.15* will do: The relevant sort is day_c, the unique time in its extension which contains the utterance time is the day on which the utterance takes place and the unique time within that time which belongs to the extension of the predicate *16.15* is the time of that description on the day of utterance. Similarly, if it is the Wednesday of a given week w and you want to refer to the Friday of that weeks you can do that by using the designator *Friday*. The relevant sort is then $week_c$, the unique member of the extension of that sort containing the utterance time is w and the unique time within that member which satisfies the descriptive content of δ is the Friday of that week. (Implicit in this analysis is that the bare NP *Friday* can be analyzed as having the predicate *Friday* as a constituent, just like NPs *the Friday, a Friday* or *every Friday*.) For a third example, suppose that it is the 10-th of the month and you want to refer to some other day of that month, for instance the fifteenth. Then all you need to say to establish that reference is *the fifteenth*. The relevant sort is then $month_c$, and the referent, the unique day satisfying the description *fifteenth (day of $\mathcal{M}$)* if we assign to $\mathcal{M}$ the month that includes the utterance time, is the day to which you meant to refer.

Besides the time designators exemplified in the last paragraph, in which there is only one calendar predicate, there are those which involve several such predicates, as in *Monday, at five minutes to midnight; the third of July, 1995* or *Wednesday morning at half past ten*. These expressions pose some additional problems. To begin with there is a syntactic question: Should they be seen as single syntactic constituents or as accidental groups of constituents, which from a syntactic point of view have no better claim to being analyzed as a single unit than the adjuncts *at Trafalgar Square* and *Sunday morning* in the sentence *The march will take place Sunday morning at Trafalgar Square?* This is not the sort of question that is of interest to us here in its own right, and we do not wish to go into a lengthy discussion about it. (According to our intuitions the answer to the question may be different for different cases.) However, the semantics of the cited expressions requires, as we will see presently, that all their different parts be considered in conjunction, whether or not they should be seen as forming a single syntactic whole. What we will need, in particular, is a function which assigns to each such expression (or group of expressions) the set of all time-designating constituents it contains. We will refer to this function as the *Granularity Spectrum* of the expression or expression group in question. In the absence of an actual syntax we are not in a position to define this function explicitly; but we will give, by way of example, the Granularity Spectrum for the three expressions just mentioned; that should give a fairly clear impression of how the function could be characterized in general.[10]

For the first two of the expressions mentioned in the last paragraph, *Monday, at five*

[10] We trust the reader will be prepared to agree that a formal definition should be unproblematic once an explicit syntax is in place. Note, however, that in cases where the syntactic analysis would not treat the relevant part of the clause as a single constituent, we will have to assume as preparatory to or part of the computation of the Granularity Spectrum for that clause some process which collects the different relevant constituents together.

minutes to midnight and *the third of July, 1995*, the Granularity Spectrum consists of two elements. For *Monday, at five minutes to midnight* it is the set { *Monday, a quarter to midnight* }; for *the third of July, 1995* it is { *the third of July, 1995* }. What the Granularity Spectrum is for the third example, *Wednesday morning at half past ten*, depends on whether we treat *Wednesday morning* as a single NP with a descriptive content that is true of every period which is the morning of some Wednesday, or as a combination of two designators, the first with a descriptive content satisfied by all Wednesdays, and the second with a descriptive content satisfied by all mornings. On the first assumption the Granularity Spectrum is the two-element set { *Wednesday morning, half past ten* }, on the second it is the three-element set { *Wednesday, morning, half past ten* }.

The point of the Granularity Spectrum of an expression (or group of expressions) δ is that the sorts of its members determine a substructure of the CGH. For many δ the correspondence between the members of GS(δ) and their sorts is one-to-one. This is so in particular for all the examples we have considered so far. But this condition isn't always satisfied. One notable exception are complex designators such as *Monday, the third (of July)*, in which both *Monday* and *the third (of July)* are of the sort *day*. In the interest of perspicuity we will ignore such cases in the formulation which we give in (3.4) below of the interpretation principle for elliptical calendar time designators and limit ourselves to the cases for which there is such a one-to-one correspondence. In addition the formulation in (3.4) assumes that among the sorts of the members of GS(δ) there is always one that is more general (that is: farther to the left in the sense of (3.1)) than all the others. As far as we can see, there are no exceptions to this constraint. When the first assumption is dropped, the interpretation principle is a good deal more difficult to state. We have relegated the more complicated version of the principle to a footnote.

Definition 3.4 Let $s_1, \ldots, s_m$ be the sorts (in the sense of CGH) of the members $\delta_1, \ldots, \delta_m$ of GS(δ) and suppose that s_m is the unique most general member of $s_1, \ldots, s_m$. Then for each designator δ_i in GS(δ) with $i < m$ there must be a sort s_j among $s_1, \ldots, s_m$ such that s_j and the descriptive content of δ_i stand in the relation $\lambda s'.\lambda\delta.\ \mathrm{UI}(s', \delta)$ of (3.2); the referent of δ_i is then the unique element of δ_i's descriptive content within the referent of δ_j. The designator δ_m is interpreted contextually in the manner described in (3.3) for non-compound designators.[11]

[11] The general statement of definition (3.4) must take account of the possibility that GS(δ) may contain different members with the same sort. In such cases sameness of sort partitions GS(δ) into groups. (In the cases familiar to us the groups never have more than two members, so our formulation of the problem has a flavour of spurious generality. Still this seems to be the simplest and cleanest way of describing the issue.). In order that such designators denote properly all members of any given group must pick out the same instance. (If this condition is not satisfied, then we get failure of denotation, and thus some kind of failure of presupposition.) Thus consider the expression *Monday, the third*. Here we have a single group, with the sort *day* and the expressions *Monday* and *the third* as members. Suppose the designator is used on some day late in the second week of May 1998, a month in which the third was a Sunday rather than a Monday. The mechanisms we described in (3.4) will select the third of May 1998 as referent for the token of *the third*,

According to (3.4), each member of GS(δ) whose sort is not maximal within GS(δ) must be interpreted with respect to the "local context" provided by the relevant "superordinate" designator in GS(δ). Only the sortally most general member gets a context-dependent interpretation. (Unless its sort is $year_c$ and the expression specifies a particular year in a context-independent way, as e.g. the designator *1995*.).

The definition of the function GS can of course be applied straightforwardly to the simple calendar time designators discussed earlier. For such an expression δ GS(δ) will be a singleton, its sort will be trivially the unique most general one within its set and interpretation will proceed as described above. It is not difficult, moreover, to verify that the procedure sketched in (3.4) delivers, for the cases we have been discussing, the correct interpretation also for expressions for which $\mid GS(\delta) \mid > 1$. Suppose for instance that you want to refer to 11.45 on the Monday of the week in which you are contemplating an utterance. Then you may use the expression *Monday, at a quarter to midnight*. Its constituent with most general sort is *Monday*, which will, according to (3.3) be taken to refer to the Monday of the week in which the utterance is made, and the other expression in the Granularity Spectrum, *at a quarter to midnight*, whose sort is *time-of-day*, will refer to a quarter before midnight of that Monday.

We can, if we want, extend the analysis based on the Granularity Spectrum also to expressions like *Monday, the third* and *Monday, the third of July*. For such expressions the Granularity Spectrum will not contain a unique designator of most general sort. In order that the interpretation mechanism returns the right result for such expressions as well, we must add the clause that where there is no unique designator of most general sort, the interpretation proceeds in the way described in the last footnote.

---

which this utterance contains. But as this day does not satisfy the predicate *Monday*, it fails to qualify as referent of the first component of our expression and the complex *Monday, the third* fails to have a proper denotation.
After this preamble the frills which (3.5) adds to 3.4 should be fairly easy to understand. For the sake of a more compact formulation, we use $\equiv_s$ to denote the relation which holds between two members of GS(δ) iff they correspond to the same sort and $[\equiv_s]_{s'}$. to denote the group of expressions corresponding to sort s'. Since the members of $[\equiv_s]_{s'}$ are required to be coreferential, it is legitimate to speak of the "referent of $[\equiv_s]_{s'}$".

Definition 3.5 Let $s_1, \ldots, s_n$ be the sorts of the members $\delta_1, \ldots, \delta_m$ of GS(δ) and suppose that s_n is the unique most general member of $\{ s_1, \ldots, s_n \}$. Then in each group $[\equiv_s]_{s_i}$ in GS(δ) (where $i < n$) there must be a member δ_i and a sort s_j among $s_1, \ldots, s_m$ such that

(i) s_j and the descriptive content of δ_i stand in the relation $\lambda_{s'}.\lambda\delta.$ UI(s', δ) of (3.2); and

(ii) the unique satisfier t_i of δ_i within the referent of $[\equiv_s]_{s_j}$ satisfies the descriptive content of each of the members of $[\equiv_s]_{s_i}$.

When (i) and (ii) are satisfied, then t_i is the referent of $[\equiv_s]_{s_i}$. In all other cases the referent of is undefined.
The referent of the group $[\equiv_s]_{s_n}$ is interpreted contextually in the manner indicated for δ_m in (3.4).

3.2 Anaphoric Uses of Elliptical Calendar Time Designators.

The interpretation strategy we laid out in the last section is all right in relation to those contexts where the speaker intends to refer to a time which shares the relevant conventional calendar period with the utterance time. But this doesn't mean that the strategy correctly captures the semantics of these expressions in general. For in fact they have many uses besides those on which we have concentrated up to now. Take the bare NP *Wednesday.* When on Monday you say:

(4) Bill will come on Wednesday.

then it is likely that you mean the next Wednesday and that is the Wednesday of the week in which your utterance occurs. But if instead you say, again on Monday,

(5) Bill arrived on Wednesday.

then you are likely to be understood to have referred to the Wednesday of the preceding week; you couldn't have referred to the Wednesday of the week in which you are, for that Wednesday is still in the future of the time at which your utterance occurs, and the past tense of your utterance indicates that you are talking about the period before that time. Indeed, it takes but little experimentation to see that the principle we stated in (3.2) is far too restrictive. The Wednesday referred to by an utterance of *Wednesday* can be just as easily the Wednesday of the week immediately preceding or the Wednesday of the week immediately following the one in which the utterance is made as the Wednesday of that week.[12] The principle that underlies all the cases we have so far considered is rather something like (3.6).

Definition 3.6 Find some interval containing the time of utterance in which there is a unique satisfier of the descriptive content of the constituent δ_m of GS(δ) (and which may have to satisfy additional conditions associated with the utterance that is being interpreted, typically those connected with the tense of the clause of which δ is part). This unique satisfier is then the referent of δ_m. The referent of δ is then determined on the basis of this referent of δ_m in the way stated in (3.3).

But (3.6) is not general enough either. For besides the "utterance-time-oriented" interpretations there are also those in which the reference determination process has been shifted away from the utterance time to some other time. For instance, analogously to the examples (4) and (5) we find cases like those in (6) and (7)

(6) On Friday, the third of July Fred delivered his lecture to the Academy. He left town again on Monday.

[12] In fact, there is one situation in which there is a special tendency to take the Wednesday referred to be not the Wednesday belonging to the week of the utterance. This is when the day on which the utterance takes place is itself a Wednesday. Weekday names are allocentric in much the same way in which third person NPs are allocentric: just as it is unnatural to use a third person NP to refer to oneself (instead of using the pronoun *I*) , so it is unnatural to use *Wednesday* on a Wednesday to refer to that Wednesday (in stead of using *today*).

(7) On Friday, the third of July Fred delivered his lecture to the Academy. He had arrived on Monday.

Let us, for the sake of definiteness, assume that both (6) and (7) are uttered on the first of September. (6) has an interpretation on which *Monday* refers to the first Monday after the third of July and (7) one on which *Monday* refers to the last Monday before that date (and it seems to us that these are the preferred interpretations, even though (6) seems to allow also one according to which *Monday* refers to the last Monday before the time of utterance). In fact, these and other discourse examples clearly show that in principle any time that is contextually salient can play the same role in the determination of an expression of the kind under discussion that the utterance time plays in (3.6).

To see how we should formulate the general principles that govern the reference of the elliptical calendar time designators, we must consider a few more cases. A relevant example is (8)

(8) 1994 was a pretty bad year for us. In February/ (On) the fifth of February Bill had his car crash. Then, on the seventeenth/ on the seventeenth of March Mary was made redundant.

The two designators of the second sentence of (8) demonstrate a variant which we have not yet encountered. Here the context (i.e. the first sentence) provides a period of the sort that is appropriate to each of the two alternatives given in the second sentence – appropriate since within each year there is exactly one February and one fifth of February; and indeed, the times to which these designators refer in the context are precisely these unique satisfiers within 1994.

The two alternatives given in the third sentence of (8) provide evidence of two further variants. *the seventeenth of March* is naturally interpreted as referring to the seventeenth of March 1994. In this case there are two ways to arrive at this result. We can either compute the reference in the same way as that of *the fifth of February* in the second sentence: The contextually salient year 1994 contains exactly one seventeenth of March and that is the time to which the phrase refers here. The second possibility is to take a time introduced by the second sentence – this may be either the month of February of 1994 or else the fifth of that month – and to chose a period including that time which includes exactly one satisfier of the predicate *seventeenth of March*. As usual there are two solutions to this problem, one catching the immediately preceding satisfier and one capturing the one that immediately follows. The choice between them is, here as elsewhere, a matter of tense, rhetorical structure and or further factors (such as, here, the presence of the word *then*).[13]

[13] The presence of this word, together with the circumstance that both the second and the third sentence are in the simple past, indicate (for reasons on which we will not expand here) that the event mentioned in the third sentence follows that of the second. So of the two options just mentioned it is the second that is to be chosen here. Again, this gives us as referent the seventeenth of March 1994.

Note that we have a kind of overdetermination in this case: two alternative strategies lead to the same result. Indeed, we could see the two as acting in consort, along the following

When the designator of the third sentence is *the seventeenth* rather than *the seventeenth of March*, the situation is different. In this case it becomes crucial to distinguish between the discourse in which the designator of the second sentence is *the fifth of February* and that where it is *February*. First, suppose it was *the fifth of February*. In this case our only option for interpreting the *seventeenth* is to find a period containing the time t which the second sentence has introduced (i.e. the fifth of February 1994) and which contains a unique satisfier of the predicate *seventeenth day of month* $\mathcal{M}$ (for some suitable choice of $\mathcal{M}$). For the reasons indicated above this satisfier should lie in the future of t. In this way the referent we get for *the seventeenth* is unequivocally the seventeenth of February 1994. If the designator in the second sentence is simply *February*, then we can take its referent, i.e. the month of February 1994, as instance of the relevant predicate for the designator *the seventeenth*, that is, as an instance of the predicate $month_c$. The referent of the seventeenth then is the unique satisfier within this month. So once again we obtain the seventeenth of February 1994, as we should[14]

In order to formulate what we will present as our final statement of the interpretation principle for the expressions that have preoccupied us in these last two sections, it is helpful to introduce one additional notion. The examples of this section all involved

---

lines. The time introduced by the second sentence can be used as starting point for computation of the reference in the manner described by (3.6) together with the complementary principles described in the last section; but while applying this computation strategy we must also do justice to the rhetorical structure of the discourse we are interpreting, and this structure clearly involves in the present case seeing the third sentence, like the second, as standing to the first sentence in the relation usually referred to as *elaboration*. That is, the third sentence is taken to provide more information about the "topic" introduced in the first sentence, viz. what happened in the year 1994. Thus locating the referent of *the seventeenth of March* within that year is an additional constraint that the computation of that referent ought to satisfy.

14 We could have analyzed this last case so that its analysis more closely resembles that of the preceding one. That is, we could have said: extend the period which we take as starting point, i.e. the month of February 1994, so that the extension contains a unique satisfier of the predicate *seventeenth day of month* $\mathcal{M}$. Since the time with which we start in this case already contains such a satisfier, it is itself already an extension of the kind we want; and no extension containing an additional satisfier – such as, say, the seventeenth of March 1994, would do. In this way we can unify this last case, as well as the cases presented by the designators of the second sentence, with the strategy that was already in place before we started with the discussion of (8).
Note the following complication for an account which includes this last strategy among the different interpretation strategies for elliptical calendar time designators. If the designator of the second sentence of (8) had been the 25-th of February instead of the 5-th of February, then the strategy just described would give us as referent for the 17-th in the third sentence the 17-th of March. To us this interpretation seems marginal – to our ears it would seem much more natural in this case to use the phrase the 17-th of March, as a clear indication that one is now talking about a different month. While an interpretation of the 17-th as the 17-th of May in this variant of (8) doesn't seem impossible, it appears to us to involve some kind of coercion. If this is right, it may well be necessary to restrict application of this particular interpretation strategy, or even eliminate it entirely. We leave this as one of many questions for further discussion.

some time or times distinct from the time of utterance to serve as "starting times" for the process of referent determination. All but one of the examples were all very simple in that the starting times were introduced by the immediately preceding sentence. Only in the last example we saw that this time can also come from a sentence before the last one. That example also raised the question which of the times that had been previously introduced into the context should be chosen as "starting time". And this is the general situation: The context in which elliptical calendar time designators are to be evaluated for reference may contain any number of earlier introduced times and thus pose a serious selection problem. We find an obvious parallel here with the problem often referred to as "pronoun resolution": which among a number of possible candidates is the intended anaphoric antecedent of a given pronoun? In fact, the analogy seems to be more than superficial: In either case the choice is governed by a complex set of factors, some of which, such as those having to do with what counts as the "topic" of the discourse at the point where the choice is to be made, are still incompletely understood and certainly aren't fully operational from a formal point of view. Starting time selection is therefore a problem for which we have no substantive solution; and rather than trying to say something about it here we prefer to address it on some other, later, occasion. Here we make do with the simplifying assumption that the choice is provided by oracle.

For the present purpose of stating our interpretation principle we will make do with an extremely simple notion of context[15], which consists of a single time t_c. Often t_c will be the time of utterance, but as we have seen, it can also be some other time, and in that case it can belong to any one of the different sorts represented in the CGH. In such cases it is the sort of t_c which determines whether interpreting a new temporal expression with t_c as starting time is possible at all. Sometimes it isn't. For instance, if t_c is of the sort $year_c$ and the new designator is, say, *Wednesday*, then there is no way of expanding t_c to give a period in which there is a unique satisfier of the descriptive content of the designator, for it itself contains many satisfiers already. Consequently referent determination fails.

At last we are ready to state our general interpretation principle for elliptical calendar time designators:

Definition 3.7 Suppose δ is a temporal expression (of the kind under discussion) uttered in a context t_c. As before, let δ_m be the member of GS(δ) of uniquely determined most general sort s_n. Then the referent of δ_m will be the unique satisfier within t_c or some extension of t_c; and the referent of δ will be determined on the basis of that of δ_m in the way described in (3.4) of Section 3.1.

Henceforth we will refer to the time t_c which serves as input to the process of reference determination δ that (3.7) describes, as the *origin of computation (of the referent of δ).*

Inspection of (3.7) and (3.4) shows that reference determination can fail for two different reasons – either because of a mismatch between δ_m and t_c or because of a similar mismatch involving some other member δ_i of GS(δ), for which no suitable partner δ_j in the sense described in (3.4) can be found.

---

[15] For more on the notion of context see Section 5.

As we saw, the procedure outlined in (3.7) is often ambiguous in that two different extensions of t_c will contain distinct unique satisfiers of the descriptive content of δ_m. In those cases, however – we have noted this already but make the point once more – ancillary constraints, especially those connected with tense, usually force a choice between these two possibilities.

3.3 Elliptical Calendar Time Designators and Indexicals.

One of the most important insights of the theory of reference is that natural languages are equipped with certain so-called *directly referential expressions*, expressions which contribute their referents directly to the truth conditions of the sentential utterances of which they are part, thus making these utterances into predications of some property P of the referents (rather than general propositions which say that an object with such and such properties satisfies P as well). According to the received views within the philosophy of language there are two main types of such expressions, on the one hand the proper names and on the other the indexicals and demonstratives. For the former the locus classicus is Kripke 1970, for the latter it is Kaplan 1971, which finally appeared in print as one of the contributions to Almog 1989. While proper names, demonstratives and indexicals have in common that they are all directly referential, there is another respect in which names and the other two categories occupy opposite extremes. Proper names are the device *par excellence* to make reference – and therewith the truth conditions of the sentences in which proper names occur – independent of context. Once a name has been assigned to an object, it can be used to refer to that object in more or less any context whatever. Demonstratives and indexicals on the other hand are referential devices of an essentially ephemeric nature; what such an expression refers to always depends on the context in which it is used and thus varies from one context to the next. But demonstratives and indexicals differ from each other in the way in which their reference is contextually fixed. In the case of demonstratives this is some kind of demonstrative act, a *demonstration*, which serves to pick out, from within the setting in which the utterance takes place, some particular object as the demonstrative's referent. The user of the demonstrative can, by opting for one demonstration rather than another, steer the demonstrative expression in the direction of this referent or that; so even within a given context he has a certain freedom to make the demonstrative refer to one of several things. Indexicals are not like that. Their reference is, in any given context fixed by linguistic rule.

It is the indexicals, and more specifically the temporal indexicals, which are of interest to us here. It is part of the classical theory of indexicals that English and similar languages contain a number of these; in English we have for instance: *now*, *yesterday*, *tonight*, *last week*, *next month* and quite a few more. For *now* the classical reference rule is particularly simple: in any context of use *now* refers to the time of utterance, n_c. The other temporal indexicals differ in that they refer not to the utterance time n_c itself but to some time that stands in some systematic relation to it (in general one that entails temporal proximity). Thus *yesterday* refers in an utterance context c to the day immediately preceding the day which includes the utterance time n_c, *tonight* to the evening of the day including n_c, etc.

Since the theory of indexical reference was first formulated it has become general wisdom that temporal indexicals do not behave in full accord with what the theory seems to predict. For instance, *now* can, when used in the right setting, be interpreted as referring to times lying in the past of the utterance time. An example is the following:

(9) On the fifteenth of May Bill came home at a quarter to eight. He got out of his shoes, threw off his suit, put on a dressing gown and slumped in front of the TV with a glass of whiskey. The day had been one uninterrupted hassle, with a thousand different things to think of. But **now** at last he could relax and enjoy what was left of the evening.

In such passages *now* can be used perfectly naturally to refer to the past time that the passage is about, and not to the time at which the passage was written or communicated.[16]

Similar observations can be made about the other temporal indexicals. For instance, (9) could be continued with

(10) But alas, tomorrow was going to be an equally gruelling day.

Here *tomorrow* would be understood as referring to the 16-th of May, not to the day after the one on which the discourse was produced.

Such examples show that, at last in the temporal domain, the interpretation rule for the indexical need not be applied to the relevant component of the context of use (i.e., for temporal indexicals, the utterance time) but may in certain contexts make use of some other time instead. Or, to put this in the terminology introduced in the last section, the rule describes a process of reference determination; normally the origin of computation for this process is n_c, but under certain circumstances it can also be some other contextually salient time.

The upshot of this assessment would seem to be that the temporal indexicals have much in common with the elliptical calendar time designators discussed in the last two sections. In either case we are dealing with expressions whose reference is sensitive to context and the process which determines the reference often uses the utterance time as origin of computation, though it may also make use for this purpose of some other temporal element that is salient in the context. Are we to conclude from this that we are talking about a single, unitary phenomenon, the individual instances of which vary in that they are associated with different constraints on possible shifting of the origin of computation, but which nevertheless can all be seen as members of one family?

The answer to this question is an emphatic no. The similarity between temporal indexicals and elliptical calendar time designators is a superficial one and as soon

[16] These uses are more common in writing than in speech, but they can occur in spoken language as well. For instance, if (9) is transposed into the 1-st person it could come across as a quite natural report of how the speaker was feeling when on that 15-th of May he came home from a particularly strenuous day at work. See Kamp and Rohrer 1983 for further discussion.

as we look a little more closely, it becomes apparent how deep the semantic and pragmatic difference between these two types of expressions really is.

One immediate and quite striking difference is that the shifts away from n_c which are found with indexicals are almost without exception shifts towards the past, whereas for elliptical calendar time designators the origin of computation can just as easily lie in the future of n_c. In fact, it isn't just that with indexicals the shift appears to be possible only in the direction of the past (or, better perhaps, in sentences that are in the past tense). One also has a strong sense that where such a shift occurs, it is not just a shift in origin of computation for one temporal expression but a general shift of perspective, of the angle from which the clause or larger passage in which the shifted indexical occurs presents the information it conveys. This is especially notable in those passages in which shifted interpretations of indexicals are the most common, passages of free indirect discourse, in which the thoughts of a protagonist are represented from the perspective of that protagonist at the time when he is supposed to have had them. In these passages it is especially clear – but it seems that this is true generally for contexts in which temporal indexicals get shifted interpretations – that the shift away from the utterance time is not incidental to the interpretation of the indexical, but that the perspectival time, besides serving as origin of computation for the indexical, supplants the utterance time n_c more generally, so that n_c is no longer available as origin of computation for other expressions within the clause or passage either. For instance, suppose that (9) (with *the fifteenth* as opposed to *the fifteenth of May*) was uttered on the 22-nd of that same month and that this utterance was continued with one of the following two alternatives:

(11) He knew it was important too. For his biggest challenge would come in a week/ tomorrow and then he would need all his wits about him.

In the given context it seems virtually impossible to interpret *tomorrow* as referring to the day after the utterance time (i.e. to the 23-d of May), the only natural interpretation appears to be that where it refers to the 16-th. Similarly with *in a week*; here the natural interpretation returns the 22-d (or thereabouts), not the 29-th. Apparently, when the perspective shifts to the 15-th, the utterance time is thereby "eclipsed".

Nothing of the sort happens in standard cases where an elliptical calendar time designator is interpreted with the help of an origin of computation other than n_c. In such cases the utterance time remains nonetheless available for the interpretation of temporal expressions (e.g. indexicals) in immediately following sentences. For instance, consider the following two texts:

(12) I am not looking forward to the time ahead of me. For instance, next week I will have to go to Brussels on Wednesday and Paris on Thursday. But in fact, the worst day in the near future is **tomorrow**, when I will have to be in Berlin in the morning and in Amsterdam in the afternoon.

(13) We have a tough time behind us. Two weeks ago we had eight concerts within a single week. At least we had a rest day on Saturday. But then we played again on Sunday, Monday, Tuesday, Thursday and Friday. And both **yesterday** and the day before **yesterday** we had a matinee in one town and an evening concert in another.

In (12) the NPs *Wednesday* and *Thursday* are interpreted as referring to the Wednesday and Thursday of next week and the origin of computation here is the time denoted by *next week*. (To clinch the point we may assume that the utterance took place on a Monday.) Nevertheless *tomorrow* can be interpreted unproblematically (and only!) as referring to the day after the utterance time. (13) illustrates the same point, but now in relation to the past. All weekday names are interpreted as relating to the week preceding that of the utterance. Here, too, it is not the utterance time n_c which serves as the origin of computation to determine the referents of all those NPs. And yet n_c is available to assign *yesterday* its default interpretation of "day preceding the day of utterance".

These differences suggest that the similarity between elliptical calendar time designators and indexicals is only superficial. The analysis of indexicals proposed by Montague, Kaplan and others may not have been right in full generality. But it may nevertheless be considered right in spirit: In those cases which apparently contradict the predictions of this account always involve a "perspectival" shift, which is typically in the direction of the past, and which renders the utterance time inaccessible as origin of computation. In contrast, elliptical calendar time designators behave, in those cases where the origin of computation differs from the utterance time, more like anaphoric expressions: Any time that has a sufficient degree of contextual salience can serve as their origin of computation, and it doesn't supplant the utterance time, which remains available as origin or computation for other expressions in the same sentence.

If it is true that elliptical calendar time designators display anaphoric behaviour, it is natural to ask whether there are other anaphorically usable expressions which they resemble. There is a positive answer to this question: We find very similar properties with many definite descriptions.

It is by now a familiar observation that definite descriptions have both anaphoric and non-anaphoric uses. Moreover, in both uses their reference is, it seems, determined by the following principle:

Definition 3.8 The referent of a definite description *the a* in a context *c* is determined as follows: choose, by using an admissible delineation, a certain subset from the set of contextually accessible entities such that this subset contains exactly one satisfier of the descriptive content *a*. This unique satisfier is then the referent of *the a*.

As stated (3.8) is seriously underspecified, for what is a "contextually admissible delineation"? Only when this notion is clarified, will we be able to claim that we have gained significant insight in the way in which definite descriptions refer. But even so, (3.8) is not vacuous even as it stands; and it has a definite bearing on our question about elliptic calendar time designators. Whatever the details of delineation, it is clear from what (3.8) says that the chosen subset must on the one hand (i) be large enough to contain a satisfier, while on the other (ii) it must be small enough to contain no more than one. These two constraints are clearly exhibited by the mechanism of reference determination for elliptical calendar time designators we have described in (3.7): choose an interval around the contextually determined origin of computation which is big enough to contain one instance of the relevant calendar predicate but not so big that it contains more. In fact, elliptical calendar time designators form a special subcategory of the expressions governed by (3.8) in that for them the delineation constraints can be stated quite precisely. That this is so, is apparently connected with the clear and simple structure of the relevant domain – the linearly ordered medium of time – which has allowed for a high degree of conventionalisation of the way in which the referents of calendar expressions are determined.

Definite descriptions and elliptical calendar time designators also resemble each other in that the domain within which delineation takes place may either be determined by the utterance context, or, in the anaphoric case, by the discourse context. For the elliptical calendar time designators this is the difference between (a) choosing as origin of computation the utterance time and (b) identifying it with some time made salient in the antecedent discourse; for definite descriptions generally the difference is that between (a) selecting a subset from the environment in which the description is uttered and (b) selection of a subset from a set of salient discourse entities.

A further similarity is that elliptical calendar time designators are, like definite descriptions, essentially allocentric. To refer to the utterance time with a calendar time designator, such as, say, *10 o'clock*, is infelicitous or "marked". In the cases where this happens it is either because the speaker is unaware that the term he uses does in fact refers to the time of utterance or else because he wants to imply that he is referring to the time in some other capacity and that its coincidence with the utterance time involves some kind of contingency – as for instance when someone says: *He said he would be here at 10 o'clock. It is 10 o'clock now. So where is he?*. The same is true for calendar expressions referring to days – as we noted earlier, we normally do not refer to the day on which we say something as *Wednesday*, even if it is a Wednesday – and so on. Much the same also holds for descriptions referring to non-temporal entities. If the President refers to himself as *the President*, his utterance must be taken as referring to himself in his role or capacity of president, and even then it would probably come across as a little strange.

In this last regard calendar time designators clearly differ from temporal indexicals like *now* or *today*. These preferably refer to times directly related to the utterance context – to the time of utterance or to the day containing that time. Perspectival shift allows them to refer occasionally to other times as well, but these cases are marginal and derivative. With the calendar time designators the situation is the reverse, with reference to utterance-related entities such as the utterance time or the utterance day being the exception rather than the rule.

Let us briefly recapitulate the gist of this section. The task we set ourselves was to locate the elliptical calendar time designators within the general landscape of referring expressions. The main result of this undertaking was the distinction to which we have been led between two types of reference mechanisms: on the one hand the mechanism governing the genuine indexicals and on the other the one which determines the reference of definite descriptions. Elliptical calendar time designators, we found, pattern with the latter, not with the former. This shows that the appearance of indexicality that attaches to certain uses of elliptical calendar time designators is appearance only.

4 Reference to Times and Reference to Amounts of Time

4.1 Ordinals and Cardinals

We have encountered ordinals – the words *first, second, third, ...,* – before, but only within the conventionalized use exemplified by *the third of July* and the analogous use of elliptical NPs like *the third*. These uses are special, we saw, in that they are understood as referring to the third day of some month, even though the concept *day* is not explicitly mentioned (and, in the second case, the concept *month* is not either). But in fact there is still a further respect in which their interpretation involves convention. We can see this element of convention when we compare e.g. *the third of July* with the NP *the third day in July*. What is the referent of *the third day in July*? Well, that depends. *the third day in July* could perfectly well be interpreted as referring, just like *the third of July*, to July 3 of the year in question. But it need not. Here is a context in which it can be interpreted differently. Suppose we have made a list of days that are possible candidates for a certain purpose, e.g. a one day workshop of a project that involves several university and industrial partners. Our list, which only contains dates for the summer of 1998 – for that is when we want the workshop to take place – has six dates in July. In an attempt to come to a final decision about the date of the workshop you say, looking at the list of which all of us have a copy in front of us,

(14) The third day in July seems to be the best option.

Here it is perfectly possible to interpret the phrase *the third day in July* as referring to the third date in July appearing on the list, even if that day is, for instance, the 18-th.

The point of this example is as follows. In general the interpretation of NPs of the form *the n-th* α, where α is any common noun phrase, involves (i) a contextually restricted subset A of the extension of α, and (ii) a contextually salient order of A. NPs that refer to times are no exception to this principle (except for the special use which preoccupied us throughout Section 3). (14) is a case where the set A consists of the days mentioned on the list and the order is (we were assuming) the order in which these dates appear on the list. This, in other words, is a case where the order on A is given in some other way than by the intrinsic order of time, even though what is being ordered, the elements of the set A, are times.

Nevertheless a case like this is rather *cherché*. The typical situations in which we use ordinals are of two kinds: (i) situations in which the set A is a subset of a domain for which there is an intrinsic order (such as the natural numbers, or the letters of the alphabet; and (ii) situations in which the order is not intrinsic, but involves temporal order in some way or other – e.g. it may be the order in which the members of A are presented to an observer, or the order in which they were born or came about. In particular, when the set A itself consists of non-overlapping times (such as, e.g., when A is a set of dates), then the intrinsic temporal order of A is a very strong and rarely overruled default (though, as our example has shown, it is one that *can* be overruled).

It follows from these observations that ordinals should be seen as context-sensitive modifiers of common noun phrases α, and that the extension of the resulting predicate *n-th* α (for arbitrary α and n) is determined as a function of (i) the meaning of α, (ii) the relevant contextual information about the intended extension of α, (iii) an ordering of this extension, and (iv) the number n. In general this involves the utterance context in ways not previously encountered in this paper: The context c must provide, for the given common noun phrase α (i) a subset A of the extension of α, and (ii) an order $<_A$ of A. We take it that ordinal interpretation requires $<_A$ to be a pseudo-ordering and to be well-founded. Given A and $<_A$ it is not hard to state the extension of *n-th* α.

Determining the values of time designators containing ordinals involves, not surprisingly, an element of recursion. This is shown in (4.1), where we specify the values of such expressions in the model-theoretic format for which we laid th foundations in section 2. We will make use of this format also in the remaining subsections of section 4 and discuss it further in the (all too brief) section 5.[17] To simplify matters we represent all ordinals as having the form *n-th*. Expressions such as *first*, *second*, *third* are treated as synonyms of the "canonical" *1-th, 2-th, 3-th*.

Definition 4.1 Suppose that $\mathcal{M}$ is a model, that u is an utterance containing the expression *n-th* α, that c is the utterance context of u and that c determines, besides the utterance time n_c, (i) a subset A of $[[\alpha]]_{\mathcal{M},n_c}$ (the contextually relevant extension of α) and (ii) a well-founded linear order $<_A$ of A; and further that $\mid A \mid \geq k$. Then for $n < k$, $[[\alpha]]_{\mathcal{M},n_c}$ is defined as follows:

(i) $[[$ 1-th $\alpha]]_{\mathcal{M},n_c}$ is the first element of $<_A$;

(ii) $[[$ n+1-th $\alpha]]_{\mathcal{M},n_c}$ is the first element of $<_A$ following $[[$ n-th $\alpha]]_{\mathcal{M},n_c}$.

The applications of (4.1) which are of interest here are those where (i) $[[\alpha]]_{M,n_c}$ consists of periods; (ii) the standard ordering between periods is well-founded on the contextually determined subset A of $[[\alpha]]_{M,n_c}$.

The contextual selection of A is often subject to the principle which we have seen at work in connection with elliptical calendar time designators. For instance, the extension of the common noun phrase *second Wednesday in July* will typically consist of a single day. This is what (3.7) predicts for any use of the phrase in which the mechanism of section 3 assigns a single month to the calendar predicate *July*.

The meaning of an expression of the form D *n-th* α, where D is some determiner, is to be computed from the meanings of D and *n-th* α in the usual way. Since, according to (3.7) the extension of *n-th* α is a singleton, one would expect that the determiner D can be the definite article *the* and nothing else. This is more or less true but not quite. For even if there may be a predominance of occurrences of *the*, NPs

[17] Though we haven't made explicit use of this format up to now, we trust that it is reasonably clear how the semantic analyses proposed in earlier sections could be recast within it. Note that in (4.1) below we have ignored the intensional dimension of our models, omitting all reference to the world w at which the expression is being evaluated. We will stick with this simplification also in the earlier parts of the next subsection, to return to our original, explicitly intensional format towards the end of that section.

like a *second Wednesday in July* and *every second Wednesday in July* are perfectly grammatical and meaningful. At first this may suggest that (4.1) cannot be right, but a short reflection shows that there really is no conflict. First note that the phrase *the second Wednesday in July* also allows for functional interpretations, as discussed in Section 2.6. (An example would be a sentence like *the second Wednesday in July isn't a holiday in Baden-Wuerttemberg*). As we suggested in Section 2.6, such a use of a phrase gives rise to a functional denotation, which maps months of July onto their second Wednesdays. To get this interpretation the term *July* must be interpreted as an ordinary common noun, rather than as a calendar term that triggers a contextual mechanism for selecting one particular July. Each of the different months in the extension of this common noun determines a separate set A of Wednesdays, from which application of (4.1) then selects the second one.

Much the same happens, it must be assumed, in the interpretation of *a/every second Wednesday in July*: The common noun phrase of this NP, *second Wednesday in July*, is to be interpreted as if it said: *second Wednesday in some July*: We get the denotation by taking any j in the extension of the predicate *July*, computing the set $A_j = \{ z : Wednesday(z) \wedge z = j \}$, applying *2-th* to A_j to get the singleton set $\{ z_j^2 \}$ consisting of the second member of A_j, and then forming the union of the sets $\{ z_j^2 \}$ over the extension of *July*:

$$[[\text{ second Wednesday in July }]] = \bigcup_{j \in [[\text{July}]]} \{ \{ z_j^2 \} : z_j^2 \in [[\text{z-th}(A_j)]] \}$$

This isn't the whole story. There are two matters that need careful spelling out. One has to do with the occurrence of the term *July* in these phrases. That *July* can be interpreted in these contexts as a *July* is probably one of the many syntactic idiosyncrasies connected with calendar times, but the issue requires closer attention. The other issue is the matter of scope: What licenses carrying out the different semantic operations in the order in which they must be carried out to provide the desired interpretation, and what is responsible for the fact that they must be executed in this order? The last question is probably the easiest to answer: the determiners *a* and *every* both come with a presupposition that the extension of the common noun phrase with which they combine potentially contains more than one element. This presupposition rules out the other interpretation which we have seen to be possible, in which *July* is taken to refer to one particular month. (Note that the phrases *a/every Wednesday in July* are ambiguous between the two interpretations of their common noun phrase because even a single July contains several Wednesdays.)

That the common noun phrase *can* be parsed in a way which yields the desired order of semantic operations is not surprising but this too requires a closer look. In particular, it isn't immediately clear why the phrase has the interpretation it has, but not the one which assigns it the unit set consisting of the very first day that was the second Wednesday of any July. (This interpretation won't be available with *a* or *every* because of the non-uniqueness requirement just mentioned, but should be combinable with *the*. However, we do not think that the second Wednesday in July has this reading, while for instance, it does seem available for the phrase *the second Wednesday in any July*.)

The class of calendar-related expressions involving ordinals contains other curiosities besides these. First we find phrases like *the second Wednesday of the month*. This phrase seems ambiguous in the same way as *the second Wednesday in July*, with a reading where *the month* refers to some particular month and the one where *the month* is to be understood generically and, consequently, the entire phrase as well. In the light of what we have seen above, this isn't surprising; but it points towards another delicate issue which we have ignored in this paper: which noun phrases permit generic interpretations, and which articles (*the*, *a*, $\emptyset$) should be used when a generic interpretation is intended. Another curiosity are phrases like *every second Wednesday*. The default interpretation of this NP is that of a quantifier whose restrictor consists of all the even members in some enumeration of Wednesdays. To fit a compositional account of this interpretation into the framework established by (4.1) together with the semantic principles of earlier sections, is not straightforward. But it seems to us that the account would have to involve the following. We argued that the generic interpretation of *the second Wednesday in July* and the standard interpretations of *a/every second Wednesday in July* involve determining the second elements of the sets A_j consisting of the Wednesdays contained in any given month of July j. That is, we start with a set of j's and to each we applied the operations (i) of forming this set of Wednesdays included in it, and (ii) applying the meaning of *2-th* to that set. It seems that the interpretation of *every second Wednesday* requires that we compute days by applying the meaning of *2-th* to different sets. But in this case the choice of these sets is itself dependent on the application of *2-th*. Intuitively the procedure is a dynamic one. We start somewhere in time, form an interval by moving from that time towards the future until we have found within it a satisfier of the predicate *second Wednesday*. Then we start over, forming a new interval, again proceeding as far as is necessary to find a new satisfier of *second Wednesday* within it, and so forth. While we think that this is the computational strategy behind the interpretation of this phrase – quite probably this is not really controversial – we do not know how this strategy might be formalized within the general framework we have developed here.

We have seen in this section that the basic principles involving the interpretation of ordinals seem fairly straight forward, both within a general and within a specifically temporal setting, but that there are nevertheless many temporal phrases containing ordinals which present additional problems of detail. Some of these seem to have to do primarily with the semantics, others with the syntax, but in all cases we are confronted with obstacles to a smooth and simple syntax-semantics interface. Such idiosyncrasies are of course found in all compartments of natural language. But in that of temporal reference they appear to be uncommonly common. This is, no doubt, due to the high degree of conventionalization of our talk about time, and especially about time as structured by the calendar. It makes the systematic treatment of the meanings of such expressions an unusually difficult affair.

4.2 Cardinals and Measure Phrases

Phrases of the type *five minutes*, *three days*, *18 months* or *a hundred years*, in which a temporal noun is preceded by a cardinal (i.e. by the name of some natural number) are very different from the combinations of temporal nouns and ordinals which we considered in the last section. How they differ is perhaps most clearly visible in uses such as that in (15)

(15) The task will take five hours.

The function of the phrase *five hours* in this sentence is that of denoting a quantity of time, the amount of time that the task is predicted to require. Its denotation is not some particular period of five hours, for someone may be able to assert (15) without having any idea exactly when the task will be carried out, or even whether it will be carried out at all. And what goes for (15) goes for all uses of such phrases: What they denote are *quantities* of time.

The first thing, therefore, that we must do before we can give an satisfactory account of the semantics of such phrases is to clarify the notion "quantity of time". More generally, what we need is a characterization of the concept of "quantity". Given that, "quantity of time" may be expected to fall out straightforwardly as a special case.

The notion of a quantity belongs to the general theory of measurement. This is not the place for a disquisition on the subject of measurement theory, however, and we will draw attention to just those elements of it that we need for our present purposes[18].

The concept of measurement that is relevant for our present needs is additive measurement. We follow the standard theory in assuming that an *additive* measurement function is a function which maps objects of certain sorts – those to which the given type of measurement is applicable – to measurement "results", objects from some ordered additive structure of possible outcomes.

Definition 4.2 Let A be any set and let m be a function from A into a structure $\langle R, <, + \rangle$, where R is some set, $<$ a linear order of R and $+$ a 2-place operation on R which satisfies the usual axioms of an additive group and is monotonic with respect to $<$ (that is, we have for all $x, y, u, v \epsilon R$, if $x \leq u$ and $y \leq v$, then $x + y \leq u + v$). Then m is a *measurement function* for A.

The idea behind this definition is that m represents some "measurement procedure" which, when applied to an object a of A yields the measurement value m(a). As a rule the structure $\langle R, <, + \rangle$ will be some familiar number structure, most often that of the real numbers. We will limit our attention to this special case.

In order that m can be thought of as representing an *additive* measurement procedure, the addition operation on the measurement results (the elements of R) must have an empirical interpretation: There must be some way of combining the things measured so that the result of measuring the combination is the arithmetical sum of the results of measuring the things combined: If a+b is the result of "combining" the objects a and b, then m(a+b) = m(a) + m(b). For many types of measurement, such as e.g. that of measuring temperature, the mode of "combination" is quite abstract and indirect. But there are also measurement functions where the empirical interpretation of additivity involves the possibility of combining objects from A into "compound objects" in a quite literal sense. A particularly straightforward case of

[18] For a detailed account of the general theory of measurement see Krantz et al. 1971.

this is weighing (with a scale balance): We can form "compound" objects out of two objects a and b just by putting them jointly on the scale. The result of weighing this compound object will be the sum of the results one gets when weighing the two objects individually. Another familiar example is that of measuring length. Any two rods r_1 and r_2 whose lengths (in cm, say) are l_1 and l_2, can in principle be put end to end so that their combination forms a straight line. The length of this compound "rod" will then be $l_1 + l_2$.

The case of spatial measurement – this is true not only of measuring length, but also of its 2- and 3-dimensional counterparts, measuring area of surfaces and volume of bodies – has a further property, which it shares only with measurement of time. We think of the objects for which spatial measurements are possible as occupying parts of space. These parts of space can be measured – and thus assigned lengths, areas or volumes – just as the objects occupying them. In fact, it is a crucial feature of spatial measurement that the results of measuring an object and measuring the space it occupies are necessarily identical. Moreover, parts of space have their own modes of combination. Here the combining isn't something we can do, the combinations are already there as part of the overall structure of space. For instance, the parts of space occupied by two rods that have been put end to end form a compound part of space whose length is the sum of the lengths of the component parts.

The case of measuring time is in certain respects like that of measuring space. Here too, measurement is applicable not only to things that are thought to be "in" time, viz events and states of affairs, but also to the stretches of time occupied by them. Also, just as in the case of space, we think of the structure of time as including an operation of combining parts of time into one. For time this is the operation we already discussed in Section 2, the one which assigns to any pair of abutting intervals t_1 and t_2 the "sum" interval whose starting time coincides with the starting time of the first and whose end coincides with the end of the second.[19]

What is a quantity? The notion of quantity we are concerned with here is one that is dependent on the notion of measurement. With each measurement function m we can associate a corresponding notion of quantity Q_m. The guiding idea is that two objects in the domain A have the same "m-quantity" iff their m-values coincide. (Thus two objects are thought to have the same "weight quantity", or mass, if they yield the same result when weighed. And two events or two periods of time have the same "time quantity" if we get the same result when clocking them; "time quantity", in other words, is what we normally call "duration".)

Still, what *is* a quantity? To some extent this is a matter of stipulation. Any kind of abstraction from the relation $\equiv_m$ of "having the same m-quantity" will do as long as it assigns entities to the objects of A in such a way that two elements of A are assigned the same entity just in case they stand in the relation $\equiv_m$. The standard construction for this purpose is that of forming equivalence classes under $\equiv_m$.

[19] An important difference between the case of time and that of space is that the temporal objects cannot be moved from one temporal location to another. Connected with this is the very different physical basis of time measurement as compared with the measurement of space. We must refrain, however, from going into this here.

Definition 4.3 Let m be a measurement function for A. The *quantities associated with* m are the equivalence classes of elements of A under the equivalence relation $\equiv_m$, which holds between a and b iff m(a) = m(b).

According to this definition each measurement function uniquely determines a corresponding notion of quantity. The converse of this is not true. The same notion of quantity is generated by any one of an indefinite number of measurement functions. In fact, whenever two functions m and m' satisfy the condition that for all a, b in A $m'(a) = m'(b)$ iff $m(a) = m(b)$, then the set of quantities associated with m' will be identical with the set of quantities associated with m; and we may note in passing that this condition is entailed by the one that m' is a *strict monotonic transform* of m, i.e. if for all a, b in A, $m'(a) < m'(b)$ iff $m(a) < m(b)$.

A many-one relationship between measurement functions and quantity concepts is not just a theoretical possibility; in the actual practice of measurement it is the rule rather than the exception. Measurement of time is a case in point. The conceptual machinery to which it has given rise involves a single notion of quantity in conjunction with a number of different measurement functions, each of which generates the same set of quantities. For our daily needs and purposes we assess amounts of time in terms of the number of years, months, weeks, days, hours, minutes or seconds they last; moreover, science has, to suit its purposes, extended this scale of units considerably at either end. Any two of these ways of assigning numbers to amounts of time are represented by distinct measurement functions – obviously so, for they must assign distinct numbers to the same period. (For instance, a period of two calendar years will be assigned the number 2 when counted in years, the number 24 when counted in months, 730 when counted in days, etc.)[20]

We must add right away that in this as in other cases the relationship between the different measurement functions that all generate the same set of quantities is quite tight. The functions of which we have just spoken differ from each other only in that they make use of different *units of measurement*. The *unit of measurement* of a measurement function m is that quantity 1_m from the set of quantities associated with m to which m assigns the number 1. We say that m and m' *differ only in the choice of the unit of measurement* if they are multiplicative transforms of each other: there is some fixed number r such for all a in A m'(a) = r. m(a). Evidently when m and m' are related in this way, then $r = m'(1_m)$.

That different measurement functions which are used in conjunction with the same notion of quantity are related by these tight connections is of course no accident. In fact, it is a consequence of their being additive with respect to the same operation(s) of combination on the elements of A in conjunction with two further assumptions:

(i) the set A is Archimedean with respect to the combination operation: for any two elements a and b of A it is possible to combine a with itself (or with things of its size) so often that the resulting compound has a size exceeding that of b; and

(ii) each a in A can be partitioned into two parts of equal size).[21]

---

[20] See next footnote.

In the light of these observations it follows in particular that the models we have defined in Section 2 will generate unique measurement functions that assign real numbers to arbitrary intervals $(t,\ t')$ once one interval is chosen as unit of measurement. More precisely, they generate unique measurement functions (modulo the choice of unit) that are additive with respect to the combination operation which turns two adjacent intervals $(t,\ t')$ and $(t',\ t'')$ into the interval $(t,\ t'')$ provided the time structure of the model satisfies the conditions (i) and (ii). It follows from the assumptions we made in Section 2 about time structures that each model $\mathcal{M}$ determines a unique measurement function $m_{\mathcal{M}}$ which assigns a real number $r = m_{\mathcal{M}}(p,q)$ when applied to any pair of intervals p and q of $\mathcal{M}$; intuitively r is the size of p when q is used as giving the unit of measurement. We will denote as $m_{\mathcal{M}}(q)$ the function which assigns to each interval p the number $m_{\mathcal{M}}(p,q)$.

Division of A by some equivalence relation is the standard method for obtaining abstract entities. But often there are alternatives. Among these there is the possibility of using the same equivalence relation to do division on some natural subset of A which contains representatives for each of the equivalence classes obtained by the standard method. In the case of time such an alternative seems particularly natural. Among the entities that can be measured for temporal size we find events, processes and states of affairs as well as periods of time. However, if we assume that each event, process or state of affairs occupies a temporal interval which (for this very reason) must be of the same duration, then intervals of time are all we need to generate the full complement of quantities of time, that we also get when events, etc. are included within the division into equivalence classes as well. This implies that we may define quantities of time as equivalence classes of intervals without having to worry that this might leave some quantities out. And that in turn implies that we can define quantities of time for the models described in Section 2 just in terms of their time structures. These considerations motivate the following definition:

Definition 4.4 Let $\mathcal{M}$ be a model as defined in Section 2.5 and let $\langle T, <, \equiv\rangle$ be its time structure. Then the *quantities of time* in $\mathcal{M}$ are the equivalence classes generated by $\equiv$ on the set of all periods of T.

One consequence of this last definition is that the measure predicates ($year_m$, day_m, $minute_m$, etc.) which were discussed extensively in Section 2 can be construed as

[21] In outline the argument goes as follows: Because of assumptions (i) and (ii) we can, for each pair of elements a and b of A and natural number n, find an element c of A and numbers $r, s > n$ such that $m(r.c) < m(a) < m((r+1).c)$ and $m(s.c) < m(b) < m((s+1).c)$, where r.c is the result of combining c $r-1$ times with itself, etc. Because of additivity, this last condition is equivalent to: $r.m(c) < m(a) < (r+1).m(c)$ and $s.m(c) < m(b) < (s+1).m(c)$, and this entails that $(r/(s+1)) < m(a)/m(b) < (r+1)/s$. Since the fractions on the right and the left converge to the same number with growing r and s, the "m-ratio" m(a)/m(b) of a and b is completely fixed by the operations of combination and division. Thus once a unit of measurement has been fixed by stipulating for one element of A that the result of measuring it is 1, then this also fixes the measurement values for all other elements of A. Consequently any two measurement functions which are additive with respect to the same combination procedure will differ at most by a multiplicative constant.

names of quantities. For they all have the property that their extension consists of (all and only the members of) exactly one quantity; so we can see the predicate as a designator of that quantity. For this reason measure predicates naturally assume the role of units of measurement. For instance, choosing *second* as the unit of measurement amounts to stipulating that the elements of its extension are assigned the value 1 (which is consistent given that all members of it stand to each other in the relation $\equiv$); this then fixes the values to be assigned to all other intervals as well. To reflect the possibility of using measure predicates to fix the unit of time measurement, we write, for any model $\mathcal{M}$ and measure predicate P, "m($\mathcal{M}$,P)" for the measurement function on the intervals of $\mathcal{M}$ which assigns 1 to the members of the extension of P.

At last we are ready to return to the semantics of measure phrases such as *five hours* which prompted this extensive excursus. The first matter that needs clarification concerns the semantic type of such expressions. On the one hand they seem to function as predicates – just as $hour_m$ acts as a predicate the extension of which consists of all periods which are of the duration of exactly one hour, so *five hours* would seem to be a predicate whose extension consists of all and only the periods that are five hours long. But as we have already noted, many typical uses of such phrases indicate a different logical type. For instance, in (15) *five hours* would seem to play the role of a term that refers to a particular quantity, and a similar impression is associated with (16), in which the same phrase appears as subject.

(16) Five hours is a long time.

In the light of what we have said on the subject of quantities, this appearance of type ambiguity should not be surprising. For as we saw, quantities are naturally construed as equivalence classes and so one might well have expected that there are expressions of which it need not be fixed once and for all whether they should be construed as designators of certain quantities or as predicates that have these quantities as their extension.

Nevertheless, the syntactic environments of measure phrases seem to suggest quite strongly that we are dealing with noun phrases, and not with common noun phrases. And if this is how we analyze measure phrases syntactically, then semantically they should be treated as referring expressions, whose referents are quantities. This is the line we will take.

The rest is straightforward. Given our last decision and all that we have said about quantities above, the principle which determines the reference of measure phrases is easy to state. It is given in (4.5).

Definition 4.5 Let ν α be a noun phrase consisting of a cardinal expression ν and a temporal measure predicate α and let $\mathcal{M}$ be a model. Then $[[\nu\alpha]]_M$, the denotation of $\nu\alpha$ in M, is that quantity of time q of M, such that if $i \in q$ and $j \in [[\alpha]]_M$, then $m_{\mathcal{M}}(i,j) = [[\nu]]_M$.

A few comments are necessary. First, we are assuming that the denotations of α and ν in M have been defined. For α this assumption is unproblematic: We had already made the assumption that the interpretation function of M maps measure predicates onto quantities. The case of ν is more problematic, for one thing because we haven't said yet what we mean by a "cardinal expression." So far we only had examples of phrases of the form $\nu\alpha$ in which ν was a true cardinal, e.g. *two* or *three*. But in fact there is no reason why we should limit our attention to expressions denoting the natural numbers. For on the one hand measure phrases $\nu\alpha$ in which ν denotes a non-integral number – as in *three and a half years*, *7.93324 nanoseconds*, etc. – are as much part of natural language as phrases with "true" cardinals. And on the other the function $m_{\mathcal{M}}$ will in general return non-integers as well as integers as values. Nevertheless we have to be careful. For nothing we have said so far warrants that the function $m_{\mathcal{M}}(a)$ (where $a = [[\alpha]]_{M,w}$) is onto the set of the real numbers. So in principle it could happen that ν denotes a number which isn't in $m_{\mathcal{M}}(a)$'s range. In that case there won't be a q in the sense of (4.5), and $[[\nu\alpha]]_{M,w}$ will remain undefined. As a matter of fact, it is a natural and widely accepted assumption that this can't happen: the function which assigns intervals of time to their duration is supposed to be such that each real number is realised as the duration of some temporal period. If we adopt this assumption as a further constraint on models – $\langle T, <, \equiv \rangle$ must be such that the function $m_{\mathcal{M}}$ is onto $\mathcal{R}$ – then this problem is out of the way. But there is still the linguistic issue as to what expressions should actually be allowed. And there is also the question which numbers are denoted by which expressions, i.e. what in general is $[[\nu]]_M$? Intuitively this seems a rather different sort of question as compared with most others of the form "what is $[[\beta]]_M$?" The grammar of number expressions is pretty much a world to itself, and a "model-theoretic semantics" for such expressions is hard to distinguish from an algorithm for translating the expressions into some kind of canonical notation. We won't go into this any further but simply assume that for all expressions that can take the place of ν in measure phrases, the denotations $[[\nu]]_M$ have somehow been defined.

This concludes what we have to say about the denotations of the measure phrases themselves. But our discussion of such phrases cannot yet be considered complete. For one of their most common use is in prepositional and quasi-prepositional phrases such as *after five hours* or *three years later*. These and some other temporal phrases of (quasi-)prepositional form are the topic of the next section.

4.3 Measure phrases as parts of larger temporal expressions

Measure phrases may occur as argument phrases to verbs and adjectives, as in

(17) a. The film lasted two hours.

b. It took him three days to correct the paper.

c. He stayed two hours longer than she did.

d. Thirty five minutes is too short.

But they also occur as part of other time referring expressions, and it is mostly to those occurrence types that this section will be devoted. Here are some examples:

(18) a. in two hours

b. for two hours

c. twenty minutes before midnight.

d. half an hour after the show started

e. half an hour after the show

f. a few minutes later

g. two days ago

About the occurrences in sentences like those in (17) there is little we have to say. As regards (17a) and (17b) the only point we wish to make concerns the verbs *last* and *take*. Given our decision to interpret measure phrases as denoting quantities we are more or less forced into an analysis of these verbs according to which they express relations between (i) quantities and (ii) whatever can occur as their logical subject. It seems that their subjects are always "eventualities" (i.e. states, processes or events). *Last* is therefore to be analyzed as a 2-place relation between eventualities and quantities of time. In fact, *last* is a "quasi-logical" relation, which can be defined provided we assume that each eventuality determines an interval, as the period during which it was going on. Thus, given a function Dur which assigns to each eventuality ev the interval of time occupied by ev in w, the extension $[[last]]_{M,w}$ of *last* in any model M at a world w should be the relation which holds between any eventuality ev and quantity of time q iff $\mathrm{Dur}(ev)(w) \in q$. The story about *take* is similar, except that here it seems that we are dealing with an essentially 3-place relation, the terms of which are (i) an eventuality, (ii) a quantity of time and (iii) a "beneficiary". In order that the relation which *take* expresses hold between ev, q and some individual x it is again required that $\mathrm{Dur}(ev)(w) \in q$, but there is a further meaning component to the effect that ev must be an action performed by x which x succeeds in bringing to a close; q gives the amount of time from the beginning of the action till the time when it was accomplished. To capture this second meaning component we need a much more powerful framework than we have given ourselves here (one that allows for the representation of propositional attitudes and which can deal with presupposition). So this is an issue that doesn't belong here.

Last and *take* belong to a small class of verbs which express relations between an eventuality, a quantity of time and, in certain cases, something else which also contains, for instance, *spend* and *waste*. For many, and perhaps all of these verbs the condition "$\mathrm{Dur}(ev)(w) \in q$" is part of their application conditions.

In the two other examples in (17) the measure phrase occurs as an internal argument to an adjective. Adjectives for which this is possible have sometimes been called "measure adjectives" in the literature, a term which draws attention to the fact that the adjective expresses a property of its subject that can be assessed by some kind of measuring procedure. What kind of measure phrase can appear with such an adjective depends, obviously, on the kind of quantity that the associated measuring procedure determines. In particular, adjectives which combine with the measure phrases that concern us here must say something about the duration of their subjects. So their subjects must be the sort of things that have duration. It seems that here we find as possible subjects not only eventualities but also intervals (as in *this interval is two hours longer than that one*). But again we find that the meanings of the predications in question involve an important "quasi-logical" component. For instance, (17c) speaks of two eventualities, ev_1, the state of his being at the party, and ev_2, the state of her being at the party, and it asserts about these eventualities that the duration of the first exceeds that of the second by two hours. Thus $\langle ev_1, ev_2 \rangle$ is in the extension of *two hours longer* (in M at w) iff there is an interval i such that $\mathrm{Dur}(ev_2)(w) + i = \mathrm{Dur}(ev1)(w)$ and $i \in [[\textit{two hours}\,]]_{M,w}$. On the role of measure phrases in comparatives much has been written elsewhere (viz. the literature on comparatives), so this too isn't a topic to be dwelt on here.

The examples in (18) can be divided into two classes. Those in (18a) and (18b) act essentially like measure phrases. The prepositions in these PPs only serve to relate the denotation of the measure phrase they govern to the rest of the clause. The other examples in (18) are all "temporal locating adverbials", which denote intervals, not quantities of time.

About the difference between (18a) and (18b) much has been written in the literature on aspect. (18a) only combines with telic eventuality types, which place an intrinsic condition of termination upon the eventualities that instantiate them. Moreover, (18a) and (18b) are in complementary distribution: (18b) can only be combined with eventuality types which lack such an intrinsic termination condition. Important as this difference is, it is not one that need detain us here and we mention it only in order to separate it from the issue that does matter in the present context. This is the question what temporal relation the preposition establishes between the quantity denoted by the measure phrase and the eventuality described by the clause. Again this temporal relation is the duration relation we already encountered and it is the same for both *in* and *for*. Thus, both in (19a) and (19b) the PP expresses that the eventuality ev of which the sentence speaks is related to the denotation q of the measure phrase *two hours* by the condition "$\mathrm{Dur}(ev)(w) \in q$".

(19) a. Bill crossed the lake in two hours.

b. Bill swam for two hours.

(That the PP in (19a) acts as a predicate of an independently identified event, while the PP of (19b) plays a part in delimiting the eventuality contributed by the rest of the sentence is, as we said, an important difference, but one whose analysis belongs elsewhere.)

The phrases in (18c) – (18g) are all of the kind that denote particular times (or, sometimes, the corresponding generic objects, see Section 2.6). In all these cases the measure phrase serves to "compute" the time referred to by the entire phrase from that given by the embedded NP or *that*-clause, with the connecting preposition indicating the nature of the computation. The simplest cases are those like (18c), where the embedded time denoting expression refers to a moment rather than an interval. Suppose that t is the time that the embedded expression b refers to and that the preposition is, as in (18c), *before*. Then the denotation of the entire phrase is the time t' such that (t', t) belongs to the quantity q denoted by the measure phrase a. The rule for phrases of the form *a after b* differs in that the relevant interval is $\langle t', t \rangle$ rather than (t, t').

(18d) is in essence just like (18c). This time the embedded time denoting expression is an event-describing clause. But if we assume that startings are instantaneous events, then the principle that we gave for the denotation of (18c) will work again, provided we let $\mathrm{Dur}(e)(w)$, where e is the event of the show's beginning, play the role of t. A little different is (18e), as there the denotation of the embedded phrase is not a moment. What seems to count in such cases is the boundary of the duration of the event in the direction indicated by the preposition; in the case of *before* it is its beginning and in that of *after* it is its end.

(18f) and (18g) differ from the preceding cases in view of their context sensitivity. Like the elliptical calendar time designators we discussed in section 3 they demonstrate the need for a systematic treatment of context – any reasonable semantics of temporal phrases will have to include some such notion of context.

Regretfully we had to abandon our original plan of concluding this paper with an explicit model-theoretic account of a fragment of English temporal phrases. Nonetheless we want to devote, in the next and penultimate section, a few words to the question what notion of context such a model-theoretic account might incorporate.

5 Denotation in Context

We must bring context into our model-theoretic analysis of denotation. So we need a suitable, formal notion of context. But it is not easy to decide what that should be, not so much because it is hard to find reasonable suggestions, but because suggestions are all too easy to come by. The structures of contextual information that are available to human interpreters tend to be very rich, far too rich to permit transparent formalization in full (if they permit full formalization at all). Theory has to make do with something much more modest and surveyable. What this should be depends on the individual theory's needs. But even when it is well-understood what the theory's needs are, there remain non-trivial questions about the form in which the required contextual information is to be represented.

From our discussions of context sensitivity in Section 3 it is clear that a notion of context that is to serve our purpose must have two components. First, we need a component which contains the necessary information about the utterance context. Since here we are just concerned with temporal reference, the only bit of information we need about the utterance context is the utterance time. But what is to be understood by an "utterance time" within the kind of model-theoretic setting we have been using? To answer this question we should remind ourselves that as soon as we are dealing with the context-sensitive aspects of denotation, we are always talking about the reference of (uttered) expression *tokens*, not expression types. These tokens are produced within a certain world. We will assume that this world can be identified with some member of an intensional model M, and thus that it is within one of the worlds of M that the utterance is produced, and that it is produced at some particular time. That time, a time from the time structure of M, is then the utterance time for the utterance in question. This means that contexts, or at least their utterance context parts, are supervenient on models: Given some model M we can speak of contexts relative to M. We cannot speak of contexts absolutely.

Secondly, we need, in order to be able to provide a creditable analysis of temporal anaphora, a certain amount of information about the discourse context (as well as, possibly, information from the common ground between speaker and audience which is in place when the discourse starts). Again, we are only interested in contextual elements that might be relevant in connection with temporal reference. So we need to include in the context only elements which can play the part of anchor points for the context-sensitive interpretation procedures – elements, that is, which can serve as antecedents for anaphoric temporal expressions or as origins of computation. It has been argued that among the antecedents for temporal anaphors we find not only times but also events. Nevertheless we will assume here that the discourse context consists just of times, and that in those cases where it may be theoretically preferable to say that the anaphoric antecedent is an event, it is the period occupied by the event that acts as its proxy in the role of antecedent.

The discourse part of the context, then, will consist of a set of times. But what times? As in the case of the utterance time, the general set-up of our analysis seems to leave us little choice: The times should be elements from the model M that we take as point of departure and on which the context supervenes. But now there is reason for uneasiness. The elements which the earlier parts of a discourse introduce and leave behind as anchor points for the interpretation of the later parts typically have a conceptual status: They are *representations* of things, often with not enough information attached to them to determine the things they represent uniquely, but which may nevertheless suffice for their role as anchor points, guiding the interpreter towards the right anchor point for a given anaphoric purpose. In a dynamic-representationalist approach such as Discourse Representation Theory, where it is the discourse representation itself which acts as discourse context, this difficulty for the "classical" model-theoretic approach we have chosen here does not arise. Therefore, as the results of this paper will eventually be integrated into a more comprehensive DRT-based semantics in any case, the problem we are facing now is a strictly local one, which will simply disappear at that later stage. Yet, even if the problem is only a local one, it still deserves a local solution.

The compromise solution we want to suggest is the following. The times of the context will be times from the model M – for again: what else could they be? But at the same time we assume that these times have certain information attached to them, which is instrumental in the role they play in subsequent interpretation. For simplicity we will assume that the information is correct, i.e. that it is true in M (at the actual world w_0) of the times to which it is attached, although this is unrealistic and also isn't essential to the contextual purpose they will serve.

What information should come with the discourse times? We propose that it should be of two kinds. On the one hand it should include information about the temporal relations between the times themselves as well as between them and the utterance time. On the other, and this is especially important for our purposes here, there ought to be information about temporal "granularity" – information about whether a given time t is, say, a week, or a year, or a day or a clock time (i.e. a time introduced by a phrase like *at ten o' clock* or *at 7.15*), etc. We will assume that information of the former kind is given by conditions familiar from the DRT literature on temporal reference (conditions such as "$t < t'$", "$t \bigcirc t'$", etc.). As regards information of the latter kind, we will assume, in order to keep things simple, that it takes the form of a calendar predicate or of a measure phrase "nP_m" where n is a cardinal and P_m a measure predicate. (We could be more exact, but in view of the provisional status of the context notion we are introducing, this seems hardly worth the trouble.)

In this way we get to the following concept:

Definition 5.1 Let M be a model, $w \in W_M$. A *context relative to M and w* is a pair $C = \langle U_C, Con_C \rangle$ such that

(i) U_C is a set $\{ t_0, t_1, \ldots, t_k \}$ (with $k \geq 0$; t_0 is the utterance time of C); and

(ii) Con_C is a set of conditions; these are either of the form "$t_i R t_j$", where $i, j \leq k$ and R is one of the relations $<$, $\bigcirc$, $\supset\!\subset$, $\subseteq$, or of the form $P(t_i)$, with $0 < i \leq k$ and P is either a calendar predicate or a measure predicate as defined in earlier sections.

Denotations can now be made dependent on a context as well as on a model and world. In other words, the formal denotation concept with which we are now dealing is $[[a]]_{M,C,w}$, not simply $[[a]]_{M,w}$. (Note that all previous analyses of $[[a]]_{M,w}$ involved a's whose reference was independent of the context C. For such a, $[[a]]_{M,C,w} = [[a]]_{M,C',w}$ for any contexts C and C' and $[[a]]_{M,w}$ can be regarded as short for $[[a]]_{M,C,w}$ with arbitrary C.)

We hope that this last section, which is much more sketchy than the one we had planned originally, nevertheless gives a glimpse of the way in which the informal analyses of Sections 3 and 4 can be integrated within a more comprehensive formal semantics of temporal expressions and, more generally, of natural language fragments in which the temporal aspects of interpretation freely interact with other aspects, such as quantification, negation and intensionality. As we indicated in the discussion of the notion of "discourse context" in this section, our own preferences for such a more comprehensive semantics go in the direction of a DRT-based – or, more exactly, a UDRT-based – approach. But to the extent that the results of this exploration are right, they should be equally valid in a classical model-theoretic setting of the sort we have adopted here.

6 Conclusion

Two matters are on this conclusion's agenda, (1) what (in our own opinion) has been accomplished in this paper; and (2) what remains to be done.

1. The main points with which this paper has tried to deal are:

i. A formally elaborated assessment of the topological and metric structure of time as it is available to a typical speaker of English.

ii. An analysis, from the same perspective, of a number of calendar concepts and of the metric concepts corresponding to them, all of which play an evidently prominent role in the way we represent and handle temporal information about mundane affairs.

Although the conceptual analyses of i. and ii. have been presented in model-theoretic terms, they can for the most part easily be recast in axiomatic terms. In this latter form they can be used as the theoretical basis for the development of theorem provers designed to handle temporal information of the kind that humans use when they reason about their daily lives. Recall our conjecture that all deductive inferences in this domain which people can be expected to draw follow from the axioms.

iii. The next theme, discussed in Section 3, is the semantics of complex expressions which comprise several calendar terms of adjacent granularity (the so-called calendar time designators). As the main results of this discussion we see (a) a general formulation of the reference principle that governs all uses of such expressions, and (b) the subsumption of this principle under a more general one, which governs the denotation of definite descriptions (of which some refer to times but most to other things). An important implication of this subsumption is that elliptical calendar time designators are, superficial impressions notwithstanding, crucially different from indexical expressions, given how indexicality has been described in the literature; and the discussion throws, we think, some new light on the nature of indexical and non-indexical reference generally.

iv. Section 4 addresses the roles that are played in time-denoting expressions by (a) ordinal terms (*first*, *second*, *third*, ...) and (b) cardinal terms (*one*, *two*, *three*, ...). It is striking how radically different these roles are. Cardinals typically occur as constituents of expressions denoting quantities of time, while ordinals appear as constituents of expressions denoting moments or periods. A modest, though we think useful, service which this section provides is a systematic development of the notion of a quantity of time. There is nothing much about what we have to say on this matter that is original: The development we present is culled from the established literature on the theory of measurement. Still, in writings which focus on temporal reference in natural language the notion is normally not made explicit; yet for an integrated model-theoretic account of reference to both times and quantities of time an explicit account of it is indispensable.

2. Things that remain to be done:

In Section 5 we recognized the necessity of enriching our model-theoretic framework with a notion of context and made a start with incorporating such a notion. But we noted that a model-theoretic treatment of context can't be more than a stop gap measure which will have to be replaced by a dynamic treatment of context dependence, of the sort that is offered by DRT and other versions of Dynamic Semantics. (It is only within the setting of such a dynamic treatment that certain aspects of the interpretation of time-denoting expressions can be treated correctly; one example is the way in which origins of computation are determined for tokens of elliptical calendar time designators.)