Times and Events. Notes for a course at the University of Texas. Spring 2014

Hans Kamp

University of Stuttgart
University of Texas, Austin

June 5, 2014

1 Introduction

The topic of these lectures is 'times and events'. But what are 'times' and what are 'events'? Our first task is to give some answer to these two questions. But as soon as we start thinking about those answers, we find ourselves in the midst of hard problems. One of the problems is that the question 'What is time?' can be approached from different angles – from the angle of physics (or, more generally, from that of the natural sciences) and from the angle of psychology or linguistics. The first approach focusses on time as a physical phenomenon, as an aspect, or 'dimension', of physical phenomena. The second is concerned with how human beings think about time and how do they talk about it? There are striking differences in what these perspectives have to tell us about time. But both perspectives are important, and we must pay attention to each of them.

With regard to 'events' the situation is somewhat different. In the natural sciences the term 'event' does not play a very prominent role. though it is encountered sometimes, for instance when one speaks of distant events in Relativity Theory or of atomic events in Quantum Mechanics. In linguistics, on the other hand, the term has become well-established. In particular, events have become an essential ingredient in theories of Tense and Aspect, which are concerned with the means that human languages make available for talking about what happens in time and more generally about what is true at some times but not at others. But it is not easy to say what the events that are spoken of in these theories really *are*. The theories have much to say about the roles that events play (as part of what they have to say about the meanings of sentences from the languages they deal with). But what they say about the roles that events play doesn't fix *what* the things are that play these roles (or not at least in any obvious way). This is one of the big questions about the events that have come to play such a central role in linguistic theories (and, to a somewhat lesser extent, also in psychology).

In this introductory part of exploration we will therefore concentrate mostly on times. Times will be the theme through almost all of it. Only in the last subsection will there be a few more remarks about events.

Before we turn to real issues, one more general remark about these lectures. This is a course sponsored by the Linguistics Department. That suggests that the principal emphasis should fall on the linguistic aspects of our general theme. There is also another – also accidental, but less respectable – reason for focussing on these aspects: The linguistic side of the issues is the one that I am most familiar with. I would be the last to suggest that the physical aspects of our general topic, about which I know much less, would be of lesser interest. But my limited competences being what they are, it would be foolish to focus on parts of our topic that lie beyond them.

1.1 Informal remarks about time and times; (a) the psychological perspective

The question 'What is time?' we said can be approached from two different perspectives, that of physics and that of psychology or linguistics. The most important difference between these two approaches is this: From a psychological perspective, time – how we humans experience and understand time – has a *deictic center*: From my personal, experiential stance there is one time which is singled out as having a unique significance, which sets it apart from all other times. This is my current 'psychological now', or 'psychological present', the time at which I am right now and experience the things that are happening around me. It differs on the one hand from the times which make up the past – my ‘present past’, to be precise – those times when the things happened that now lie behind me and about which nothing can be done any more; and on the other from the times that are still to come, those which constitute my future, the times at which things will happen. Some of these are still under my control, or under the control of other agents. Because of that the future appears to be ‘open’, or ’indeterminate’ in a way that that the past is not. This is a difference between past and future about which more will be said later. In other words, my current psychological now divides the totality of all times into two halves, the (present) past and the (present) future, and acts as a kind of transition point (or ‘hinge’) between those two halves.

One of the striking and puzzling things about the psychological now is that as soon as it is there it is gone, and replaced by another psychological now, a time which until this moment had been one belonging to the (then present) future. And as soon as that time has become present, it too will be gone and replaced by yet another psychological now, that until that very moment, was still part of *its* future, and so on. But talking about the fleeting character of the psychological now in these terms – as an endless succession of distinct psychological nows, one after the other – doesn’t seem quite right. For as will be argued later on, there are good reasons for assuming that among the times that are in the future of my current psychological now there isn’t a first one, which will replace my current psychological now as the next psychological now after it. Between any such time and my current now there ail be other times, which should;d have become my current psychological now before that one does. How an agent’s psychological nows are to be thought of, from the purely internal point of view of his own psychology and from that of his place in the outside world and his interaction with it, is a challenging conceptual issue, to which we will return.

Whatever the details about how one psychological now replaces another, it seems clear that somehow such shifts to new psychological nows do take place. Every agent runs through an indefinite number of such psychological nows in the course of his or her lifetime, and when the life of one agent a comes to an end there will be other agents whose psychological nows will continue the line of psychological nows of a. And before a there were other agents whose psychological nows extend backwards from when a came into being and a series of psychological nows started. This personal and interpersonal range of psychological nows extending into the past and the future impose an order on times in general. Take me and my current psychological now t_{ps}. t'_{ps}, we said, divides the totality of times (excepting t'_{ps} itself) into my current past and my current future. But there is more than just this binary division. The succession of psychological nows also imposes a certain order on this totality. This follows more or less straightforwardly when we assume that all times in my current past were themselves psychological nows and that all times in my current future will be psychological nows. Let us begin by making this assumption. Let t_1 and t_2 be two times from my current past. Consider the situation in which t_2 played the part of psychological now. Then there are three possibilities for t_1: (1) t_1 was in the then current past; (2) t_1 was in the then current future; and (3) t_1 was also the then current psychological now. In the first case we can conclude that t_1 precedes t_2; in the second that t_2 precedes t_1 and in the last that t_1 and t_2 coincide, i.e. that they are one and the same time. It seems reasonable to assume that the times in my current future will be ordered in like fashion: if t_3 and t_4 are times belonging to my current future, and in the situation when t_3 plays the part of psychological now, t_4 belongs to what is the current future then, then that establishes that t_3 precedes t_4 (and thus, by the same token, that t_4 follows t_3). Furthermore, my current psychological now t_n counts, by the same kind of consideration, as preceding all times in my current future, and as following all times in my current past. And by implication all times in my current future follow all times in my current past.

One problem this raises has to do with when I have in a state that enables me to experience a psychological now. What when I am asleep, or knock out in some other way? Has there been no time (in the sense of psychological time) until I wake up, for come to? That would seem quite arbitrary, quite apart form the fact that it is often hard to tell whether or not I am asleep. But this problem is you might say dwarfed by another one. And it will be solved if the other one can. This other problem is that, at best, the characterization I have just given of time in terms of *my* psychology covers times that fall within my life span. (I am ignoring here the possibility that the soul survives the goody and that when it is not incarnated, it can be aware of time in a manner comparable to that in which we mortal are aware of it.) But surely there are more times than just those during which I am alive. There were times before I was born, and there will be times after I will; have died. And in comparison with all those times, those that at some point play the part of psychological now for me, in the course of the 'spec of time' that is covered by my life as a human mortal, appear to be no more than a vanishing few.

Arguably we can do better. I am not the only one around. There are all my fellow human beings. Some of those were born before me and some will outlive me. Suppose you were around before I was born. At any point when we are both alive, it might be argued, your psychological now and mine coincide. This assumption seems especially plausible when you and I communicate: we look at something that is happening right now, and may talk about it to each other. Essential to such situations of joint observation and communication about what is jointly observed is that you and I think of each other as having the same psychological now as we have ourselves. Everything in the social psychology of such interactions points toward the conclusion that our psychological nows are shared; there is just one 'socio-psychological' now that we both partake in.

If that is true, then it seems legitimate to take my current past to include not only times that once were psychological nows for me, but also those that once were psychological nows for you. And by the same token, not only your psychological nows but also the psychological nows of people that were around when you were born, and so on. Going down this path we can, with some slight of hand, take my current past to include times that were psychological nows of any agents capable of temporal awareness (i.e. those which entertain, at any time of their consciousness, a psychological now).

We can argue similarly with regard to my current future. You may outlive me. In that case some of your future psychological nows will not be future psychological nows of mine. But for the reasons above we may be consider them nevertheless part of my current future. And so on.

In this way my current past and my current future will be substantial extended. But will they now cover all of time there is? That depends. If we believe in a God, who has existed from the beginning of the world and will exist until the end of it, and who is *in* the world in the sense of having, at each point of its development a psychological now corresponding to that point, then the answer is 'yes': all times are psychological nows of God at some point, and the psychological nows of mortal creatures like us will cover minute stretches of that very long series of psychological nows of the divine. But if we assume instead that only humans are agents that are capable of temporal awareness, and furthermore that the human race hasn't been around for very long and that for all we know it will come its end long before the universe as a whole, then the answer is clearly 'no'. If this is your view, then there appears to be no way of grounding an ordering relation that extends to all times, including those preceding the advent of man and those following its demise, in the aspects of human psychology that we have been referring to. For times lying outside the range of collective human experience we have to rely on different information sources, those provided by the natural sciences and in particular of physics. (There is also a third possibility – that we share our universe with aliens and that aliens were around long before we were and will still be around after the human race has become extinct. Since so far at least we haven't made any direct contact with aliens, it isn't all that clear what sharing a psychological now with all or any aliens would come to. But in any case, even if the aliens cover between them a longer stretch of time than we do, it is not plausible that they should have been there from the very beginning or will be there to the very end. And of that is so, then roping the aliens into our socio-psychological account of the order of time won't really be a good enough, even if it helps some.)

1.2 Informal remarks about time and times; (b) the perspective of Natural Science

For Physics, and for Natural Science generally, time is an inalienable dimension of the phenomena with which it deals. For Physics many of the phenomena that it studies and for which it tries to find explanations are *changes* of some kind or another – things that happen with the result that after they have happened the world is in a different state than the one it was in before. It is probably to even say such a thing – to say what a change consists in – without referring to time, explicitly or implicitly. But that is just an indication of how fundamental time is to anything having to do with change, and this to those things that physics is about. 'Time' and 'change' are inseparable concepts; there can be no change without time, and many have argued that there can be no time without change.

Of particular importance among the changes that are studied in physics are *movements* – movements of some object o which starts at one place A and ends up at some other place B, after having followed some path that bridges the space between A and B. To understand movement space is as indispensable as time. (The description that was given of movement makes that plain enough.) In fact, you m,gut say that in physics movement is where time and space meet and get correlated with each other. But for now let us set space aside and just take it as established that time is one essential aspect of physical phenomena, and that it is therefore reasonable to expect that the study of physical phenomena will reveal all or most of the structural properties that should be attributed to time.

In particular, physical phenomena – changes – impose a certain order on time. They do so because often changes come in some natural, irreversible sequence. When change C1 changes the world from state S1 into state S2 and C2 is a change of a kind that leads from a state of the kind of S2 to a state of some third type of state but one that could not occur if the world were in a state opt the kind of S1, then that imposes an order on these two changes: C2 can occur after but not before C1, and if both changes did occur, then C2 must have happened after C1, since it could have happened only after C1 had changed the state S1, in which C2 would have been impossible, into the state S2 in which C2 was possible. Movements provide simple illustrations of this sort of situation. Movement of o from B to C is possible only when o can start at B, and for that o has to be *at* B. So if we consider a movement of o from B to C and a movement of o from A to B, then (barring any other movements in between) it is clear that the last-mentioned one must have happened before the first. For o first has to get to B bedford it can move on from there. Many physical phenomena involving change are like this: one can occur only after the other has created the conditions that make it occurrence possible. And in such cases the time of the other phenomenon must of necessity have come before the time of the first one.

These intrinsically physical determinants of temporal order are of course in agreement with the psychological based order of time: If one change C1 precedes a second change C2 in the sense of physical time and both changes are observed by you, then your observations will be ordered in your psychological time in the same way that the observed changes are ordered physically. You may, when observing C2, *remember* your observation of C1 and that complex experience – remembering that you observed C1 while observing C2 – is your psychological testimony to the fact that your observation of C1 must have been before your observation of C1. And since the time of your observing C1 must have been the same as the tike of C1 itself, and likewise for C2 and your observation of C2, you will conclude, on the basis of this experience, that the time of C1 is before the time of C2. And so the psychological and the physical criteria are in line.

There is a somewhat different way of looking at this last point. We ourselves are part of the physical world, we are among the 'physical phenomena', and thus in principle among those that physics could and, arguably, should study. (It is just that we are among the phenomena that physics hasn't had much of interest to say about (not at any rate so far).) But nevertheless, it might be held that your observation of C1 and your observation of C2 are physical events of sorts, or at a minimum that these observations are accompanied by physical changes, and that these physical changes could not have happened if there hadn't been the events C1 and C2 of which they were (or accompanied) the observations. And it could then perhaps also be shown that remembering your observation of C1 is the kind of phenomenon that could not happen before the observation of C1, in much the same way, and on similar physical grounds, as why the movement from B to C could not have happened before the movement from A to B.

If this kind of story could be made to stick, then our psychological criteria for temporal order would turn out to be just a special case of the physical criteria. But of course that would not really be surprising, for what we are speculating about is some kind of whole-sale reduction of the psychological to the physical – the reduction of 'mind' to 'body'. We have learned how exceedingly difficult it is to make such a story stick. Perhaps it cannot be made to stick, in which we case we still have our psychological indicators for the order of time and our physical ones and the best we hope and argue for is that they line up.

More could be said about the psychological and physical basis of temporal order and about the ways they dovetail. But we won't get much farther this way. Time has come to look at the notions of order and temporal order in a more formal way, and to see what the formal properties are that any structure that deserves to be described as an 'order' should have, what properties differentiate between different types of orders and which of those properties we should assume acre true off the order of time

2 Formal Properties of Time.

2.1 Topological Properties

- We assume that time is an ordered set of instants. That is, we represent time as a structure $< T, \prec >$, where T is the set of instants and $\prec$ is the ordering relation between them. (We will often refer to this relation as the *earlier-later relation.*)

  Here are three potential properties for relational structures like $< T, \prec >$.

  (1) a. For every $t, t' \in T$, if $t \prec t'$, then not $t' \prec t$. ('asymmetry')
  b. For every $t, t', t'' \in T$, if $t \prec t'$ and $t' \prec t''$, then $t \prec t''$. ('transitivity')
  c. For every $t, t' \in T$, if $t \neq t'$, then either $t \prec t'$ or $t' \prec t$. ('linearity')

  A structure $< T, \prec >$ with properties (1.a) and (1.b) is called a *partial order*. If $< T, \prec >$ has in addition property (1.c), then $< T, \prec >$ is called a *linear order*. (So linear orders are spcial kind of partial orders, according to this terminology.)

- For the moment we will focus on linearly ordered structures. For these we can define a number of further properties that are also relevant to the question we are pursuing right now: What is the structure of natural time?

  (2) a. for every $t, t' \in T$, if $t \prec t'$, then there is a $t'' \in T$, such that $t \prec t''$ and $t'' \prec t'$. ('density')

  b. (1) For every $t \in T$, if there is a time $t' \in T$ such that $t \prec t'$, then there is a 'first time t'' after t' – formally:

there is a $t'' \in T$ such that (i) $t \prec t''$, and (ii) for any $t' \in T$ such that $t \prec t'$, $t'' \preceq t'$.

(2) for every $t \in T$, if there is a time $t' \in T$ such that $t' \prec t$, then there is a 'last time t'' before t' – formally:
there is a $t'' \in T$ such that (i) $t'' \prec t$, and (ii) for any $t' \in T$ such that $t' \prec t$, $t' \succeq t''$.

('discreteness')

c. (1) For any non-empty $X \subseteq T$, 'if X has an upper bound in T, then X has a least upper bound in T'; formally:
if there is a $t \in T$ such that for all $t' \in X$, $t' \preceq t$, then there is a $t'' \in T$ such that (i) for all $t' \in X$, $t' \preceq t''$ and (ii) for any $t \in T$ such that for all $t' \in X$, $t' \preceq t$, $t \preceq t''$.

(2) For any non-empty $X \subseteq T$, 'if X has a lower bound in T, then X has a greatest lower bound in T'; formally:
if there is a $t \in T$ such that for all $t' \in X$, $t' \succeq t$, then there is a $t'' \in T$ such that (i) for all $t' \in X$, $t' \succeq t''$ and (ii) for any $t \in T$ such that for all $t' \in X$, $t' \succeq t$, $t \succeq t''$.

('continuity')

In addition to the question whether time might have any of these properties, there is the further question whether it has any of the following. ('Does or doesn't time have a beginning? (Is there a first instant of time?) Does or doesn't time have an end? (Is there a last instant of time?).

(3) a. There is a $t \in T$ such that for all $t' \in T$, if $t \preceq t'$. ('time has a beginning)

b. There is no $t \in T$ such that for all $t' \in T$, if $t \preceq t'$. ('time has no beginning)

c. There is a $t \in T$ such that for all $t' \in T$, if $t \succeq t'$. ('time has an end)

d. There is no $t \in T$ such that for all $t' \in T$, if $t \succeq t'$. ('time has no end)

2.2 The metric of time

All the potential properties of time that we have reviewed so far are *topological* properties: they can be defined just in terms of the ordering relation between instants. But as time is understood and handled in natural science, it not only has a topological structure but also *metrical* properties: Intervals of time can be assessed for their size - for how long thy last. (The 'duration', or 'length', of an interval of time can be an hour or a day or a year or a light year, or a second, or a nano-second and so on.)

In order to be able to assess the metrical properties of time we need to assume that it has a metric as well as an order. We think of the metric as a function that assigns real numbers to sets of instants. The sets that are primarily of interest in this connection are the *intervals* of the topological structure $< T, \prec >$ of time. The general notion here, which can be applied to arbitrary linear orders is that of a *convex subset* of T

(4) (Def. of *convex subset of* T)

Let $< T, \prec >$ be a linearly ordered structure. Then a *convex* subset of T is a subset C of T such that for all t, t' and t'' in T, if t, t' in C and $t \prec t'' \prec t'$, then t'' also belongs to C.

When $< T, \prec >$ is continuous, then the convex subsets C are all of one of the following nine possibilities:
(i) T $((-\infty, \infty))$;

(ii) There is a t in T such that C is the set of all t' in T such that $t \prec t'$ $((t, \infty))$;

(iii) There is a t in T such that C is the set of all t' in T such that $t \preceq t'$ $([t, \infty))$;

(iv) There is a t in T such that C is the set of all t' in T such that $t' \prec t$ $((-\infty, t))$;

(v) There is a t in T such that C is the set of all t' in T such that $t' \preceq t$ $((-\infty, t])$;

(vi) There are t and t'' in T such that C is the set of all t' in T such that $t \prec t' \prec t''$ $((t, t'))$;

(vii) There are t and t'' in T such that C is the set of all t' in T such that $t \preceq t' \prec t''$ $([t, t'))$;

(viii) There are t and t'' in T such that C is the set of all t' in T such that $t \prec t' \preceq t''$ $((t, t'])$;

(ix) There are t and t'' in T such that C is the set of all t' in T such that $t \preceq t' \prec t''$ $([t, t'])$.

For the present discussion the convex sets of T that matter are the sets of types (vi)-(ix). These are the so-called *bounded intervals* of T. According to the standard terminology: sets of type (vi) are called *open*, those of type (vii) are called *half closed, half open*, those of type (viii) are called *half open, half closed* and those of type (vii) are called *closed*.

For the present discussion we assume that time is continuous and we will focus on the bound intervals, for it is these that should be assessable for size.

That is, we want to investigate the properties of time structures $< T, \prec >, |\ |$, where $< T, \prec >$ is a topological structure and $|\ |$ is a *measure (function)* which assigns, at a minimum, positive real numbers to each of the bounded intervals of $< T, \prec >$.

Is it only to the bounded intervals that $|\ |$ should assign a 'duration'? No, there is at least one other type of subset of T for which we should expect $|\ |$ to return a value, viz. those subsets which consist of a single instant in T only. Instants are assumed to be durationless – 'instantanous' – ; so a set consisting of just one instant should have the duration of that instant, and that means: zero duration. So for any *singleton set* $\{t\}$, where t is an instant from T, we should demand that $|\ \{t\}\ | = 0$.

But now it seems that we are faced with a paradox: For every bound interval I, $|\ I\ |$ is to be a positive number. But I is a set of instants for each of those we have that its duration is 0. But if all th instants in I have duration 0, how can a set consisting just of such instants have a size greater than 0? Isn't 0 such that if we 'sum up' a collection of things all of size 0 the result will necessarily be 0?

It is a remarkable and somehow mysterious result of mathematics that functions $|\ |$ with these properties so exist: Even when $< T, \prec >$ is dense and continuous, and in particular if it is like the real line (the standard assumption mad in classical mathematical physics), there are functions $|\ |$ which assign positive numbers to all the bounded intervals and 0 to all the singleton sets. This is a deep and surprising result of modern mathematics, which for some people (me among them) never quite loses it sense of mystery. Part of what is behind the mystery is that bound intervals contain VERY MANY points, even when their end points t and t' lie very close together. They are *superdenumrable* (or *uncountable*) sets – sets that are not only infinite but of a higher degree of infinity than, for instance, the natural numbers 0,1,2,. . . . But that by itself doesn't explain the mystery away.

In fact, there are measure functions with these properties and which also have another property that we would want them to have, viz. *additivity*. An *additive measure* $|\ |$ is one with the property that if I and I' are subsets of T for both of which $|\ |$ is defined and I and I' have no point in common, then $|\ |$ will also be defined for the *union* of I and I' – i.e. the set I'' the points of which are just those of I together with those of I' – and $\mid I'' \mid = \mid I \mid + \mid I' \mid$. More formally:

(5) (Def. of *additive measure on* $< T, \prec >$)

A measure $\mid \mid$ on an ordered structure $< T, \prec >$ is *additive* iff for any two subsets A,B of T if $A \cap B = \emptyset$, and $\mid A \mid$ and $\mid B \mid$ are both defined, then $\mid A \cup B \mid = \mid A \mid$ and $\mid B \mid$.

To ask of $\mid \mid$ that it be defined for the unions of all sets for which it is defined is stronger than to ask that additivity holds for bounded intervals. But it is possible to show that measures of the kind we want can be additive in this strong sense.

Moreover, such measures have the following property: for any sets A, B such that $\mid A \mid$ is defined and A and B differ only by a finite number of points, $\mid A \mid = \mid B \mid$. This entails in particular that for any $t \prec t'$, $\mid [t, t'] \mid = \mid (t, t'] \mid = \mid [t, t') \mid = \mid (t, t') \mid$.

Exercise: Prove this!

Note that this is a natural demand on $\mid \mid$. Consider two intervals $(t, t']$ and $(t', t'']$. These have no point in common and their union is the interval $(t, t'']$. Clearly we want in this case that the duration of the union is the sum of the durations of its part intervals.

To go into the reasons why this is the case – to prove that there structures of time $< T, \prec, \mid \mid >$ where $< T, \prec >$ is dense and continuous and $\mid \mid$ is an additive measure with the mentioned properties – requires quit a bit of sophisticated mathematics and so that is something we will not do here. We shall just take this to capable of rigid proof.

What evidence does natural science provide for the structural properties of the topological and metrical structure of time?

The main evidence from natural science we will consider here comes from physics, and more specifically from the theory of motion. Since motion is motion of objects through space – motion along a spatial path the successive positions of which the object reaches at successive times – motions correlate positions along their paths with times during which the motion takes place. This suggests that there must be at least as many times as there are distinct positions along paths of motions and also that the structure of time mimicks the structure of spatial paths.

But it is not only the topological structure of time that we should assume mimicks the topological structure of spatial paths; there should also be a correlation be a correlation between their metrical structures. The reason for this is that according to physics since Galilei motion plays a central part in the way things move through space, and it is widely assumed in natural science that many aspects of the physical world have to do with how things move – from the very big things, such as Galaxies, their stars and their planets and their moons to the very small things such as molecules, their atoms and their subatomic parts.

In particular a key concept of modern physics is *uniform* motion – motion at constant velocity. But uniform motion mans that the moving object o covers equal distances in equal times. So to be able to verify that o moves at uniform velocity over period of time I we must be able to determine on the one hand the amounts of distance covered by o during different parts of its motion and on the other hand the amounts of time that it it takes o to cover those distances; and these measurements must be independent from each other.

So then, how do we measure spatial distances and temporal durations and how can we be sure that the methods we use are not arbitrary, but reveal authentic properties of the actual world? This is actually a deep question about physics, even if the answer may seem simple and straightforward and we tend to take it for granted and treat it as self-evident. The short answer is: We can measure spatial distances using rulers and we measure amounts of time by using clocks. But it is important to reflect on why these procedures are not arbitrary, why thy can be seen as revealing the 'true' lengths of spatial distances and temporal durations.

First, the measurement of distance in space. What is a 'ruler' and how do we use rulers? Simple answer: rulers are objects with 'rigid length': a ruler has two points Beg and End marked on it such that when Beg and End are aligned with two points A and B in space, then the distance between A and B can be taken to be the distance embodied by the ruler. One reason why this is not arbitrary is that when we carry out this procedure at some time t_0, then take the ruler away and then repeat the procedure at some later time t_1, and we have independent evidence that A and B have not moved in relation to each other, then we will get the same result in that it is possible once more to align Beg with A and End with B.

We can use such a ruler also to assess the length of distances between points A and B that are not equal to its own distance. First, w can use it to determine distances that are some integral multiple of its own length: We can determine the distance between A and B ask times its own length by 'dividing the distance between A and B into k equal parts': We mark points p_1, .., p_{k-1} on the straight path from A to B such that it is possible to align Beg and End with A and p_1 and possible to align Beg and End with p_1 and p_2 and .. and possible to align Beg and End with p_1 and B. Again, if we repeat this procedure and there has no relative movement between A and B in the meantime, then the result will be the same.

We can also assess distances less than that of our given ruler R by constructing other rulers that are shorter than R. For instance, we can construct a ruler R' that is half the length of R in that we can mark a point Half on R such that it is possible to align Beg and Hal with Beg' and End' (where Beg' and End' are the beginning and end points of R') and also possible to align Hal and End with Beg' and End'. And then we can construct a ruler R" that is half the length of R' and so on. (In fact, the rulers that we normally use are really sets of rulers in the sense just discussed, insofar as they have 'evenly placed' marks on them. Any two marks on a given ruler R constitute a ruler as defined above and the numbers on R indicate how the lengths of these different 'rulers are related to each other.

As described, these procedures apply only to the (shortest) distances between two point A and B, i.e. to assess the spatial distance between A and B along a straight path between them. But we can also measure the distance between A and B along any curved path between them, for instance by using 'measuring tapes'. What I mean by an ideal measuring tape is roughly speaking a non-elastic piece of string: it has end Big and End such the distance between them along the length of the path defined by the string between them is always the same. but at the same time the string is completely flexible, in the sense that it can be made to follow any path, straight or non-straight, that is at least as long as it is. Such measuring taps give the same results as the rulers discussed above when paths are straight, but their range of application os of course much wider.

What lends further strength to the claim that there procedures for assigning length is that rulers can be made out of all sorts of different materials. All that matters is that the material is rigid. That of course opens up a new set of questions: How can we tell whether a material is rigid or not. But here too physics and chemistry ('materials science') provide independent criteria.

This is only a tiny part of the story of measurement of space that natural science has to offer. Modern science has availed itself of an enormous range of different methods for measuring distance, and the methods used for measuring very large distances and those for measuring very short ones are very different from the ruler-based methods just described (and also from each other). That there exists such a variety different methods, and that they give consistent results whenever more than one of them can be applied. is a further testimony to the regularity of the physical world. But even the little we have said gives some indication of the reason why rulers get at an intrinsic feature of space – at an intrinsic spatial metric.

Temporal distance, is measured by clocks. But what is a clock? Here is on way of understanding the fundamental feature of the physical world that can be seen as the foundation of everything that deserves to be considered and used as a clock:

> The world contains a large number *mutually periodic types of events*

What is meant by *mutually periodic types of events* can be described as follows:

A *type of events*, or *event type* as we will also say, is what the word says: a property E of events that particular events e will either instantiate or not instantiate – e either will or will not be of type E. For example '(the event type E_{diur} of) the sun rising above the horizon at some particular point P on Earth' is an event type. It is also an example of a *periodical* event type in that it has an 'unending' discrete sequence of instances: individual sun risings at P, one for each day in the history of our earth. Of course the sequence won't be literally unending. For if we believe physical theory on this point, the number of days – of full rotations of the Earth – is finite (and there are no grounds for thinking that physics might be wrong on this point). But so long as the earth keeps rotating E_{diur} can be used as a clock (if one that is a fairly rough and not fully correct).

What justifies us in regarding and using E_{diur} as a clock is that it is mutually periodical with an open-ended range of other periodical event types. We say that two event types E_1 and E_2 are *mutually periodical* iff, roughly speaking, for a given number n of successive instances of E_1 there will always be roughly the same number k_n of successive instances of E_2 occurring in the period covered by the n instances of E_1, irrespective of which n successive instances of E_1 we consider.

More precisely, there is a number r – the 'ratio' between the number of successive occurrences of E_1 and the number of successive occurrences of E_2 – such that for any interval of time $I = [t-1, t-2]$ if there are n successive occurrences of E_1 within I, then there are between [r.n]-1 and [r.n]+1 successive occurrences of E_1 within I. (Here [r,n] is the greatest whole number *leq* r.n.)

This correlation between E_1 and E_2 becomes the more significant as I gets bigger. For 'long' I, i.e. for I such that the number n is large, the difference between [r.n]-1 and [r.n]+1 will be small in comparison with the 'amounts of time' – n and some number between [r.n]-1 and [r.n]+1 – that the 'clocks' E_1 and E_2 assign to I.

The use of such clocks in measuring amounts of passing time is of course something that is overly familiar to us all. But let us try to be explicit nevertheless. Suppose that we have an independent way of identifying the beginning point t_0 and end point t_1 of some temporal interval I; for instance, let I be the duration of the movement of object o from position A to position B and that t_0 can be identified as the time when o leaves A and t_1 as the time when o reaches B. We can use one of our clocks E to get some approximation of the duration of I by counting how many instances it has within the interval from t_0 to t_1. When E is a 'slow ticking' clock relative to I then the result will not tell us very much about the duration of I. (Perhaps there is no complete instance of E within I, so that the count is 0; but of course that doesn't mean that the duration of I *is* 0. But the result of the procedure become more revealing as the clock E we use ticks faster, i.. as the number of instances of E within I becomes larger. Suppose that the number of successive instances of E within I is some large number N. Then any other interval I' within which there are also N instances of E will in fact be close in actual duration to I. More precisely (and assuming that E is completely accurate): if there as N successive instances of E within I and also in that N successive instances of E within I', then the actual durations of I and I' can differ from each other only by a factor 2/N. So the faster ticking our clock E, the better the approximation w get when we use it to measure the duration of some interval I.

To make this a little more concrete, suppose that our system of mutually periodical event types contains two types E_{min} and E_{sec} such that E_{sec} 'ticks' 60 times as fast as E_{min} – that is, within the period of time between the beginnings of two successive instances of E_{min} there are always between 58 and 60 instances of E_{sec}. Then, when we measure our given intervals I and I' with both E_{min} and E_{sec} the results will be that we get, for either of them, roughly 60 times as many counts from E_{sec} as we get from E_{min}. And when I and I' contain the same number N_{sec} of instances of E_{sec}, then their durations can differ by no more than a factor of $2/\mathrm{N}_{sec}$. Suppose that I and I' also contain the same number N_{min} of instances of E_{min}. Then on the basis of this count we can assert that their durations differ by no more than a factor of $2/\mathrm{N}_{min}$. But since N_{sec} is approximately $60.\mathrm{N}_{min}$, the results obtained with E_{sec} give us a more accurate assessment of how close I and I' are in actual duration than those obtained with E_{min}: According to the results obtained with E_{sec} the percentual difference between the two durations cannot be more than $2/\mathrm{N}_{sec}$, while the results obtained with E_{min} only tell us that the two durations differ by no more than a factor $2/\mathrm{N}_{min}$. Since approximately $\mathrm{N}_{sec} = 60.\mathrm{N}_{min}$, the approximation obtained with E_{sec} is roughly by a factor 60 better than that obtained with E_{min} ($2/\mathrm{N}_{sec} \approx (1/60).2/\mathrm{N}_{min}$).

All this is just a very elaborate and rather contorted way of saying that, for instance, measuring time with a clock that measures in seconds gives more accurate results than measuring with a clock that measures in minutes.

In general, each 'clock' in our system of mutually periodical event types measures time at its own rate, and thereby defines its own *unit of time*. The actual practice of time measurement has been to select a small number of such units – hours, minutes, seconds, micro-seconds, nanoseconds,.., years, lightyears, siderial years,.. – such that the rates between those units are precisely fixed and to 'calibrate' every clock we use in terms of on or more of these units. (The clocks and watches we use in our daily lives are calibrated in terms of the units 'hour', 'minute', 'second' and sometimes also have ways of indicating the (less accurate) units 'day' and 'month'. For scientific purposes, where the intervals we want to measure can be very short or very long, calibration will usually be in terms of other units.)

Just as we argued in connection with the measurement of distance in space, the evidence that our methods of measuring durations of temporal intervals by means of clocks get at the 'true' metric of time rests on the fact that the system of mutually periodical event types in our physical world is so large and so diverse. Moreover, the two systems – that for measuring spatial and that for measuring temporal distance – are closely connected. One way in which they are connected is through the phenomena of motion. I already observed that the evidence for the structure of time that is provided by physical phenomena is as much evidence for the structure of space as it is evidence for the structure of time and that the salient phenomena that establish the connection are those that involve physical motion. In particular, uniform motion plays a central role kinematics - that part of physics which deals with the motions of bodies with and without forces acting on them. And as we noted, in order to make concrete sense of uniform as opposed to non-uniform motion – in order to be able to verify that certain motions *are* in fact uniform – we must be able to measuring spatial and temporal distances reliably and independently. That is what our systems based on rulers and clocks deliver.

But once we have been able to convince ourselves, by applying our methods for measuring spatial and temporal distances to certain kinds of motions, that these motions are uniform, we can then use this result in its turn to use devices that initially are just to measure one of the two kinds of distances as a way of measuring the other. (For instance, in radar the distance of the object that a radar station detects, when it registers the reflection off the object (the 'echo' returned by it) can be determined by measuring the amount of time that passes between the sending of the beam and the reception of its echo: the nearer the object, the shorter the time that the beam has to travel on its way to and return from the object and so the shorter therefore the amount of time between the time of sending and the time of reception.) In such ways the devices we use for measuring space or time may themselves be based on our confidence in physical laws for which we gathered evidence by other measuring devices and methods. This is how and why physical measurement and physical theory are in general so intimately interwoven.

And it isn't just that the system of mutually periodical event types is so large and diverse; it is also that there is just one such system – one system consisting of a large number of different possible clocks from which we can choose and any two of them will yield consistent results (in terms of their respective units which stand in a constant rate to each other). It is the results from this large system of alternative clocks that present themselves inescapably as approximations to the 'true' durations of the intervals we measure in this way. Had there been two competing systems of mutually periodical event types, then that would have suggested that there were two distinct systems of physical regularities regulating our world, without a simple correlation between them. It is hard to imagine what a world governed by two such systems of regularities, with no apparent law-like correlations to connect them, could possibly be like. (Just try to imagine a world of some such sort!)

According to these considerations, the evidence from natural science for the structure of time cannot be separated from the evidence for the structure of space. What we are getting is a package deal: If we take any of this seriously as evidence for the structure of time, we better also take it as evidence for the structure of space, and conversely. This is the conclusion that physicists from the sixteenth and seventeenth century saw themselves confronted with and the way they drew it has been fundamental for the way in which physics and many other branches of natural science have been conducted ever since.

The form the conclusion took was due to one of the most important breakthroughs in the history of mathematics, vs. the development of the differential and integral calculus – roughly that which today is referred to in English as 'The Calculus'. Differential and integral calculus is essential to the theories of motion in the presence or absence of force that Newton and Leibniz were formulating and that they developed this cluster of mathematical tools for. And what these tools were being applied to in those intended applications were, among others, functions that describe the amount of spatial progress an object is making along a path as a function of the amount of time it has been travelling. In other words, this kind of function, that maps temporal distances as arguments to spatial distances as values, is the kind that the differential and integral calculus must be applicable to. So the domains from which their arguments and values are drawn – that is the structures of points along spatial paths and the instants constitute the structure of time – must have the kind of structure that application of the Calculus presupposes.

So that, then was what space an time were assumed to be like: their structure must such that the Calculus is applicable to them. Whether that means quite the same thing for space as it does for time may not be clear right away, for time has just on dimension, whereas space is a more complex structure with all sorts of one-dimensional 'subspaces' (Viz. the paths that we have been talking about so far). But it is not too hard to show that when we assume that space is 3-dimensional, in the sense that each of its points can be identified by *by three coordinates* – numbers that give its distance to some given fixed point (the *origin*) along three different directions, or *axes* – then that will give us just what we need.

So this then seemed to give the kinematicists of the seventeenth century all they needed: one suitably structured order for time and a compositum of three similarly structured orders for space. The structures in question were assumed, moreover, to be ('isomorphic') copies of the 'real line', the structure of the 'real numbers' with their intrinsic 'smaller than' relation as order and with the metric given by numerical difference: $|(\mathrm{r}_1,\mathrm{r}_2)| = \mathrm{r}_2 - \mathrm{r}_1$, where '-' is the subtraction operation on real numbers. In short, and using '$\mathcal{R}$' to refer to the real line, the structure of time is that of '$\mathcal{R}$' and the structure of space is that of '$\mathcal{R}^3$', the structure made up of all points $< r_1, r_2, r_3 >$, where r_1, r_2 and r_3 are points on three orthogonal copies of the real line.

That this is what time and space are supposed to be is perhaps clearest in the writings of Newton, who assumed that a space-time structure of this kind formed the God-given receptacle within which all physical events actually take place. (Leibniz's position was somewhat different. We will come to this below.) So, even if all or some of the events that make up the history of this world are *contingent* – they might have been different, or not happened at all –, the space-time within they unfold would still have been the same.

So far so good. But what really *is* the real line? It might be thought that if there is something that mathematicians knew at last once the Calculus had been properly developed and it is quite possible that the great mathematicians of the seventeenth and eighteenth centuries who took an active part in this development thought so, or would have given such an answer had anyone asked them. But in hindsight it seems possible that at that time no one was really in a position to understand the various implications of that question in the way that we understand them now.

We owe our current understanding of those implications to those 19-th century mathematicians – some of the names to mention are *Cantor*, *Cauchy*, *Dedekind* and *Weierstrass* – who recognized the need to put all of mathematics, but most particularly the Calculus, on a more secure conceptual and formal footing. Such a thorough revision of the foundations of mathematics had become necessary because of various paradoxes that had emerged and that seemed to threaten the entire edifice of mathematics with inconsistency.

It was only through those investigations that it became fully clear what is involved in 'building' a structure like the real line from a solid, contradiction-free basis. And this is the foundation that one came up with. The basic structure is formed by the *natural numbers*: the numbers 0,1,2,3,.., which can be obtained by starting with the number 0 and then getting the other numbers by adding 1, repeating this procedure indefinitely. From the natural numbers we can obtain the integers, i.e. the natural numbers together with the negative whole numbers, by making 'negative copies' of the positive natural numbers (as we all do, distinguishing the negative copies from their positive originals by putting a '-' in front of them).

We can then extend the structure of the integers to that of the *rational* numbers - all those numbers that can be written as fractions k/n, where k is an integer and n is a positive natural number. The structure thus obtained is beginning to look like the real line, but we are really still quite far removed from where we want to be. The rationals have all sorts of gaps – not only at places that seem fit to receive irrational numbers such as $\sqrt{2}$, $\sqrt[3]{5}$ and so forth. (The gap within the rational number structure where $\sqrt{2}$ fits is between the set of all rational numbers r such that $r^2 < 2$ and the set of all rational numbers r' such that $r^2 > 2$; and likewise for other irrational numbers of this this sort.) But filling all these gaps won't be nearly enough;

there are many more gaps within the structure of the rationals than that. In fact, there are - and here we hit on another of those paradoxical properties that one has to be prepared for when dealing with infinite collections – many more gaps within the ordering of the rational numbers than there are rational numbers. The only way to close all those gaps is in one fell swoop, and not one at a time (starting, say, with the gap of $\sqrt{2}$, then the gap of $\sqrt{3}$ and plodding on in this vein). That there is a way of closing all the gaps within the rationals in one go is not self-evident, but we have as much evidence that this can be done without running into new paradoxes as we can hope to get.

It is by way of such a simultaneous-gap-filling operation that we get from the rational numbers to the real numbers, and with that to a structure that is suitable for the Calculus since it has the property that limits of bounded sequences always exist. In fact, it is the smallest such structure that contains the integers, and so a natural candidate for each of the four dimensions thzat make up time and space according to the conception of Newton and his fellow scientists from the 17th and 18th century. So, it might be said in hindsight that it was the topological and metric properties of which we now know that they are the properties of real line that Newton and Co. were assuming time and space to have. But some of the distinctions that it has become standard practice within mathematics to distinguish between, and that we also defined earlier in these notes – in particular, the distinction between density and continuity.

In short, then, it seems not unreasonable to say that as fat as the topology of time and space is concerned the evidence from physics points to the topological structure of the reals, and that it has been taken to do that since the seventeenth century even if at that time the different properties that we now know the reals to possess could not be clearly articulated.

As we have seen, the time and space of mathematical physics have not only topological but also certain metric properties. These too are fully determined once it is assumed that time and space are like the real numbers in the way explained above. The metric of the real line is given by the way that it is built from the natural numbers: The integers run through the reals as a kind of discrete spine with a very simple metric: all distances between any two successive integers (15 and 16, -273 and -272 and so on) are the same. And from that all other metrical properties follow: if we take the distance between two successive integers as our unit, then the distance between two integers n and m (with n < m) is simply m-n; the distance between rational numbers n/m and k/l with n/m < k/l is equal to (k.m - n.l)/m.l and the distance between two numbers at least one of which is irrational are also fixed by the way in which the reals are constructed out of the rationals. (We cannot go into the details of that here.) In short, and repeating a point already made, for any two reals r_1 and r_2 such that $r_1 < r_2$, the distance between them is the arithmetical difference r_1 - r_2.

One consequence of this is that time is infinite both ways in both a topological and a metrical sense. The topological sense was defined earlier, in (3.d): there is neither a first nor a last point in time. But this topological sense of time being infinite both ways – before any time, no matter how early, there is a time even earlier than it; and for any time no matter how late there is an even later time. But besides this topological sense there is also a metrical sense in which the reals are infinite in both directions. It is this: take any unit of temporal measurement u, no matter how big, and take any point in time t no matter how far back. Then there will always be a point t' before t at a distance u from t; likewise for any t, no matter how far in the future, there will always be a t' in the future of t at a distance u from it.

It should be clear that topological infinity and metrical infinity are not the same. The open interval (-1,1) of the real line, for instance, is infinite in the topological sense – being topologically infinite and being open really com down to the same thing – but it is obviously not metrically infinite.

While the real line is metrically infinite both ways it also has a property that keeps, you might say, this infinity within certain bounds. This property is called – in honour of the Greek mathematician Archimedes – the 'Archimedean' property. This is its definition:

(6) (Def. of the *archimedean property* of ordered structures with a metric)

Let $\mathcal{T} = < T, \prec, | \, | >$ be an ordered structure with a metric $| \, |$. Then for any two elements t_1, t_2 of T (no matter how far apart) and any unit of measurement u (no matter how small) it is possible to bridge the distance between t_1 and t_2 by a finite number of steps of size u.

That the real line is archimedean in this sense is fairly obvious. Take any two numbers r_1 and r_2 that are as far apart as you like and any 'unit' u –

that is: any positive real number u – no matter how small. Then you can get from r_1 to r_2 in no more than k steps of size u, where k is an integer > (r_2 - r_1)/u. When r_2 - r_1 is big and u, then k may be very big. But it will always be a finite number. (To put the same point in more intuitive terms: No matter how long an interval of time I you take and no matter what clock C you choose (no matter how fast C ticks), C can in principle measure the duration of I in that it will perform a finite number of ticks during the time that is covered by t.)

In pursuing the question what natural science can tell us about the structure of time, we have so far said almost nothing about events. Inasmuch as events have entered into the discussion at all, it has been only as motion events; these, we argued, play a primary part in connecting tom with space and connecting topological with metrical properties. Since that was all we needed for the argument, we never considered the question what other events there might be and whether the totality of events displays any interesting structural properties, in a similar sense in which we have been exploring the structural properties of the totality of times.

As a matter of fact the general concept of event doesn't play a very important part in classical physics, not at any rate under this name. But that is no longer true. In Einsteinian physics events do play a crucial part, at least at the level of an intuitive explanation of its central concepts and ideas. And the events that play a key role in this explanation are not motions (although these are also central to the theory, as they are to any essentially kinematic theory), but quasi-instantaneous events in one particular place, such as the emission of a light pulse or the disintegration of a radio-active nucleus or the absorption of an incoming photon or other particle. It is of such events that the Theory of Relativity says that if two of them happen at a great spatial distance from each other and close in time, then there will be no non-arbitrary answer to the question which of the two came before the other or whether they happened at exactly the same time.

(This aspect of the Theory of Relativity entails that time and space cannot be separated from each other in the way that seems self-evident to most of us and as is assumed in classical physics (see above), but that time and space form a single four-dimensional structure which dos not permit this simple kind of unscrambling into a temporal and there spatial components. What

Relativity Theory can tell us about such metaphysically fundamental categories as time, space and causation is an important and fascinating question, but it is not one we will pursue here.)

But are there events apart from the two kinds we have now identified – the motions and the local changes in the structure or constitution of matter? Or is anything one might want to consider an event either one of these two kinds or else some complex made up from component events that belong to these kinds? What natural science might have to say on this question is another matter we won't pursue here. But I mentioned it because it is connected with a way of thinking about the structure of time and space that is important not only in the context of the present discussion, but also for the remainder of the class, which will be concerned with time as it is conceived and expressed by the speakers of natural languages. Although there is much in common between the contributions that Newton and Leibniz made to the theory of motion and its mathematical foundations, they differed quite radically in their understanding of the nature of space and time. For Newton, as we saw, space and time are ike a receptacle within which the events of our world unfold but which has its identity and form independently of what happens within it. In contrast, Leibniz thought that space and time were nothing over and above the events that happen within them. Their structure may therefore depend on what happens within them: they are a reflexion of the structure of the totality of actual events.

A full explication of this idea became available only with Einsteim's General Theory of Relativity. This is also something we cannot go into here. (I am not the right person to tell you about this, and eve if I was it would take up far more time than could be justified in a course whose central topic are times and events in language.) But there is a much simpler story about how the temporal relations between events can be assumed to determine the topological structure of time. For us this is a story worth telling because it will become important also in connection with conceptual time and with time in language.

It is a story that can be told without going into details about what kinds of events there are. All we need to assume is that events are temporally ordered in the following sense: For any two events e_1 and e_2 either (i) e_1 wholly precedes e_2 in time ($e_1 \prec_{ev} e_2$) or (ii) e_1 and e_2 temporally overlap ($e_1 O_{ev} e_2$)

or (iii) e_2 wholly precedes e_1 ($e_2 \prec_{ev} e_1$). In addition, we assume that the relations $\prec_{ev}$ and O have the following intuitively plausible properties, stated in (7).1-6)

(7) (Properties of the relations $\prec_{ev}$ and O_{ev})

1. For any two events e_1 and e_2, if $e_1 \prec_{ev} e_2$, then not $e_2 \prec_{ev} e_1$ (asymmetry of $\prec_{ev}$);

2. For any three events e_1, e_2 and e_3, if $e_1 \prec_{ev} e_2$ and $e_2 \prec_{ev} e_3$, then $e_1 \prec_{ev} e_3$ (transitivity of $\prec_{ev}$);

3. For any two events e_1 and e_2, if $e_1 O_{ev} e_2$, then $e_2 O_{ev} e_1$ (symmetry of O_{ev}):

4. For any event e, $eO_{ev}e$ (reflexity of O_{ev});

5. For any two events e_1 and e_2, if $e_1 \prec_{ev} e_2$, then not $e_1 O_{ev} e_2$ ($\prec_{ev}$ and O_{ev} are mutually exclusive);

6. For any four events e_1, e_2, e_3 and e_4, if $e_1 \prec_{ev} e_2$, $e_2 O_{ev} e_3$ and $e_3 \prec_{ev} e_4$, then $e_1 \prec_{ev} e_4$ (transitivity of $\prec_{ev}$ and O_{ev}).

7. For any two events e_1 and e_2, either $e_1 \prec_{ev} e_2$ or $e_1 O_{ev} e_2$ or $e_2 \prec_{ev} e_1$ (exhaustivity of $\prec_{ev}$ and O_{ev}):

The only one of the postulates (7.1-6) that may not seem obvious at first sight is (7.6). To see the plausibility of this try to draw a picture that would disprove it: a picture of four events e_1, e_2, e_3 and e_4 standing in the relations $e_1 \prec_{ev} e_2 O_{ev} e_3 \prec_{ev} e_4$, but so that it is not the case that $e_1 \prec_{ev} e_4$. You'll see that it just isn't possible to do that.

The one postulate that may not seem obvious even upon reflection is the exhaustivity principle that was mentioned before (7) and repeated as (7.7). We will have more to say about this principle below.

Suppose that E is a collection of events satisfying the postulates in (7). Then we can construct from E a linearly ordered structure $< T_E, \prec_E >$ of 'instants of time' and a relation 'is going on at' between E and T_E such that for each e from E the set of t such that e happens at t is a convex subset of T_E.

To show this we proceed as follows. We construct the members of T_E as *maximal subsets of E of pairwise overlapping events*. More formally, a *set of pairwise overlapping events from E* is any subset X of E such that for any two events e_1 and e_2 from X, $e_1 O_{ev} e_2$. Furthermore, X is a *maximal set of pairwise overlapping events in E* if

(i) X is set of pairwise overlapping events in E, and

(ii) there is no subset Y of E of pairwise overlapping events that is bigger than X (i.e. for any event e in E that is not in X there is some event e' in X such that not $eO_{ev}e'$).

Let T_E be the set of all maximal subsets of E of pairwise overlapping events, and let $\prec_E$ be the relation that holds between members t_1 and t_2 of T_E iff there exist e_1 in t_1 and e_2 in t_2 such that $e_1 \prec_{ev} e_2$. Then we can show that $< T_E, \prec_E >$ is a linear order in the sense that it satisfies the postulates in (1). (Showing this is left as an exercise.)

Next we have to define the relation 'e is going on at t', for e belonging to E and t belonging to T_E. To see how this relation should be defined, not that the intuition behind the definition of the members of T_E as maximal sets of pairwise overlapping events is that such sets can be thought of as instants of time at which all the events in the set are going on. If this is right, then the relation 'e is going on at t', which we will denote as 'AT', can be simply defined as:

e is going on at t iff e belongs to t.

To show that for any e in E the set T_e of points t in T_E such that $eATt$ is a convex set, we need to show that if $eATt_1$ and $eATt_2$, and $t_1 \prec t_3 \prec t_2$, then $eATt_3$. Assume that $eATt_1$, $eATt_2$ and $t_1 \prec t_3 \prec t_2$. Suppose that it is not the case that $eATt_3$, in other words, that it is not the case that e belongs to t_3. Then, since t_3 is a maximal set of pairwise overlapping events, there must be an e' in t_3 such that not $eO_{ev}e'$. So by (7.7), either $e \prec_{ev} e'$ or $e' \prec_{ev} e$. Suppose that $e \prec_{ev} e'$. Since by assumption $\prec t_3 \prec t_2$, there are e_3 in t_3

and e_2 in t_2 such that $e_3 \prec_{ev} e_2$. Since e' and e_3 both belong to t_3, it must be the case that $e'O_{ev}e_3$. So we have: $e \prec_{ev} e'O_{ev}e_3 \prec_{ev} e_2$. Therefore, by (7.6), $e \prec_{ev} e_2$. But e and e_2 both belong to t_2, so $eO_{ev}e_2$. But by (7.5) this contradicts the supposition that $e \prec_{ev} e'$. So this supposition must be false. In the same way we can infer that the supposition that $e' \prec_{ev} e$ is false. But if both of these are false, then the assumption that $eO_{ev}e'$ must be true. So it cannot be that e does not belong to $\prec t_3$. q.e.d.

. This is one way in which we can construct for a given set E of events that satisfies the conditions in (7) a linearly ordered structure of 'temporal instants' such that each of the events e in E occupies a convex set of instants. But is there anything more we can say about the properties of this instant structure? The answer is: not unless we know more about the properties of E. When we focus on the properties we have defined earlier – density, discreteness, continuity, being with or without a beginning or an endpoint; see (2) and (3) – we see that we can without too much trouble translate what it means for the instant structure to have any of these properties back into corresponding conditions on E. But it seems hard to form any intuitions about whether E should be thought of as meeting any of these conditions, so long as our intuitions about what the events are that make up our world. And even if our views on this point were more definite, it would probably still be hard to find evidence that any of these additional properties are fulfilled. Presumably this would be so in particular for the properties of density and continuity.

For someone with a Leibnizian conception of the relation between times and events this presents a certain complication: if the totality of actual events isn't rich enough to guarantee that the instant structure it determines is dense and continuous, how then are we to apply the differential and integral calculus to the explanation of kinematic phenomena? One possible reply would be that it should be possible to embed the instant structure obtained from the totality of events in a formal structure of time and space which has the properties needed to apply the Calculus, and to then apply the results of its application to the instants of the embedded structure. (N.B. There have been attempts to extend the method of building instant structures from event structures described above to the construction of structures of spatial points on the basis of spatial relations between events. This proves much trickier and at least until fairly recently the proposals were rather sketchy. See e.g.

A.N Whitehead 'Process and Reality'). But ideally it is this that the Leibnizian program requires, not just the construction of a structure of temporal instants.)

This is where we leave the matter of the structure of time as part of the natural world and the question what natural science can offer by way of justification of the structure of that time. We now turn to the central topic of this class: time in thought and time in language.

3 Time in thought and language

3.1 Past, present and future

The most dramatic difference between time as we think of it as part of our lives and time as it is used and elucidated in the natural sciences is that we think of time as divided into past, present and future. That difference is all important to us, for we think of that which is in the past as immutable, as the things that nobody can do anything about and that we have to live with, whether we like it or not. At best we can derive pleasure or satisfaction from the things that went well for us, or that we feel we did right, and learn from what went wrong. But what we learn things for is the future - we learn so that we can make the right decisions and take the right actions, shaping, or helping to shape, what the future is going to be like.

This difference, between a future towards which we are going and whose form and content we can push in one direction or another, and a past of things that are given and beyond intervention, couldn't have been more dramatic and more deeply interwoven with who we think we are and what is the purpose of our being here.

Between these two polar opposites of past and future there is the present. In one way the present is like the past in that what we see happen right now is beyond our control as soon as we can observe it. But on the other hand the present is where the future begins, where we see the first impact of the actions we undertake and of the calamities we have seen coming but that we are powerless to do anything against.

As against this, the present is also often seen as something quite different – as a category in its own right, rather than as the fleeting transition between two others. On this view the present consists of all that is directly accessible to us (or would be if we happened to be in the right place and position): only what is (in the) present is what we can directly perceive. In fact, some philosophers have gone so far as to suggest that only the present is 'real' – everything else is just the content (actual or potential) of our memories or of our projects, plans and expectations. On this vow we live in an ever changing reality which in our minds may be colored by the shadows of things past and the foreshadows of things future.

When we attend more closely, however, at what is going on in perception, then we see that this picture of the present, as consisting of all and only what is real, is quite problematic. For perception, the causal process that reaches from the things or happenings perceived to the content of perceptual awareness – to the thought 'I am seeing/hearing/feeling/smelling/tasting such and such' –, takes its time however brief and by the end of it, when the observer thinks his 'I am seeing/hearing/feeling/smelling/tasting such and such', the such and such is past already. So these two views of the present, as fleeting transition from future to past and as the one real ground between the equally shadowy realms of past and future, are hard to cleanly separate from each other. This curious ambivalence of the present manifests itself in how we speak of things as currently present as opposed to things present or past. In some languages – English is one of them – this difference, we shall see, is manifest in the particular forms of speech we must use.

It isn't all that hard perhaps to see how these two notions of time – the scientific and the conceptual – can be reconciled. We ourselves are complex bundles of events, or, if you prefer, we are involved in complex bundles of events, that take place in the time of the natural world, and among those events that make up any one of us (or anyone's history) there are in particular the thoughts that we form and entertain at successive points of the time in and through which we live. One of the features of our thoughts is that they relate the events and states of affairs they are about to the times at which they themselves occur - the events or states we think about can stand in one of the three relations of complete precedence, overlap or complete succession, just as we assumed for any two events whatever when we discussed the possibility of constructing instants out of events, and in this instance those relations are almost always of special importance to us – in those ways in which the distinction between past, present and future is of crucial importance – so the temporal relations between thoughts and the events and stats they are about is almost always an integral and explicit part of the thought's content.

It is in this manner, then, that the distinction between past, present and future is built into the thoughts we have, and also, overtly or covertly, in the things we say. And because we are (or are involved) in successions of events and because many of those successive events are thoughts – anyone of us has thoughts at many different, successive times – the way in which we divide up the contents of our thoughts into past, present and future are forever shifting: Our memories are largely filled with what we remember we remember; we realize that some of the things we (remember we) were expecting are happening now, or should have been happening now, or have or should have come to pass. Our cognitive systems are on the whole remarkably effective in updating the temporal relations in which the contents of our thoughts stand to those thoughts themselves, as a function of how earlier thoughts persist in (natural) time, or are replaced by other thoughts.

. That is roughly how we human beings work and that is why the tripartite division of time into past, present and future is real in a psychological but not on a physical sense, and why the presents of human consciousnesses forever shift from earlier to later instants of physical time – because those consciousnesses are themselves successions of mental events happening in physical time and each of those mental events determines 'its' present, which is nothing other than the time at which it occurs.

The last thing to note about past, present and future is that it isn't just a distinction that each one of us makes and remakes on his own. At any time when we are part of an interacting group, our interactions are events in time, and as often as not those interactions don't take up much time. For instance, when you say something to me, then your act of producing your utterance and my act of understanding what you are saying are experienced by us as *simultaneous*: We both take your act to take place in what is at that point your *and* my psychological present. Likewise when we both hear a cry or an explosion. I take you to hear the sound at the same time as I do, and so do you. Likewise again when we are doing something together: I pour, you stir and this we make our sauce. I see how you stir what I am pouring into the mixture that is already in the bowl, you see how I keep pouring. Our actions are coordinated, we make an effort to coordinate them and we see them as coordinated – as happening our shared current present.

Of course, shared presents shift in physical time just as does the present of each separate individual consciousness. But the awareness that we share our presents (and thus also our pasts and our futures) with others lends the distinction an additional sense of (subjective) reality.

We may feel that arguing along such lines we can explain the distinction between past, present and future as a psychological quirk of consciousness, with its function of guiding us in making decisions that improve our chances in our future on the basis of information we have gathered in our past, and that in doing so we can explain the distinction away – as a distinction that is real only for the consciousnesses that entertain it. But that doesn't alter anything to the fact that the division into past, present and future is a very real feature of tim as we experience it and any theory of time as conceived by us should deal with it and accord it the central place that it deserves.

One attempt to do this is that of *Tense Logic*. Tense Logic was intended not only as an abstract way of accounting for the temporal dimension of the content of thought and of the role it plays in reasoning, but also as a tool for linguistic analysis, viz. as a tool for investigating the semantics of th tenses of the verb (hence the name 'Tense Logic') and other natural language expressions that have to do with temporality. As a tool for linguistic analysis Tense Logic soon proved to be quite inadequate and it is now no longer used in this capacity. It can still serve though, up to a point at least, as an abstract model of how the temporal dimension of our thoughts affects the ways in which we can use them in reasoning. And in any case, Tense Logic has been important historically important, both within a philosophical and a linguistic context. That alone justifies us in having a brief look at its main ideas and features.

If anyone can b called the father of Tense Logic it is no doubt the New Zealand philosopher and logician Arthur Prior. Prior's best known system of Tense Logic is that in which classical propositional logic (with, let us assume, the proposition letters p, q, r,.. and with the full complement of propositional connectives $\neg, \&, \vee, \rightarrow, \leftrightarrow$) is enriched with two 1-place propositional connectives P ('P' for 'Past'; '$P\phi$' can be read as 'it was the case that ϕ') and F ('F' for 'Future'; '$F\phi$' can be read as 'it will be the case that ϕ'). So among the formulas of this system we find for instance: p, q, $\neg q$, $p \,\&\, \neg q$, $p \rightarrow q$, $(p \rightarrow q) \rightarrow (\neg q \rightarrow \neg p)$, Pq, $F(q \,\&\, \neg p)$, PPq, $P\neg F(p \,\&\, q)$, $P(q \,\&\, \neg r \,\&\, Pr)$, and also, using '$H\phi$' as short for '$\neg P \neg \phi$' and '$G\phi$' as short for '$\neg F \neg \phi$', Hp, $H(p \rightarrow q)$, $GF(p \vee q)$, $(H(p \rightarrow q) \,\&\, Hp) \rightarrow Hq$; and so on.

(N.B. in virtue of what H and G are short for, the natural reading for '$H\phi$' is 'it has always been the case that ϕ', and for '$G\phi$' is 'it will always be the case that ϕ'.)

The intuition behind this so-called '(P,F)-calculus' is the following. In any application of the calculus the letters p, q, r,.. are supposed to stand for sentences, or, in other words, to denote the propositions expressed by those sentences. And the sentences that it most natural to think of as instantiations of the letters in such applications are present tense sentences like 'It is raining.', 'The sun is shining.', John loves Mary.'. 'Mary is in Paris.' and such like. Suppose for instance that p is taken to stand for 'It is raining.' and q for 'The sun is shining.' Then Fp can be read as 'It will be raining.', Pq as 'It was raining.' $G(q \rightarrow Fp)$ as 'whenever the sun will be shining, it will be raining at some later time', $\neg P(p \,\&\, q)$ as there never was a time when it was both raining and the sun was shining.' etc.

It should be fairly clear from these few examples that by combining the tens operators with the classical connectives $\neg, \&$, etc. we can build quit complicated propositions out the basic propositions p, q, r,.. . Some of these formulas, moreover, are true irrespective of what sentences their proposition letters are taken to stand for, or what the facts of the world happen to be like. An example is the formula $(H(p \rightarrow q) \,\&\, Hp) \rightarrow Hq$. It just follows from the meaning of the operator H and the conntctives $\&$ and $\rightarrow$ that what this formula says must be true no matter what propositions are denoted by p and q: If it has always been the case that p and it has also always been the case that if p then q, then at all times in the past it must have ben both be the case that p and that if p then q. But at any point at which it is both the case that p and if p then q it must necessarily also be the case that q. Since by assumption this applies to all times in the past, q must have been true at all times in the past, i.e. it must always have been the case that q.

But how can we tell in general which formulas are true no matter how their sentence letters are interpreted and no matter what the facts of the world are like? That is a non-trivial question. And equally non-trivial is the question which formulas logically follow from which ones. (Here by 'formula B logically follows from formulas A_1,.., A_n' we mean that irrespective of how the proposition letters in A_1,.., A_n, B are interpreted and irrespective of what the facts of the world, then, *if* A_1,.., A_n are true given such an interpretation and such facts, then so will be B.)

There are two main methods for answering these questions, the *proof-theoretic* and the *model-theoretic* method. The proof-theoretic method is much older than the model-theoretic method, and goes in its essence back to Aristotle (384-322 BC). The model-theoretic method dates from the middle of the last century, with some antecedents in the first half of the century. So for most of the history of logic the proof-theoretic method was th only one around.

The proof-theoretic method can be implemented in a number of different ways. One of these – and the only one we will say anything about here – consists in specifying a certain set of basic logical truths (called *logical axioms*) and one or more *inference rules* which permit the formal derivation of some formulas from certain others. Perhaps the best-known inference rule of all is the so-called rule of *Modus Ponens* or *Detachment*, which says that from formulas of the forms A and $A \to B$ one may derive the formula B. Intuitively, this rule is sound: if the premises A and $A \to B$ are both true, then B must be true as well; from true premises the rule can only lead to conclusions that are also true.

One way to capture the notions of a logically true formula and of a logically valid argument of classical propositional logic is to combine the rule of M(odus) P(onens) with a number of logical axioms. This system enables us to verify that B logically follows from A_1,.., A_n by deriving B from A_1,.., A_n and the logical axioms, using M.P. as inference rule. And to verify that B is a logically true formula it is to be derived from the logical axioms alone.

We don't go into the question how the axioms for classical propositional logic should be chosen. As it turns out, there are a number of compact and non-redundant choices, each of which gives in conjunction with the rule of M.P. a proof system for classical propositional logic that is both *sound* and *complete*. Intuitively this means that only and all logically true formulas and logically valid arguments are verified by the proof system. We will return to the notions of soundness and completeness below and define them more precisely.

For the (P,F)-calculus we can proceed in much the same way, and that is what Prior did. He added a number of axioms that essentially involve P and F and also two new rules. The axioms are given in (8) and the rules in (9).

(8) (Axioms for Linear Tense Logic)

a. $H(p \to q) \to (Hp \to Hq)$

b. $G(p \to q) \to (Gp \to Gq)$

c. $FHp \to p$

d. $PGp \to p$

e. $PPp \to Pp$

f. $FFp \to Fp$

g. $(Pp \; \& \; Pq) \to (P(p \; \& \; q) \vee P(p \; \& \; Pq) \vee P(q \; \& \; Pp))$

h. $(Fp \; \& \; Fq) \to (F(p \; \& \; q) \vee F(p \; \& \; Fq) \vee F(q \; \& \; Fp))$

The new rules that are needed differ from the rule of M.P. in that their application is more restricted: they can be applied only to formula that have already been established as logical truths, which in the context of the present system means: have already been derived from the logical axioms alone. (As logical terminology has it, they are *Rules of Proof*, whereas M.P. is an *Inference Rule*.)

(9) (Rules of Proof for Tense Logic)

a. if $\vdash H\phi$, then $\vdash \phi$

b. if $\vdash G\phi$, then $\vdash \phi$

Together with M.P. and a sound and complete set of axioms for classical propositional logic the axioms and rules in (8) and (9) form a proof system that is sound and complete for linear tense logic. But what exactly does this mean? In order that we can be more precise about this we must turn to the other method for characterizing logical truth and validity, viz. the model-theoretic method. The model-theoretic method is based on a very simple idea. Intuitively, we said, a formula is a logical truth iff it is true irrespective of how its proposition letters are interpreted and irrespective of what the facts are that determine the truth values of the interpreting propositions. Similarly, B follows logically from A_1,.., A_n iff irrespective of interpretation of the letters and the truth-determining facts when the premises A_1,.., A_n are true, so is B. Model theory makes this intuition precise by

(10) a. defining what the range of different possibilities is of interpreting the propositional letters as denoting certain propositions and of the facts that determine their truth values; this range is given by the *class of models* of the theory;

b. spelling out formally which formulas are true in which models. (The definition that spells this out is known as the *truth definition* of the given model theory.)

With (10.a,b) in place we can now formally define a formula as logically true iff it is true in all models of the class specified under (10.a); and B can now be defined as following logically from A_1,.., A_n iff for any model $\mathcal{M}$ in the class, of A_1,.., A_n are true in $\mathcal{M}$, then B is true in M.

We are already in a good position to apply this idea to the case of (linear order) Tense Logic. For we have already discussed the most important part of what our models should be like in this case: Intuitively it should be clear that they ought to consist of a linear ordered time structure together with a specification of what the propositions are that interpret the proposition letters and information about where those propositions are true. We can think of this latter information as given simply in the form of a specification, for each of the propositions in question, of all the instants of the time structure at which the proposition is true, and we can do that by specifying for each proposition the set consisting of just those instants. Thus a model for linear Tense Logic is a structure $< T, \prec, I >$, where $< T, \prec>$ is a linearly ordered instant structure and I is a function which assigns to each of the proposition letters $p, q, r, ..$ a subset of T (intuitively: the set of those instants of T at which the proposition denoted by the given letter is true).

There is one further complication. We said a moment ago that a formula ϕ is logically true if it is true in a model M of the class specified by the model theory. But what does it mean for a formula to be true in a model $< T, \prec, I >$? At which instant of T should ϕ be true? Which instant of T should count as the present for the purpose of our (P, F)-calculus, which is based on the notion that time is divided into past, present and future? In the light of what we have said above about the relation between physical and psychological time, the natural answer, it would seem, is this: Any instant in T could play the part of psychological present; for if we abstract from the fact that there have been (and presumably will be again) times at which there are no human beings to have tensed thoughts or make tensed statements, each time is potential temporal location for such thoughts and utterances.

This last answer leads us to our final proposal for the models for linear Tense Logic:

(11) A *model for linear Tense Logic* is a structure $< T, \prec, I, t_0 >$, where $< T, \prec, I >$ is as above and t_0 is a member of T.

To spell out what it is for a formula of the (P, F)-calculus to be true in such a model, however, we cannot restrict attention to what it is for formulas to be true at t_0. The reason for this is that truth definitions of the sort that are used in model theory are *recursive*; they characterize the truth values of complex formulas in terms of the truth values for their immediate constituents. Consider for instance the clauses (ii), (iii) and (vii) of the truth definition (12) below. (12.ii) deals with negated formulas (formulas of the form '$\neg\phi$'). It says that $\neg\phi$ is true at t iff ϕ is false at t, thus reducing the question of truth or falsity for $\neg\phi$ to that of truth or falsity for ϕ, and more specifically to ϕ's truth or falsity at t. (12.iii) deals with conjunctions (formulas of the form 'ϕ & ψ') and reduces the question of their truth or falsity at t to the question of truth or falsity of each of the conjuncts ϕ and ψ at t. (12.vii) gives such a reduction for formulas of the form '$P\phi$'. But there is one important difference with the clauses (ii) and (iii) (and the three other clauses (iv)-(vi) for the remaining classical connectives $\vee$, $\rightarrow$ and $\leftrightarrow$). It is this: the question of truth or falsity of $P\phi$ at t is reduced not to a question about ϕ at t but about questions about ϕ *at other times than* t. Because of this we need, even when our primary interest is in the truth values of formulas at t_0, also know about the truth values at other times than t_0. This need manifests itself as soon as tense operators operators (connectives or other tense operators) within their scope. For a very simple illustration, consider the formula $P(p \,\&\, q)$. According to (12.vii) this formula is true in $\mathcal{M}$ at t_0 if there is a time t' such that $t' \prec t_0$ and $p \,\&\, q$ is true in $\mathcal{M}$ at t'. But what is it for $p \,\&\, q$ to be true in $\mathcal{M}$ at t'? That is what we are supposed to be told by (12.iii). And indeed, (12.iii) does tell us this, but it does because it reduces the questions about the truth or falsity of $p \,\&\, q$ for *arbitrary* t to questions concerning p and q, and not just for the 'present' t_0 of $\mathcal{M}$.

(12) (Truth definition for the (P, F)-calculus)

Let $\mathcal{M} = < T, \prec, I, t_0 >$ be a model for the (P, F)-calculus and let t be an arbitrary instant of T. For any formula ϕ of the (P, F)-calculus we abbreviate 'ϕ is true at t in $\mathcal{M}$' as: '$[[\phi]]_{\mathcal{M},t} = 1$' and '$\phi$ is false at t in $\mathcal{M}$' as: '$[[\phi]]_{\mathcal{M},t} = 0$'. Then

(i) $[[p]]_{\mathcal{M},t} = 1$ iff t belongs to $I(p)$, and likewise for the other sentence letters;

(ii) $[[\neg\phi]]_{\mathcal{M},t} = 1$ iff $[[\phi]]_{\mathcal{M},t} = 0$;

(iii) $[[\phi \,\&\, \psi]]_{\mathcal{M},t} = 1$ iff $[[\phi]]_{\mathcal{M},t} = 1$ and $[[\psi]]_{\mathcal{M},t} = 1$;

(iv) $[[\phi \vee \psi]]_{\mathcal{M},t} = 1$ iff $[[\phi]]_{\mathcal{M},t} = 1$ or $[[\psi]]_{\mathcal{M},t} = 1$;

(v) $[[\phi \rightarrow \psi]]_{\mathcal{M},t} = 1$ iff $[[\phi]]_{\mathcal{M},t} = 0$ or $[[\psi]]_{\mathcal{M},t} = 1$;

(vi) $[[\phi \rightarrow \psi]]_{\mathcal{M},t} = 1$ iff $[[\phi]]_{\mathcal{M},t} = [[\psi]]_{\mathcal{M},t}$;

(vii) $[[P\phi]]_{\mathcal{M},t} = 1$ iff for some t' in T such that $t' \prec t$ $[[\phi]]_{\mathcal{M},t'} = 1$;

(viii) $[[F\phi]]_{\mathcal{M},t} = 1$ iff for some t' in T such that $t' \succ t$ $[[\phi]]_{\mathcal{M},t'} = 1$.

We now have a basis for saying more precisely what it is for a formula of the (P, F)-calculus to be logically true or for it to follow logically from one or more other formulas. But there still is one further distinction we need to make. First suppose that time has some fixed structure and that $< T_0, \prec_0 >$ is that structure. Then we can consider what it means for a formula to be a logical truth (or to logically follow form others) in virtue of time having precisely that structure. The set-up above suggests a natural answer to this question: For the given time structure $< T_0, \prec_0 >$ there are many different models of the form $< T_0, \prec_0, I, t_0 >$ with I being some assignment of subsets of T_0 to the proposition letters and t_0 some instant or other from T_0. It is natural to take a formula to be *logically true in virtue of the structure* $< T_0, \prec_0 >$ iff it is true at t_0 in each of these models $< T_0, \prec_0, I, t_0 >$. (And analogously for a formula logically following from others in virtue if the structure of $< T_0, \prec_0 >$.)

But there is also another line we can take. We can ask what it is for a formula to be logically true in virtue of just the fact, or assumption, that the time structure is linear – time's linearity matters but not any other properties that the real time structure might actually have. To capture this notion we should consider the class C_{lin} of all linear structures $< T, \prec >$ and for each of these the class of all models $< T, \prec, I, t_0 >$ that can be obtained by adding to a structure in $< T, \prec >$ in C_{lin} an assignment function I and a 'present' t_0. This gives us a large class of models – many structures $< T, \prec >$ in C_{lin} and for each of those many models $< T, \prec, I, t_0 >$. It is truth in all these models that distinguishes the formulas that are logically true in virtue of (just) time's linearity.

Evidently this notion of being logically true in virtue of time being linear (and the corresponding notion of following logically in virtue of time being linear) can be replicated for other possible properties of time. For instance logical truth in virtue of tim being linear and dense can be defined as truth in all models based on the class $C_{lin,den}$ consisting of all time structures that are linear and dense, and so forth.

It is an interesting fact about the (P, F)-calculus (and an indication of its expressive power) that it has formulas that are sensitive to the properties of time in the relevant sense: Such formulas will be logically true when time is assumed to have a certain property P but not if time is not assumed to have P. Or, equivalently put, formulas that are true in all models based on time structures that have P, but that are false in certain models based on time structures that do not have P. As a matter of fact we have already encountered two examples of such formulas: the formulas (8.g,h) are true in all models based on linearly ordered time structures, but are false in some models based on time structures that are non-linear. Some other examples are given in (13).

(13) a. $Pq \to PPq$

b. $Fq \to FFq$

c. $q \to FH(q \vee Fq)$

d. $q \to PG(q \vee Pq)$

e. $(F\neg q \ \&\ Gq) \to F(Gq \ \&\ \neg PGq)$

f. $(P\neg q \ \&\ Hq) \to P(Hq \ \&\ \neg FHq)$

(13.a,b) are true in all models based on time structures that are dense, but not necessarily in models based on non-dense structures; and the sam is true for (13.c,d) and the property of discreteness and for (13.e.f) and the property of continuity.

This is the right point at which to return to the notions of soundness and completeness of proof systems that we mentioned informally earlier. The formal notion of completeness as it is used in logic presupposes that we have a model theory for our logical formalism as well as a proof system. The model theory defines the class of those formulas that are logically true in its sense (and the set of pairs $< \{A_1, .., A_n\}, B >$ such that B logically follows from A_1,.., A_n) and it is in relation to these model-theoretic characterizations that the proof system can now be assessed: It is *complete* with respect to these model-theoretic characterizations is for all B and A_1,.., A_n such that B follows logically from A_1,.., A_n in this model-theoretic sense the proof system enables the derivation of B from A_1,.., A_n. (And, as a special case, the proof system enables a derivation of B just from its logical axioms alone whenever B is a logical truth in the model-theoretic sense.) Soundness can be characterized along similar lines: the proof system is *sound* with respect to a given model-theoretic characterization of logical truth and following logically iff whenever it enables a derivation of B from A_1,.., A_n, then B follows logically from A_1,.., A_n in the semantic sense. (And agsin as a special case: if the system enables a derivation of B from the logical axioms alone, then B is a logical truth in the model-theoretic sense.)

It is in this technical sense that the proof system specified in (8) and (9) is sound and complete for the notion of following logically (and that of being a logical truth) with respect to the property of being a linear ordering. Proving soundness and completeness results like this one is not all that easy and there can be no question of going through the demonstration of these or similar results here.

We observed that the formulas in (13) stand in special relations to certain properties of the time structure - the first two to density, the next two to discreteness and the last two to continuity. In fact, the relationship in which they stand to those properties is even more intimate than we could indicate so far. Suppose we add the formulas (13.a,b) to the axioms in (8). This addition to the proof system we had will give us a new proof system (whose axioms now include the ones we have just added) which is sound with respect to the assumption that time is (linear and) dense. (This is so because the added axioms are logical truths with respect to the density property and this feature will be preserved by all derivations in the new system.) But the new system is not only sound with respect to density but also complete: whenever B follows logically from A_1,.., A_n in virtue of the assumption that time is (linear and) dense, then the new system enables the derivation of B from A_1,.., A_n.

In addition, we already noted, the formulas (13.a,b) are not always true in models that are based on time structures which lack the density property. This means that relative to a notion of temporal structure that does not include density the new system will not be sound. (The new axioms are trivially derivable in this system since they can be inferred from themselves and thus – trivially– from the set of axioms. And they are not logical truths in the sense we are now considering.)

So all-in-all the correlation between formulas and time structure property is a very strong one: If the formulas are added to a proof system that is sound and complete with respect to the assumption that time is linear, then we get a system that *sound and complete* with respect to the assumption that time is dense in addition.

The same is true for the other formulas in (13). When they are added as axioms to the proof system specified in (8) and (9) the result is a new proof system that is sound and complete with respect to the corresponding property (discreteness in the case of (13.c,d) and continuity in the case of (13.e,f)).

3.2 Tense Logics with other Tense Operators

As we have just seen, Prior's (P, F)-calculus as considerable expressive power: it enables us to express the distinctions between some of the subtle distinctions between various types of linear orderings. But even so, there are quite number of things that it cannot express and among these there are quite a few that we would expect Tense Logic to be able to express if it is to live up to its purported function as a general tool for the investigation of the time-dependent aspects of meaning in natural language. Among these are in particular the following complex propositions involving the constituent propositions p and q:

(14) a. It has been the case that q ever since it was the case that p.

b. It will be the case that q until it is the case that p.

c. There was a time when it was the case that p and such that since then it has always been the case that q.

d. There will be a time when it is the case that p and such that until then it will always be the case that q.

Perhaps the English circumlocutions in (14.a,b) are not completely free from ambiguity, but we can make them precise through the paraphrases in (14.c,d). (I think it is fair to say that (14.a,b) can be used to truthfully describe situations that are also correctly described by (14.c,d), even if there are arguably also situations that can be described with (14.a) or (14.b) but not with (14.c) or (14.d).) But let us, so as to set aside possible ambiguities, focus on (14.c,d).

(14.c) and (14.d) can each be regarded as the English paraphrase of a Tense operator. These are more complex than the operators P and F considered so far, first in that they are *two*-place operators and not *one*-place operators – they form new formulas out of two constituent formulas (here p

and q) rather than just one – and second that their semantic specifications are more complex than those for P and F (see (15) below).

The operators paraphrased in (14.c,d) have come to be known in the literature as 'Since' and 'Until' (as suggested by their paraphrases in (14)) and are formally represented as 'S' and 'U'. (15) gives the truth definition clauses for formulas which have S or U for their principal (that is: outermost) operator, in the same format as we did in (12).

(15) a. $[[S(\phi, \psi)]]_{\mathcal{M},t} = 1$ iff there is a t' in T such that $t' \prec t$ and $[[\phi]]_{\mathcal{M},t'} = 1$ and for all t'' from T such that $t' \prec t'' \prec t$ $[[\psi]]_{\mathcal{M},t''} = 1$;

b. $[[U(\phi, \psi)]]_{\mathcal{M},t} = 1$ iff there is a t' in T such that $t' \succ t$ and $[[\phi]]_{\mathcal{M},t'} = 1$ and for all t'' from T such that $t' \succ t'' \succ t$ $[[\psi]]_{\mathcal{M},t''} = 1$;

The (S, U)-calculus, in which the operators S and U replace the operators P and F, proves to be very powerful – much more so than the (P, F)-calculus. In fact it is in a certain sense maximally expressive, at least when time is assumed to be continuous. In particular it is capable of expressing pretty much all temporal properties of and relations between propositions that we are likely to find are expressible in any natural languages.

Here we cannot go into the question in what sense the expressive power of the (S, U)-calculus is maximal, as that would detract is even more from our central topic, viz. tim and events in language. But the very fact that that would be a distraction is itself something that deserves a comment. It would be a distraction because the (S, U)-calculus is, like the (P, F)-calculus, quite unsuitable for the analysis of temporal reference in English and other human languages. It is that in spite of the fact that it can, in its fashion, express pretty much any topological temporal property or relation that we can expect to come across while studying the languages of the world. The problem is with the 'in its fashion'. The (P, F)-calculus can express in particular those temporal properties and relations that may be important to the working linguist; but it will often do this in a manner that deviates so much from the ways this is done in the languages we study and whose mods of expression we want to understand better that the modes of expression offered by the (S, U)-calculus are irrelevant at best and at worst highly misleading.

One of the most serious drawbacks of Tense Logic as a tool in the study of time in language and thought is that it has no way of talking about events. This is as true for the (S, U)-calculus as it is for the (P, F)-calculus (or for that matter for Tense Logics built on any other set of basic operators). As we proceed, it will become clearer how serious this deficiency is: An account of time in human language that cannot speak in detail about events is a lost cause from the outset.

(As a side remark: Tense Logic, and in particular the (S, U)-calculus, have found a use in Computer Science, first in systems for program verification and then, more recently, as integral parts to the specification languages used in chip design. Here the manner in which the (S, U)-calculus represents temporal properties and relations - that, in other words, which makes it so unsuitable as a tool in linguistics, has proved to be a distinct advantage. There is something ironic in this: A feature of the system that makes it unsuitable for the purpose for which it was developed turns out to be just what is wanted in an application that the designers never thought of.)

3.3 The algebraic way of doing and looking at things

Let us return to our models for Tense Logic. For the following discussion the 'present' t_0 of the model will not be relevant, so we drop that constituent and consider, once again, structures of the form $< T, \prec, I >$.

Recall that in a structure $< T, \prec, I >$ I assigns a subset of T to each proposition letter. When we discussed this matter earlier, we remained noncommittal as to what the exact relation is supposed to be between these subsets and the propositions denoted by the letters according to the given interpretation (or the propositions expressed by the sentences that the letters are assumed to stand for in the given interpretation): are the subsets just the times at which the given propositions are true, or is there nothing more to the propositions than just at what instants of T they are true and at what instants of T they are false? Questions of this sort are sometimes debated at length within philosophy, but as often as not such debates are barren. What we can affirm in connection with the issues at hand is that the only thing about propositions that will matter in what follows is at which instants they are true or false. So we may as well identify the propositions with the subsets of T consisting of those instants at which they are true – the instants from T that are not in the set being those at which the proposition is false – and that is what we will do.

Given $< T, \prec, I >$, then, the possible propositions, relative to $< T, \prec, I >$, are just the subsets of T, and what I can be said to do is 'activate' some of those propositions, as assignments to the sentence letters and therewith as propositions expressed by some of the formulas of our formalism (viz. those very sentence letters, which are the formalism's simplest formulas).

When we think of I's role in these terms, then we can think of the truth definition for the formalism as extending the activation – as extending the assignment (by I) of propositions to the sentence letters to an assignment of propositions to all formulas. More specifically, each of the 'recursive' clauses of the truth definition – those that deal with syntactically complex formulas, formed out of one or more smaller formulas through the application of a connective or operator; in the truth definition (12) these are clauses (12.ii-viii) – as telling us how to transform the proposition or propositions expressed by the immediate constituent(s) of the formula mentioned on the left into the proposition expressed by that formula. Let us have another look at some of the clauses of (12) from this somewhat different perspective. (For easier reading the relevant clauses are repeated below.) For ease of reference I will write '$[[\phi]]_{\mathcal{M}}$' for the proposition denoted by the formula ϕ (i.e. for the set of instants of T at which ϕ is true; $\mathcal{M}$ now is short for $< T, \prec, I >$).

We start once again with (12.ii). This clause tells us that the times t at which $\neg\phi$ is true are precisely those at which ϕ is false, or not true. The set of these instants is thus what is called the *complement*, relative to T, of the set $[[\phi]]_{\mathcal{M}}$, the proposition denoted by ϕ. The standard way of denoting the complement of a set A relative to a set D is as '$D \setminus A$'. So $[[\neg\phi]]_{\mathcal{M}} = T \setminus [[\phi]]_{\mathcal{M}}$.

In the same vein, (12.iii) tells us that the proposition denoted by ϕ & ψ is the set of times t that belong both to the proposition $[[\phi]]_{\mathcal{M}}$ denoted by ϕ and to the proposition $[[\psi]]_{\mathcal{M}}$ denoted by ψ; that is, proposition denoted by ϕ & ψ is the *intersection* of $[[\phi]]_{\mathcal{M}}$ and $[[\psi]]_{\mathcal{M}}$. Intersection is standardly denoted as $\cap$, so $[[\phi \ \& \ \psi]]_{\mathcal{M}} = [[\phi]]_{\mathcal{M}} \cap [[\psi]]_{\mathcal{M}}$. And (12.iv) tells us that proposition denoted by $\phi \vee \psi$ is the set of those t which belong either to

$[[\phi]]_{\mathcal{M}}$ or to $[[\psi]]_{\mathcal{M}}$ (or to both). This set is called the *union* of the sets $[[\phi]]_{\mathcal{M}}$ and $[[\psi]]_{\mathcal{M}}$ and is standardly represented as '$\cup$': $[[\phi \vee \psi]]_{\mathcal{M}} = [[\phi]]_{\mathcal{M}} \cup [[\psi]]_{\mathcal{M}}$.

(12.ii) $[[\neg\phi]]_{\mathcal{M},t} = 1$ iff $[[\phi]]_{\mathcal{M},t} = 0$;

(12.iii) $[[\phi \,\&\, \psi]]_{\mathcal{M},t} = 1$ iff $[[\phi]]_{\mathcal{M},t} = 1$ and $[[\psi]]_{\mathcal{M},t} = 1$;

(12.iv) $[[\phi \vee \psi]]_{\mathcal{M},t} = 1$ iff $[[\phi]]_{\mathcal{M},t} = 1$ or $[[\psi]]_{\mathcal{M},t} = 1$;

Let us pause for a moment at this point and see what this first part of the exploration of the clauses of (12) has shown us. Together the clauses (**??**) tell us that when we restrict ourselves to that part of our formalism in which the only connectives and operators are $\neg$, & and $\vee$ (in other words we only consider formulas in which there are no occurrences of $\rightarrow, \leftrightarrow, P$ or F), then the total set of propositions expressed by this set of formulas is the *closure* of the set of propositions assigned by I under the operations $\setminus$, $\cap$ and $\cup$: For each set A it contains it also contains its complement $T \setminus A$ relative to some fixed set T; and for any two sets A, B is contains it also contains their intersection $A \cap B$ and their union $A \cup B$.

Such a structure, consisting of a set of basic elements – in our example these are the propositions assigned by I – and closed under the application of a given family of operators – here the operators $T \setminus$, $\cap$ and $\cup$ – is often called an *algebra*. The case just considered is an example of a type of algebra that is of considerable importance for logic and semantics and that is the reason for dwelling on the way we are looking at the truth definition in (12). Algebra's of this kind are called *Boolean algebra's*.

The prime examples of Boolean Algebra's, it could be claimed, are structures whose elements are sets, with the operators $X \setminus$ (for some given set X), $\cap$ and $\cup$. The example we have just considered is one of this kind. But such set algebra' are by no means the only Boolean algebra's. In general a Boolean algebra is a structure $< X, \setminus, \cap, \cup >$, where $\setminus, \cap, \cup$ need not be the set-theoretic operations that we have just been using them to denote, but that must obey certain 'laws', which the corresponding set-theoretic operations do obey. One way of stating those laws is in the form of the axioms in (16) below. (They are given here for future reference. Don't bother with them now if you don't feel like it.)

(16) (Axioms for Boolean algebra's)

(AxBo.1)	(for all x in X)	$x \cap x = x$
(AxBo.2)		$x \cup x = x$
(AxBo.3)	(for all x, y in X)	$x \cap y = y \cap x$
(AxBo.4)		$x \cup y = y \cup x$
(AxBo.5)	(for all x, y, z in X)	$(x \cap y) \cap z = x \cap (y \cap z)$
(AxBo.6)		$(x \cup y) \cup z = x \cup (y \cup z)$
(AxBo.7)	(for all x, y in X)	$(x \cap y) \cup x = x$
(AxBo.8)		$(x \cup y) \cap x = x$
(AxBo.9)	(for all x, y, z in X)	$(x \cap y) \cup z = (x \cup z) \cap (y \cup z)$
(AxBo.10)		$(x \cup y) \cap z = (x \cap z) \cup (y \cap z)$
(AxBo.11)	(for all x, y in X)	$(x \cup \backslash x) \cap y = y$
(AxBo.12)		$(x \cap \backslash x) \cup y = y$

A *Boolean algebra*, then, is any structure $< X, \backslash, \cap, \cup >$ such that $\backslash$,$\cap$ and $\cup$ satisfy the conditions (AxBo.1-12).

N.B. The last two axioms, (AxBo.11) and (AxBo.12), are a bit artificial as they stand. What (AxBo.11) really says is that for any x the union of x and $\backslash x$ the 'biggest' element of the algebra, so that the intersection with it and any other element y will always be equal to the 'smaller' element y. (This means in particular that for any x and y the union of x and $\backslash x$ is the sam as the union of y and $\backslash y$. Show this!) Likewise, the force of (AxBo.12) is that for any x the intersection of x and $\backslash x$ is the smallest' element of the algebra, so that when on forms the union of that element with any other element y the result is y.

A more perspicuous way to encode this information is to give names to the biggest and smallest element of the algebra and to include these in the specification of the algebra. It is standard practice to denote the biggest element as '1' and the smallest element as '0'. Following this practice we can characterize Boolean algebra's as structures of the form $< X, \backslash, \cap, \cup, 1, 0 >$, which satisfy the axioms (AxBo.1-10) together with the following three (AxBo.13-16). In these last axioms we have also made us of another common notational convention, that of writing the complement operator as a bar over its argument (thus writing '$\overline{x}$' in lieu of '$\backslash x$').

(17) (Boolean axioms for 1 and 0)

(AxBo.13)	(for all x in X)	$x \cup 1 = x$
(AxBo.14)		$x \cap 0 = x$
(AxBo.15)		$x \cup \overline{x} = 1$
(AxBo.16)		$x \cap \overline{x} = 0$

We return to our exploration of the clauses in (12). Note that we are far from done: $\neg, \&$ and $\vee$ are not the only operators of our formalism; there are four more; their clauses are repeated here:

(12.v) $[[\phi \rightarrow \psi]]_{\mathcal{M},t} = 1$ iff $[[\phi]]_{\mathcal{M},t} = 0$ or $[[\psi]]_{\mathcal{M},t} = 1$;

(12.vi) $[[\phi \rightarrow \psi]]_{\mathcal{M},t} = 1$ iff $[[\phi]]_{\mathcal{M},t} = [[\psi]]_{\mathcal{M},t}$;

(12.vii) $[[P\phi]]_{\mathcal{M},t} = 1$ iff for some t' in T such that $t' \prec t$ $[[\phi]]_{\mathcal{M},t'} = 1$;

(12.viii) $[[F\phi]]_{\mathcal{M},t} = 1$ iff for some t' in T such that $t' \succ t$ $[[\phi]]_{\mathcal{M},t'} = 1$.

Of those remaining four, however, the first two, $\rightarrow$ and $\leftrightarrow$, are of little interest since they do not contribute anything new to the closure of the set of propositions introduced by I. The proposition denoted by $\phi \rightarrow \psi$ is equal to $(T \setminus [[\phi]]_{\mathcal{M}}) \cup [[\psi]]_{\mathcal{M}}$ and the proposition denoted by $\phi \leftrightarrow \psi$ is equal to $[[\phi \rightarrow \psi]]_{\mathcal{M}} \cap [[\psi \rightarrow \phi]]_{\mathcal{M}}$. These sets are part of the closure described above already.

The matter is different for the tense operators P and F. The propositions denoted by $P\phi$ is the set $\{t \mid \text{there is a } t' \prec t \text{ such that } t' \text{ belongs to } [[\phi]]_{\mathcal{M}}\}$ and in general there is no reason to assume that this set is a member of the closure under the operations $T \setminus$, $\cap$ and $\cup$. The same of course applies to F.

The closure of the set of propositions activated by an assignment I under all the operations of our formalism – $T \setminus$, $\cap$ and $\cup$, together with P and F – is thus a richer structure than the closure under just the 'Boolean operators' $T \setminus$, $\cap$ and $\cup$. Algebra's of this kind, which are generated from some set of basic elements by a set of operators which properly includes the Boolean operators, will always contain a Boolean algebra as a 'kernel', but will typically contain more besides. The algebraic perspective has been very fruitful in formal logic. But for us it is important primarily because of the role played by Boolean algebra's and similar algebraic structures.

3.4 Prior's 'Tense Logic and the Logic of Earlier and Later'

[**Warning!** I have added this part to the notes as a kind of IOY, since I had included Prior's paper among the reading material on Blackboard, and then decided that it wasn't fair to foist it upon you as is and promised to (i) to produce a more accessible introduction to Tense Logic (that's what I have tried to do in Sections 3.1-3.3) and (ii) to say something about the project Prior pursues paper while relying on that introduction. In spite of my efforts, the outcome is, I think, still not easy to follow, especially for those of you who have had no previous exposure to formal logic.

I am not expecting from any of you that you will actually read this part. If some of you take an interest and feel up to it, then: good for you, and let me know if there are things you would like some help with or that you think could be explained more clearly. And those of you who do not want to bother with this intrlude: Pass straight on to Section 4.]

In all likelihood Prior would not have approved of our discussion of the semantics of the (P,F)-calculus, since he would have regarded it as a case of putting the cart before the horse. Central to his concerns was the question: 'What is time as we conceptualize it?' and he thought that the answer to that question could be found by analyzing the temporal structure of our thoughts and the way that structure reveals itself in the forms of the languages that we use to express our thoughts. Tense Logic was meant to lay that structure bare and thereby provide us with a foundation from which time as an ordered structure of instants can be built, rather than a formalism that must prove its conceptual adequacy by showing compliance with a temporal structure that is given independently and in advance.

One way in which his paper 'Tense Logic and the Logic of Earlier and Later' could be read is as a reflection on such a program for the construction of time on the basis of the propositional structure rooted in language and thought. The paper proceeds by showing the successive stages of a stepwise reduction: It starts by laying out the 'Logic of Earlier and Later', a calculus that takes an ordered instant structure and the meaningfulness of statements of the form 'p is the case at a', where p is a proposition and a an instant, for granted and shows how in such a calculus the tense operators G and H can be defined. (Prior opts in this paper for the choice of these as primitive, rather than P and F, but nothing of importance hangs on that decision.) By way of a number of successive modifications Prior then transforms this Logic of Earlier and Later into a system of pure Tense Logic, which is much like the (P,F)-calculus as we have presented it here. (One difference is that Prior does not assume the linearity of time – neither in his initial Logic of Earlier and Later nor in the subsequent transformations of that system, so that in particular the final purely tense-logical system is without the linearity axioms (13.g,h). This makes for a number of complications that are finessed by assuming the linearity of time throughout. But Prior has good reasons for not wanting to assume linearity as given. One of his reasons has to do with the status of time in Relativity Theory, about which he has a number of interesting things to say in the final section of the paper. In these comments that twist to the general story will be ignored.)

The philosophical point of the reduction of Logic of Earlier and Later pure Tense Logic is clear. This is not just a reduction of one formalism to another; it is a reduction of an essential part of the *ontology* presupposed by the first system to the (more parsimonious) *ontology* presupposed by the second. More precisely: The Logic of Earlier and Later speaks overtly about temporal instants; in the Tense Logic that emerges at the end of the successive reductions the terms that in the Logic of Earlier and Later are used to refer to temporal instants have been reinterpreted as standing for a special kind of proposition – those propositions that incorporate all that is the case simultaneously somewhere in time (i.e. either now or somewhere in the past or somewhere in the future). So the ontology presupposed by the Tense Logic at the ned of the range of transformations consists of propositions only and instants are at this point just (propositions that are equivalent to) maximal conjunctions of simultaneously true propositions. (Recall in this connection the Russell-Wiener construction of instants as maximal sets of pairwise overlapping events that we went through earlier.)

Given what has been said here about the proof theory and model theory of Tense Logic, Prior's Logic of Earlier and Later can best be described as follows: Its formulas are in essence the kinds of statements that we get when we write out the truth conditions for formulas of Tense Logic in a given model at a given time. For instance, suppose that a model $\mathcal{M}$ and a an instant t from its instant set T are given and consider, say, the formula p & $\neg q$ and the formula $P(p \mathbin{\&} F\neg q)$. When we apply the truth definition for the (P,F)-calculus to p & $\neg q$ and the formula $P(p \mathbin{\&} F\neg q)$ and write out the result, what we get is the statement in (18.a). And as statement of the truth conditions of $P(p \mathbin{\&} F\neg q)$ we get (18.b).

(18) a. p is true at t in $\mathcal{M}$ and q is not true at t in $\mathcal{M}$.

b. There is a time t' in T such that $t' \prec t$ and p is true at t' in $\mathcal{M}$ and there is a time t'' such that $t' \prec t''$ and q is not true at t'' in $\mathcal{M}$.

Prior's Logic of Earlier and Later (LEL) is a formal system the formulas of which correspond closely to sentences like those in (18.a) and (18.b). More precisely, they correspond to statements that are like (18.a,b) except that there is no mention of a model in them or anything comparable to that. (Intuitively speaking, Prior's Logic of Earlier and Later assumes that there is a temporally ordered world that one is speaking about when using its formulas to make claims about that world. That these claims are about *this* world and not about any other is understood and need not be mentioned explicitly, just as we do not need to specify that it is about the actual world – the world that we are living in – when we say things about. For instance, we say 'It is raining' or 'Mary will come tomorrow', not 'It is raining in this world.' or 'Mary will come tomorrow in the actual world.' If at all, then only a philosopher would express himself in this bizarre and awkward fashion.)

But when the reference to $\mathcal{M}$ is dropped, as in (19), then (18.a,b) turn into statements that match the corresponding formulas from the Logic of Earlier and Later down to their constituents.

(19) a. p is true at t and q is not true at t.

b. There is a time t' in T such that $t' \prec t$ and p is true at t' and there is a time t'' such that $t' \prec t''$ and q is not true at t''.

LEL uses some of the abbreviations that are standard in formal logic (though the particular notation Prior uses is now widely considered arcane and instead I will use more generally accepted conventions in my presentation of this and the other systems Prior discusses in 'Tense Logic and the Logic of Earlier and Later', as I have already done in the discussion of Tense Logic in the earlier parts of Section 3). Among the abbreviations Prior uses are: '$(\exists t \in T)$' for 'there is a t in T such that', $\neg$ for 'not' and & for 'and'. In addition, Prior uses the symbol 'U' for the relation that we have been denoting as $\prec$ and the symbol 'T' for the 'true at' relation between instants and propositions. We will here continue to use '$\prec$' for the $\prec$-relation and 'AT' for the 'true at' relation. With these conventions (19.a) becomes (20.a) and (19.b) becomes (20.b).

(20) a. $AT(p,t)\ \&\ \neg AT(q,t)$.

b. $(\exists t')(t' \prec t\ \&\ AT(p,t')\ \&\ (\exists t'')(t' \prec t''\ \&\ \neg AT(q,t'')))$.

These two formulas give a good impression of what part of LEL is like. This part can be described as a system of classical quantification theory – with the logical operators $\neg$, &, $\vee$, $\rightarrow$, $\leftrightarrow$, $\exists$ and $\forall$ ('$(\forall t)$' reads 'for every t such that') – and with two kinds of variables, t, t', t'', .. for instants of time and p, q, r, .. for propositions; and there are two basic 2-place predicates, the predicate $\prec$ between instants and the predicate AT between instants and propositions.[1]

But this is not all of Prior's LEL. Prior also admits operators on the first arguments of AT one for everyone of the sentence operators of the (P,F)-calculus. That is, we also have formulas like the following:

$AT(\neg p,t)$, $AT(\neg(p\ \&\ \neg q),t)$, $AT(\neg p,t) \rightarrow \neg AT(p,t)$, $AT(\neg p,t)\ \&\ (\exists t')(t \prec t'\ \&\ AT(p,t'))$, $AT(Gq,t)$, $AT(\neg P \neg q,t)$.[2]

[1]In technical logical jargon, the calculus is a theory of 2-sorted first order logic whose non-logical vocabulary consists of the two predicates $\prec$ and AT, where $\prec$ has the 'signature' <t,t> and AT the signature <t,p>. ('t' stands for the sort 'temporal instant' and 'p' for the sort 'proposition'.)

[2]Prior is thinking of a Tens 3ogic in which the basic tense operators are G and H rather than F and P and in which F and P are short for $\neg G \neg$ and $\neg H \neg$. Nothing of importance hangs on the choice between G and H and F and P as primitive operators.

In other words, LEL does not only have formulas corresponding to statements like those in (18), which are obtained as the final analyses that the truth definition for Tense Logic provides for what it is for a complex formula to b true at a time t in a model $\mathcal{M}$, but also for the 'input statements' to such analyses, such as those in (21.a,b), as well as compound sentences containing statements like (21.a,b) as constituents, see (21.c,d):

(21) a. p & $\neg q$ is true at t in $\mathcal{M}$.

b. Gp is true at t' in $\mathcal{M}$.

c. p is true at t' in $\mathcal{M}$ or Fp is true at t' in $\mathcal{M}$.

d. If Gp is true at t' in $\mathcal{M}$ and $t' \prec t$, then p is true at t in $\mathcal{M}$.

Prior wants it to be possible to carry out the analyses of what it is for a complex formula of Tense Logic to be true at a time in a model as formal derivations in LEL. To this end he needs additional axioms. (Axioms in addition to some standard set of axioms and rules for the 2-sorted first order logic within which LL is formulated.) Among these are in particular those in (22).

(22) a. $AT((p \to q), t) \to (AT(p, t) \to AT(q, t))$

b. $AT(\neg p, t) \to \neg AT(p, t)$

c. $\neg AT(p, t) \to AT(\neg p, t)$

d. $AT(Gp, t) \leftrightarrow (\forall t')(t \prec t' \to AT(p, t'))$

e. $AT(Hp, t) \leftrightarrow (\forall t')(t' \prec t \to AT(p, t'))$

We can think of (22.d) and (22.e) as a kind of definitions of the operators G and H. In fact these biconditionals can be used directly in the derivations that mimic the reduction of the truth of a formula $G\phi$ or $H\phi$ at a time t to the truth conditions for ϕ. Similar biconditionals are also provable from the relevant axioms for the connectives of classical propositional logic, such as $AT(\neg p, t) \leftrightarrow \neg AT(p, t)$ (which can be obtained by conjoining (22.b) and (22.c)) and $AT(p \,\&\, q, t) \leftrightarrow AT(p, t) \,\&\, AT(q, t)$.
LEL is not a system of tense logic. Its formulas do not depend for their truth values on where in time they are asserted or evaluated. That this is so for its smallest formulas should be obvious: if it is true at any time t' that

$AT(p,t)$ – that is, that p is true at t – then this must be so at all times t'; likewise, if it is true at one time that t precedes t' then that too must be so at all times; whether a time t precedes another time t' is just a fixed property of the temporal ordering to which they belong and is independent of any particular perspective somewhere along that ordering. And it should also be obvious that this property (of one's truth or falsity being independent of any particular time) is inherited from these atomic formulas by all the larger formulas that can be built out of them.

We can, however, turn LEL into a kind of tense-logical system by treating the sentence letters as standing for 'generic propositions' (i.e. propositions that can be true at one time without being true at other times). That is, we now allow the simple p,q,r,.. as atomic formulas together with the atomic formulas of the forms $AT(p,t)$ and $t \prec t'$ that we already had. And we can go beyond this by admitting as formulas that stand for such generic propositions not only the sentence letters p,q,r,.., but also the compound formulas of the (G,H)-calculus that Prior adopts as his Tense Logic in this paper.

In fact, no syntactic distinction is to be made any longer between formulas from the (G,H)-calculus, formulas of the earlier LEL and formulas that can be obtained by combining the former and the latter. In particular, even formulas of the forms $AT(\phi,t)$ and $t \prec t'$ can occur as first arguments of AT. This means that among the formulas of the new formalism we have:

p, $p \rightarrow q$, Gq, $\neg H \neg q$, $AT(p,t)$, $AT(p,t)$ & p, $(q$ & $t \prec t')$, $AT(p,t) \rightarrow p$, $p \rightarrow (\exists t)AT(p,t)$, p & q & $(\exists t)(AT(p,t)$ & $\neg AT(q,t))$, r & Hq & $(\exists t')(t' \prec t$ & $AT(\neg r \vee \neg q)$, $AT(AT(p,t'),t))$, $AT(AT(p,t')$ & $q,t)$, $AT(q$ & $(\forall t')(t' \prec t \rightarrow AT(p,t')),t)$ and so on.

Let us refer to this new formalism as 'LEL+TL' (for 'Logic of Earlier and Later + Tense Logic')
There is still a big difference, of course, between the formulas of the old LEL and many of the formulas of LEL+TL: the old formulas of LEL are, just as they are as formulas of LEL as defined above, 'eternal' in the sense that they are true at all times if they are true at one time, whereas many of the new formulas do not have this property. (Potentially this is so for any formula that contains occurrences of sentence letters that are not part of formulas of the form '$AT(..,t)$'.) Furthermore the property of being an eternal formula

can be expressed within the new formalism itself. There are several ways of doing this. One way to say that the formula ϕ is eternal is by way of the following schema, which says that either ϕ is true at all times or ϕ is false at all times.

(23) $(\forall t)AT(\phi, t) \vee (\forall t)\neg AT(\phi, t)$

That formulas of the forms $AT(p, t)$ and $t \prec t'$ and all formulas that can be built from them are eternal may be clear from the intuitive meaning of AT and $\prec$, but it is something that LEL+TL must be 'told' in some way or other. The way to do this is to add yet another set of axioms, which state certain properties of eternal formulas that have to do with their being eternal and to do that in such a way that all other such properties become derivable. Prior chooses to this end the rule of proof in (24.a) and the axioms in (24.b-d).

(24) a. $\vdash \phi \;\Rightarrow\; \vdash AT(\phi, t)$

b. $(\forall t)AT(p, t) \rightarrow p$

c. $(\forall t)AT(p, t) \rightarrow AT((\forall t)AT(p, t), t')$

d. $AT(p, t) \rightarrow AT(AT(p, t), t')$

But besides formulas that are eternal by virtue of their form there are also some that are eternal for contingent reasons. For instance, it could be that the interpretation function I of some particular model $\mathcal{M}$ assigns to the letter p the proposition T, consisting of all instants of the time structure of $\mathcal{M}$. Then p will be eternal on this interpretation. And then of course there are the formulas of Tense Logic that are valid because of their tense-logical structure, such as, say, $PGq \rightarrow q$. These are, you might say, also eternal by virtue of their form. But in these cases it may be much more difficult to detect what it is about their form that is responsible for their being eternal than it is for the formulas of the old LEL.

In any case, we need an axiom which relates the statement that says of any formula ϕ that it is eternal to the plain assertion of ϕ: if ϕ is eternal then we are entitled to assert ϕ no matter where in time we are. This axiom is given in (25)

(25) $(\forall t)AT(p, t) \rightarrow p$

In our new system we can define an operator which says that the formula ϕ to which it applies is eternally true:

(26) For any formula ϕ, '$L\phi$' is short for '$(\forall t)AT(\phi, t)$'

It is now possible to derive from the axioms and rules we have that all formulas ϕ that are eternal for reasons of form (i.e those that come out as true at all t in all models) can be shown to be eternal in the formal sense that we can give, for each such formula ϕ, a formal derivation of $L\phi$. In particular, by combining the rules and axioms in (24) and (25) with the rules and axioms of 2-sorted predicate logic within which LEL has been formulated, which remain in force for the extended calculus we are looking at right now, one can derive $L\phi$ for all formulas of the old LEL. This is also possible for the logically true formulas of Tense Logic, such as $PGq \rightarrow q$: First a derivation mimicking the analysis provided by the truth definition for Tense Logic will yield the theorem $AT(\phi.i)$ for such a formula. We may then apply the rule of *universal Generalization* (this is a rule of proof for the underlying 2-sorted 1st order logic) to obtain from this the theorem $AT((\forall t)\phi.i)$, i.e. $L\phi$.

A formal confirmation from within the new system that 'L' does indeed behave like an 'always' operator, which can read as 'it is, always was and always will be the case' is obtained by showing that all formulas of the following three forms are provable as theorems.

(27) a. $L\phi \rightarrow \phi$
b. $L\phi \rightarrow G\phi$
c. $L\phi \rightarrow H\phi$

LEL+TL has an unconventional feature that comes to the fore when we compare the formula p – the formula that consists just of the sentence letter p – and the formula $AT(p, t)$. It could be said that, formally speaking, neither of them are *sentences* of the system: both contain free variable occurrences, p of the propositional variable p and $AT(p, t)$ of both p and the instant variable t. But the status of the propositional variables in LEL+TL is quite different from that of the instant variables. The best way to think of the difference is by going back to the model theory for Tense Logic that we took as a starting point for our formulation of LEL. In a model $< T, \prec, T, t_0 >$ for the (P,F)-calculus the interpretation of the propositional variables is fixed by I. But nothing fixes the interpretation of the instant variables. In LEL the letters p, q, .. and the tense-logical formulas that can be built from them – p & $\neg q$, Fp, $H(\neg p \rightarrow Gp)$ and so on – are, one could say, reinterpreted as simple and complex predicates of times. (So '$AT(p,t)$' could be read as 'the predicate p is true of the instant t', or 'the instant t instantiates the predicate p'.) In LEL+TL, however, the propositional status of p & $\neg q$, Fp, $H(\neg p \rightarrow Gp)$ etc. has been restored: They can now be used as standing by themselves, and not only as first arguments of occurrences of ART. In other words, these expressions have now regained the status of formulas that they have in systems of Priorean Tense Logic. And when they are used as formulas the sentence letters they contains are *not* to be thought of as variables but as standing for particular generic propositions, of which it is determined, for any instant of time t, whether or not they are true at that instant. But at the same time they also still function in essence as predicates of time, viz. when they do occur as arguments of AT.

But the instant variables do not admit of such an interpretation in either LEL or LEL+TL. They are variables in the traditional sense of formal logic and a formula that contains a free occurrence of one of these variables does, on account of that, not have the status of a 'sentence' – it is not fit to express a proposition. So, in particular, p can stand on its own as a formula, and when we think of it as asserted, we have to think of it as asserted at some particular time t_0. But $AT(p,t)$ cannot be thought of as assertable at all: attempts to use it to make an assertion are doomed to failure since a particular value for t, which would determine at which time p is supposed to be true, is missing. (One could introduce a convention according to which free occurrences of instant variables in formulas of LEL+TL that are used in assertions are interpreted as taking as their values the times at which those assertions are made; but LEL+TL doesn't do this and in fact, it is hard to see how it could since matters having to do with assertion (at different possible times) are beyond its ken.)

To conclude, in LEL+TL the propositional variables and instant variables have a different status, to the effect that p and other tense-logical formulas can be used in the capacity of sentences – that is, to express well-defined propositions, whereas formulas with free instant variables cannot.[3]

[3] A glaring illustration of the ambivalent status of tense-logical formulas in LEL+TL

Arguably the most decisive step in the successive transformations of the initial formalism LEL is the next one. It consists in reinterpreting the instant variables of LEL+TL as not standing for instants of time but for a certain kind of propositions – propositions that, intuitively and roughly speaking, are true at one time only. Such a proposition can be taken to identify the time at which it is true and as 'entailing' every other proposition at that time, in the sense that if p is the first, special proposition and q any other one, then $L(p \rightarrow q)$ will be true. (It will be true because the conditional $p \rightarrow q$ must be

---

is the derivation of the formulas (27.b,c). Consider (27.b). To prove that this formula is a theorem of the system we ned to derive the consequent Gp from the antecedent Lp. Now Lp is short for $(\forall t)AT(p,t)$. From this we can infer (using the rules and axioms of the underlying system of first order logic) the free variable formula $AT(p,t')$ and from that (using principles from classical propositional logic that are part of the rules and axioms of first order logic) that $t \prec t' \rightarrow AT(p,t')$. Since this formula follows from a formula that has no free occurrences of the variable t' (viz. $(\forall t)AT(p,t)$) first order logic allows us to conclude that

(1) $(\forall t')(t \prec t' \rightarrow AT(p,t'))$

also follows from this premise. According to (22.d) this is equivalent to (2):

(2) $AT(Gp,t)$

So we have now derived (2) from our premise. We can now apply once more the principle of first order logic that we used to obtain (1) from our premise $(\forall t)AT(p,t)$: since the premise does not contain any free occurrences of t we may infer that

(3) $(\forall t)AT(Gp,t)$

Since (3) is of the form $L\phi$, we can, using (24.b), infer from it the formula Gp, which is what we wanted.

The moral of this exercise is that in order to arrive at the desired conclusion Gp, we first have to prove the formula LGp. Only then can we apply the 'bridging axiom' (24.b). But in order to get to LGp we first had to obtain the formula $AT(Gp,t)$. That may look quit close to the formula Gp that we are aiming for; but the two aren't equivalent since one has a free variable occurrence while the other does not. In order to get from $AT(Gp,t)$ to Gp we first have to pass through the intermediate step to the effect that since $AT(Gp,t)$ has been derived from a premise that does not have any free occurrences of t, this conclusion must hold irrespective of what value might be assigned to t; that is we first have to infer (3), that is LGp. Only then can the 'bridging axiom' be applied. (I am calling it a 'bridging axiom' because it permits the transition from a formula in which a formula ϕ occurs as argument of AT to an occurrence of it in which it stands on its own.)

true at all times, and that is so because at the time identified by p q is true by assumption, while at all other times p is false.) In fact, if p is a proposition of this special kind, then we will have for any q whatever that either $L(p \to q)$ is true or else $L(p \to \neg q)$ is. So the disjunction $L(p \to q) \vee L(p \to \neg q)$ will be true for every q whatever. Prior uses this fact to define what it is for a formula ϕ to denote a proposition of the special kind, using at this point quantification over arbitrary propositions, as in (28). (28) expresses that ϕ denotes a proposition of the special kind.

(28) $(\forall q)(L(\phi \to q) \vee L(\phi \to \neg q))$

The reinterpretation of instant variables as variables that range over propositions of the special kind entails that quantification over instant variables, as we have it and have been using it in both LEL and in LEL+TL, now becomes quantification over such special propositions. The formula in (28) can be employed to recast such quantifications over arbitrary propositions, which are restricted by this formula. Thus for instance, the formula in (29.a) can be rewritten as (29.b) and (29.c) as (29.d).

(29) a. $(\exists t)AT(p \ \& \ Fq, t)$

b. $(\exists r)((L(r \to q) \vee L(r \to \neg q)) \ \& \ AT(p \ \& \ Fq, t))$

c. $(\forall t)(AT(p \ \& \ \neg q, t) \to AT(p, t))$

d. $(\forall r)((\forall q)(L(r \to q) \vee L(r \to \neg q)) \to (AT(p \ \& \ \neg q, t) \to AT(p, t)))$

Let as abbreviate the schema (28) as '$IP\phi$' (for 'Instant Proposition'). Then, for instance, the formulas in (29.b) and (29.d) can be written in the more intuitive and more easily surveyable forms in (30.a,b).

(30) a. $(\exists r)(IPr \ \& \ AT(p \ \& \ Fq, t))$

b. $(\forall r)(IPr \to (AT(p \ \& \ \neg q, t) \to AT(p, t)))$

In order that the role of instant variables in LEL and LEL+TL can be transferred without loss or remainder to variables for instant propositions we must make sure that there are enough instant propositions around. More specifically we must make sure that there is an instant proposition that is true now and that it always was the case that some instant proposition was true and that this also always be the case. One way in which we can encode these requirements is to adopt the following axiom.

(31) $(\exists r)(IPr\ \&\ r)\ \&\ H(\exists r)(IPr\ \&\ r)\ \&\ G(\exists r)(IPr\ \&\ r)$

Having made the transition from instant variables to proposition variables we can now also eliminate the instant predicate $\prec$ as a primitive symbol from our formalism, defining it instead in tense-logical terms:

(32) '$\tau \prec \tau'$' is short for '$(IPr\ \&\ IPr'\ \&\ L(r \to Fr'))$'.

(Intuitively the formula on the right says that at any time at which r is true thn r' will be true, which is, when r and r' are instant propositions, precisely what precedence amounts to for the instants they identify.)

Moreover, it is now also possible to eliminate the AT predicate: If r is an instant proposition that is true (only) at instant t, then $AT(\phi, t)$ can be replaced by $L(r \to \phi)$:

(33) '$AT(\phi, \tau)$' is short for '$L(r \to \phi)$'.

Finally we can also recast L in purely tense-logical vocabulary. Since we are assuming time to be linear, we can simply 'define'as the conjunction of the three consequents of the formulas in (27), i.e. as made explicit in (34).[4]

(34) '$L\phi$' is short for '$\phi\ \&\ H\phi\ \&\ G\phi$'.

[4]Prior does not make the general assumption that time is linear. For this reason he cannot define L in the way that is don here. (Note well, L is to be such that $L\phi$ should entail ϕ for any tense-logical formula ϕ whatever, including those that involve any numbers of nested occurrences of tense operators. In particular the entailment should hold for a formula such as HGp. But in order that HGp be true at t it must be that p is true at any time t'' that is in the future of any time t' that is in the past of t. But if time is not linear – for instance, if it is branching in the direction of the future: there can be times t and t'' that are both in the future of some earlier time t', but with neither t in the future of t'' nor t'' in the future of t nor with the two coinciding – then HGp could be false at t even though $(p\ \&\ Hp\ \&\ Gp)$ would be true. In order to arrive at a tens-logic al definition of L such that $L\phi$ entails ϕ for all tense-logical formulas ϕ assuming that time might not be linear, Prior goes through some fairly fancy construction that actually presupposes that we are working within *Infinitory Logic*, which admits of infinitely long formulas. This is a complication that I definitely do not want to go into here. So that has been a technical reason for making the assumption that time is linear up front. But it seems to me that linearity is quite a plausible assumption to be made about time in any case.)

We now have all the bits and pieces in hand for a systematic translation of all formulas of LEL+TL into formulas of Tense Logic (more precisely, of the (G,H)-calculus). Let ψ be any such formula. Then we choose for each instant variable occurring in ψ a corresponding proposition variable that does not yet occur in ψ (so that we get a 1-1 correspondence between ψ's instant variables and the new proposition variables). And then we replace all occurrences of '$L\phi$', '$\tau \prec \tau'$' and '$AT(\phi, \tau)$', working 'from the inside out' (i.e. starting with the smallest subformulas of ψ, then passing to those which are no the smallest and so on up). Then the resulting formula ψ' will be a formula of the (G,H)-calculus. And it will always be possible to derive in LEL+TL that ψ' is equivalent to the formula ψ of which it is the translation result.

That means that we can now, if we want to, kick the stepladder of LEL+TL away and confine ourselves to just those formulas that are formulas of the (G,H)-calculus. Since for each of the formulas of LEL+TL there is an equivalent formula of the (G,H)-calculus, none of the expressive powers of LEL+TL are lost thereby. What we are left with is an extension of the (G,H)-calculus, in the sense of Tense Logics as presented in Section 3.1; this is an extension insofar as it includes quantification over propositions. (Remember that such quantification was needed both in the definition of IP and in the statement of Axiom (ref28).) A further reduction to a version of the (G,H)-calculus without such quantification is not aimed for by Prior in this paper and as far as I can see would not be possible. The importance of this point can hardly be overemphasized. For quantification over propositions can, unlike quantification over particulars (such as, among others, instants of time) add to a formal calculus expressive powers (and with these formal complications) that are not always easy to foresee. (In this particular instance, the effects of adding quantification over propositions happens to be known fairly precisely – at least that is the case so long as we assume that time is linear. I am not sure what can be said in case the linearity assumption is dropped. But this simply an admission on my part of insufficient knowledge on my part of the state of the art in this domain of formal logic.)

A more parsimonious solution is to adopt two sorts of propositional variables, say p_0,p_1,p_2,.. ranging over arbitrary propositions and r_0,r_1,r_2,.. ranging over instant propositions. There is then no need to define instant propositions as arbitrary propositions of a special kind, (which is precisely where quantification over propositions was needed in the proceedings described above).

Tense-logical systems of this sort are known as *Hybrid Logics*. They were developed by Patrick Blackburn.
The importance of the distinction between formulas of LEL+TL with and those without free instant variables became clear at one point when we trying to understand how LEL relates to LEL+TL and, via LEL+TL, to pure Tense Logic. That difference is of course still with us – it must be, since our translation procedure applies equally to formulas with and without such free occurrences. It now manifests itself in that in the translation of a formula ψ in which all instant variables are bound (by quantifiers), all letters for instant propositions that have been introduced to replace the instant variables in ψ will be in the scope of some occurrence of L. And when ψ is a formula of LEL, then *all* proposition letters in the translation will be in the scope of occurrences of L. But if ψ is not closed in this sense, however (i.e. if ψ does contain free occurrences of instant variables), then this will in general not be so. The semantic significance of this difference should be clear (and is easy to verify in formal detail): the translations of closed formulas of LEL (those which have no free occurrences of instant variables) are all of them *eternal* formulas: in any model for Tense Logic the formula will be either true at all its instants or else false at all its instants.

For the translations of formulas of LEL that do contain free instant variables this will in general not be so. Nor will it be so in general for formulas of LEL+TL that do not belong to LEL, even if all their instant variable occurrences are quantificationally bound. And among those formulas there are in particular those that only involve the primitive vocabulary of the (G,H)-calculus and which are therefore not changed in any way by the translation procedure we have described. For as we have seen repeatedly by now, formulas in which not every occurrence of a sentence letter is within the scope of an occurrence of L may nevertheless be eternal. Among the examples for which that isn't so are the tense logical axioms and theorems that we passed on our way in Section 3.2, such as for instance $PGq \rightarrow q$. But the difference is that for the translations of closed formulas of LEL the fact that they must be eternal can be read straightforwardly off from their form, just as that is the case for the LEL-formulas of which they are the translations. In other cases determining that a formula is eternal ned not be all that straightforward and in some cases deciding the matter can be quite difficult.

The stepwise transition from LEL to the (G,H)-calculus is Prior's way of showing how, and in what sense, the Logic of Earlier and Later can be reduced to Tense Logic. A reduction of Tense Logic to the Logic of Earlier and Later is possible as well. But in this case the reduction has a somewhat different flavor. As a matter of general principle the translations of formulas of Tense Logic into the Logic of Earlier and Later lead to formula with a free instant variable – representing, intuitively speaking, the time at which the tense-logical formula is taken to be asserted or evaluated. But of course it may be that this free occurrence is 'vacuous' in that in any model the translating formula will be satisfied by one time only if it is satisfied by all.

The moral for the languages we speak, I take it, is that while the statements they allow us to make are generally organized in the manner of the tenses of the verb – i.e. as talking about certain things as in the past, others as in the present and yet others as in the future – and while for this reason what proposition we express in making such statements will depend on when they are made as well as on the sentences used, it is nevertheless possible to make statements that are immune to such variation. And statements that correspond in form to the closed formulas of LEL – statements in which we say things about what is the case at some time or at all times. Such statements about 'earlier and later' are independent for their truth or falsity from the time at which they are made as are countless others, for many of which this may be less transparent from their form or apparent content.

4 Tense and Aspect in Natural Language: The First Beginnings

In this section we will go over some of the early work that has made a lasting impact on current theories of temporal and aspectual features of natural languages. We focus on classical papers by Vendler, Reichenbach and Bach.

4.1 Vendler: 'Verbs and Times'

The lasting contribution of this paper is its proposal of a four-fold apsectual classification of (simple and complex) verbs into (i) accomplishment verbs, (ii) achievement verbs, (iii) activity verbs and (iv) state verbs.

(35)

accomplishments	achievements	activities	states
write a letter	die	run	know
build a house	reach the top	push a cart	love
run to the store	win the race	breathe	be sick
open the door	stop running	knock	be obnoxious

Tests for classifying verbs as belonging to one of these four categories.

1. The *'in X amount of time'/'for X amount of time'*-test.

(36) a. She wrote the book in 8 months. $\surd$
b. She wrote the book for a year. ??

c. She reached the top in half an hour. $\surd$
d. She reached the top for half an hour. ??

e. She pushed a cart for 15 minutes. $\surd$
f. She pushed a cart in 15 minutes. ??
g. She was sick for three weeks. $\surd$
h. She was sick in three weeks. #

According to this test it is the accomplishment verbs and the achievement verbs go with adverbs of the form ‘in X amount of time’, but not with adverbs of the form ‘for X amount of time’, whereas activitiy and state verbs go with adverbs of the form ‘for X amount of time’, but not (or under duress or with another, here unintended meaning) with adverbs of the form ‘in X amount of time’. Verbs that go with adverbs of the form ‘in X amount of time’ but not with adverbs of the form ‘for X amount of time’ are often referred to as telic (for Greek ‘telos’, Engl: ‘goal’); verbs that go with adverbs of the form ‘for X amount of time’ but not straightforwardly with adverbs of the form ‘in X amount of time’ are called atelic.

In other words, accomplishment and achievement verbs are telic; activity verbs and state verbs are atelic.

2. The ‘at a time’-test. Achievement verbs differ from accomplishment verbs in that they go naturally with adverbs of the form ‘at t’, where ‘t’ denotes some particular instant of time. Accomplishment verbs do in general not go happily with such adverbs:

(37) a. She died/reached the top/won the race/stopped running at 5.15. $\surd$

b. She wrote the book/built the house at 5.15. ??

c. (But: She opened the window at 5.15. $\surd$)

As our intuitions about (37.c) seem to indicate, the ‘at t’-test isn’t completely reliable as a way of distinguishing achievements from accomplishments. The problem is two-fold. What the test tries to capitalize on the idea that we conceive of the events described by achievement verbs as instantaneous and of those that described by accomplishment verbs as non-instantaneous, as taking time, as having some ‘non-zero duration’. But of course in actual life no event can really be so short as taking no time at all. (If it were, we would never be able to observe it.) On the other hand there is a problem with the way we use and understand ‘instant-denoting terms’ such as *at 5.15*. Usually we treat these not as identifying the indivisible instant which qualifies as the one and only one to which ‘5.15’ applies in the strictest sense possible. So we accept (37.c) as possibly true, viz. when the window opening of which it speaks occurred in a brief interval around the exact instant of 5.15. (or some short interval in its immediate vicinity).

Another test for distinguishing between achievement verbs and accomplishment verbs is the ‘it took x amount of time for a to V’-test. Compare:

(38) a. It took her 8 months to write the book.

b. It took her half an hour to reach the top.

The difference between (38.a) and (38.b) is not that between acceptability and unacceptability, or between naturalness and strangeness. Both sentences are perfectly acceptable. But there is nevertheless a significant difference in their semantics. We understand the period denoted by the adverb in (38.a) as that during which the book was written. In contrast the period denoted by the adverb in (38.b) is not understood as the duration of the event of reaching the top, but rather of the period needed to reach the time when that event could take place (e.g. the time from the moment the agent started climbing in order to reach the top).

3. The ‘Progressive’-test. This test distinguishes between stative verbs and and the other three aspectual classes. So it is useful in particular as a means for separating the state verbs from the activity verbs. (This is the one distinction for which we do not have a test so far.) It is a test, by the way, that we have in English, which has ‘progressive’ as opposed to simple tense forms. Most languages which are in other respects much like English when it comes to matters of tense and aspect (such as for instance French or German) do not hav this distinction, so there this test doesn’t exist.

Consider (39)

(39) a. She was running/pushing a cart/breathing/knocking. $\surd$

b. She was knowing the answer. *

c. She was loving Bob/being sick/being obnoxious. */?/#

There are two aspects to the ‘Progressive’-test:

a. Certain state verbs cannot appear with any progressive tense form (or at least not without reinterpretation as some kind of activity).

b. When the present tense is used in its 'standard sense' (i.e. to describe what is going on while the speaker is producing her utterance), then non-stative verbs *must* appear in the progressive.

For instance, the replies in (40.a) are simply not right; the proper form to use here is the present progressive not the simple present.

(40) a. (A: What are you doing?)
   1. B: I eat an apple. *
   2. B: I am eating an apple. $\surd$

b. (A: What are you doing?)
   1. B: I run. *
   2. B: I am running. $\surd$

(The replies in (40.a) are typical mistakes made by nonnative speakers whose native languages are German or French, where the distinction between progressive and non-progressive forms is not found.)
The Progressive is also involved in two other aspect tests. The first of these is meant to distinguish between the telic and non-telic event descriptions (the accomplishment and the achievement verbs as opposed to the activity verbs). (So the test is supposed to accomplish roughly the same as the 'for an hour-in an hour' test.) The second is to give us a way of telling achievements from accomplishments.

4. The use of the Progressive to distinguish between telic and non-telic aspect is connected with a phenomenon known as the 'Imperfective Paradox'. The phenomenon is illustrated by the contrast between the triple in (41.a) in the one hand and those in (41.b,c) on the other.

(41) a. 1. She was running.
   2. She ran.
   3. She has run.

b. 1. She was building a house.
   2. She built a house.
   3. She has built a house.

c. 1. She was dying.
   2. She died.
   3. She has died.

The difference between (41.a) and (41.b,c) is that (41.a.1) seems to entail (41.a.2) and (41.a.3). But no such entailments seem to hold in the case of either (41.b) or (41.c). To amplify the point: Suppose that it is true that she was running and you know this, and I then ask you 'Did she run?'. It seems that you cannot truthfully answer that question with 'No'. And, it has been contended, the same holds for when I ask you 'Has she run?'. For (41.b) and (41.c) the facts are different. There are situations in which it would be true to say 'She was building a house.' but where the house she was building never got finished or even got properly off the ground. If moreover the subject didn't ever complete the building of any other house, then the answer to either 'Has she built a house?' or 'Did she build a house?' should be in the negative. Likewise, someone may have been dying without actually dying in the end. (She turned round with one foot already in the grave.) Then clearly it would be false to answer 'Yes' to the questions 'Has she died?' or 'Did she die?'.

The point should be clear: With accomplishments such as 'build a house' or achievements such as 'die' a description in the Progressive can be true without the corresponding non-Progressive description being true – you can have been involved in the building of a house but yet have failed to build a house, since the house never got completed and in fact very little of it may actually ever have been there. And you may have been in the process of dying without dying for real. But with activities the matter is different, running is running – an activity – whether you describe it with or without the Progressive form.

This test too is not without its problems. The negative part of it – there can be situations that can be correctly described using the past progressive form of an accomplishment or achievement verb but not by using its Simple past or Present perfect – seems solid enough. Less clear, from a strictly linguistic perspective, is the positive entailment claim about activity verbs that makes for the other half of the test. Are (41.a.2) and (41.a.3) really entailed by (41.a.1)? It does not seem impossible to conceive of a situation in which we would agree to the truth of (41.a.1) and yet not be prepared to approve of (41.a.2) or (41.a.3) without reservations or qualifications. Suppose for instance that the subject is in the habit of running a certain distance every day – maybe she is in training for some athletic event. But today she was running along, intent on completing her daily task, when halfway she got a cramp and had to stop. In that situation it seems fine to say 'She was running (when she got this cramp).', but it is not so clear what you should answer to my question later in the day 'And has she run?'. Simply saying 'Yes' might well be misleading in this case. A less misleading way of phrasing your reply would be something like 'Well, she did set out, but then she got this cramp so she had to stop.'

What does this show about the positive side of this test? One way to defend it against examples like the one just presented is to point out that the use of *run* that makes the answer 'She has run.' (or 'She ran.' or 'She did run.') sound wrong or awkward in the situation just considered has to do with the fact that in this situation the verb can be reinterpreted as an accomplishment verb, roughly synonymous with 'do one's daily bit of running'. But this can only be a partial defense. For even in the situation we are contemplating the activity reading of *run* is in principle still available. Presumably it is in that activity sense that *run* is understood in the statement 'She was running.', even when it is used with reference to this situation. Rather, it is the *combination* of the scenario in question *and the tenses of the b-and c-sentences* that so strongly suggests the interpretation of *run* as 'do one's daily run'. And if this is right, then that would seem to imply that a straightforward application of the test is not always possible (and perhaps that it never is).

In the course of the modern history of theorizing about tense and aspect the term 'Imperfective Paradox' seems to have undergone a shift of emphasis. Rather than a name of the difference in behavior of the progressive forms of activity verbs on the one hand and accomplishment and achievement verbs on the other, in the way that (41) was meant to illustrate, it now is more often used to point at the fact that the progressives of telic event descriptions can be true while the corresponding non-progressive descriptions are false. Semanticists have come to perceive this as a kind of puzzle. It seems intuitively clear that we understand what makes such progressive descriptions true in situations where the non-progressive alternatives do not apply on the following basis: We combine our understanding of the given verb with a general understanding of what the Progressive can do with arbitrary telic verbs to which it is applied. Or, to put the matter somewhat more formally: from a semantic point of view the Progressive appears to be an operator which among other things can be used to transform the meanings of telic event descriptions into other (apparently non-telic) descriptions. But exactly what *are* the meanings of these 'progressivizations' of telic event verbs? What is the general principle according to which we compute them from the semantics of the underlying telic verbs?

Here is a brief indication of why people have come to see this as a real puzzle. Many of the telic descriptions that give rise to the phenomenon in question – the progressive form can be true in situations where the non-progressive form is not – are descriptions of voluntary actions, as in (41.b). At least for these cases the following formula for how the progressive form is semantically related to the non-progressive form seeems plausible:

(42) The progressive form of an action description applies to all those cases in which the agent *intends* to perform an action of the sort described by the non-progressive form and *attempts* to do so; but there is no need that the attempt be successful, i.e. that the action is completed as planned.

For many progressives of action descriptions this formulation appears to be a fairly good approximation. (A classic example is that in (43).

(43) She was crossing the street when she was hit by a truck.)

But there are also counterexamples, viz. when the subject's intention seems wholly unrealistic. For instance, it has been contended that (44) cannot be accepted as true in a situation in which the subject is intercepted by a police boat half a mile from the shore from where she set off, because she would never have had a chance of making it to the other side no matter how intent she may have been of doing that and no matter how convinced she may have been that she would make it.

(44) She was swimming across the ocean when she was intercepted by a police boat.

But how are we to demarcate the 'possible cases', in which a sincere attempt on the part of the agent may be considered sufficient for the truth of the progressive form, from the 'impossible cases', in which it appears not to be? This is just on of the questions about the meaning of progressive forms that have been the topic of much linguistic debate.

5. Use of the Progressive to separate achievements from accomplishments.

The intuitive picture of the difference between accomplishments and achievements is that an event described by an accomplishment verb consists of a period during which a certain activity takes place which leads up to the 'instant of accomplishment' – the instant at which the described event is completed and thgat is part of the event described by th (non-progressivized) verb. In the case of an event described by achievement verbs there is just this instant: since the event is conceived as instantaneous, this instant is conceived as the only one at which the event happens; any process or activity that leads up to it is not part of it.

This difference manifests itself in particular when accomplishment and achievement verbs are put in the Progressive form. When the Progressive is applied to an accomplishment description, what it picks out is normally understood to be part of the event that is described by the non-progressive form of the verb (or of the event that would have resulted had the non-progressive form been properly instantiated). As opposed to this, the progressives of achievements do not describe what is part of the events described by the non-progressive forms but rather some process or activity that can be understood as leading up to such an event without being a part of it. The difference can be seen in the following two examples.

(45) a. She was tidying up the apartment for three days.

b. She was dying for three days.

c. Yesterday at five, when I rang up to find out how she was doing, she was tidying up the apartment.

d. Yesterday at five, when I rang up to find out how she was doing, she was dying.

Let us assume for the sake of argument that for each of the four sentences in (45) the event described by the non-progressive form actually happened. Then the obvious difference between the sentences involving accomplishment descriptions ((45.a) and (45.c)) and those involving achievement descriptions

((45.b) and (45.f)) is that the activity that is described in the first two sentences is part of the completed tidying-up event (the one that is completed when the apartment is at last fully tidied up). But the process described in (45.b) and (45.d) is not part of the event of dying – that event which occurs just and only when the process has run its full course.

This difference is made particularly clear when the progressive description occurs in conjunction with an adverbial that indicates an instant at which the described process or activity is going on, as in (45.c,d). In (45.c) the time of five o' clock that is contributed by the adverb is understood as lying within the complete tidying-up event described by the non-progressive form. But in (45.d) the time of five o'clock is *not* understood as part of the event that can be described non-progressively. Far from entailing that she died at five, (45.d) rather suggests that the event described by the non-progressive sentence was not going on at that time, so that 'She died at five.' would actually be false.

We conclude this discussion of the difference between achievements and accomplishments with another bit of terminology. Both achievement verbs and accomplishment verbs, we have noted, are telic in that there is a natural completion to the events they can be used to describe. In the current literature on tense and aspect this natural completion point is called the *culmination point* of the event. (So the events described by achievement verbs consist of their culmination points only – they *coincide* with their culmination points.) We will often use this term.

So much for tests that help us in classifying verbs according to Vendler's four-fold scheme. We now turn to additional qualifications for the application of these tests and in fact of the whole project of aspectual classification. The first qualification was made by Vendler himself.

One very important observation Vendler makes is that there are many verbs that according to his tests fit more than one of his four categories. Among these are verbs that according to the Progressive-test seem to qualify both as state verbs and as activity verbs. Some examples:

(46) a. 1. I think that he might still be in Paris.

   2. I am thinking that he might still be in Paris.

b. 1. Of course, we treat your case as confidential.
   2. Of course, we are treating your case as confidential..

c. 1. The statue stands in the middle of the square.
   2. The statue is standing in the middle of the square.

But as Vendler notes, verbs may bridge categories also in other ways.

One other quite common pattern is for a verb to manifest itself both a state verb and as an achievement verb. Two such verbs are *know* and *see*:

(47) a. I know that she is in Paris.

b. I see that you are right.

c. I see a raft that is slowly floating downstream.

d. And then I saw that she was right

e. And in that instant I knew what had to be the correct answer.

Vendler discusses the verb *see* at some length. He notes that it not only has the state and achievement uses illustrated by the examples in (47), but also that when used as in (48.a), it behaves as an accomplishment verb.

(48) a. We saw Carmen last night.

b. A: Where are you? What are you doing right now?
   B: Right now I am seeing my grandmother.

c. He is seeing Julie.

d. Are you still seeing Julie?

e. Weren't you seeing Julie at some point?

f. (Didn't you see Julie at some point? (?))

Vendler says at one point that *see* cannot be used as an activity verb:

'But seeing cannot be a process. *What are you doing?* can never, in good English, be answered by *I am seeing...*'

But what about (48.b-e)? All not good English?

Some further interesting observations of Vendler's:

(49) a. H thinks that Jones is a rascal.
b. H is thinking about Jones.
c. I saw him crossing the street.
d. I saw him running.
e. I spotted him crossing the street.
f. I spotted him running.
g. I saw him cross the street.
h. I spotted him cross the street. (*)
i. I spotted him run. (*)
j. I saw him cross the street. ($\surd$)
k. I saw him run. ($\surd$)

The circumstance that so many verbs have a semantics which is flexible enough to allow different instances of the same verb to display different Vendlerian categories makes aspectual classification in his sense a much more delicate and problematic matter than it would have been otherwise. In fact, this flexibility poses a serious problem for lexical semantics even when the problems of Vendlerian classification are set aside. Do we want to treat aspectually flexible verbs as *ambiguous* between the different aspectual senses with which they can be used? Or should we try to treat them as having one single sense but assume that this sense can be transformed in various ways, and then attempt to say what triggers such transformations on particular occasions of use? This is a methodological issue that arises in many other parts of lexical semantics as well, and it seems that there is no single, uniform solution – each case has to be judged on its own merit.

There is however also another aspect classification problem, which is different from anything that Vendler mentions in his paper but that turns out to have close connections with the issue about individual uses of verbs that he does raise. This is the question *which expressions* should be the targets of aspectual classification. What Vendler embarks on is a classification of verbs. But almost from the start there are among the 'verbs' that he presents as examples expressions that do *not* consist of just a single verb, but of a verb together with something else, as in *run to the store* or *run a mile*. These two examples are particularly apt to bring out the problem that this and the next few paragraphs will be devoted to. Recall that the simple *run* also figures among Vendler's examples and that he (correctly) cites it as an example of an activity verb. In contrast, *run to the store* and *run a mile* are accomplishments, as is unequivocally established by the 'in-an-hour-for-an-hour' test.

What this triple of examples seems to show is that expressions that belong to one of Vendler's classes can be turned into an expression belonging to some other class by combining it with some other expression. Moreover, the ways in which aspectual category can be modified by such means is clearly something of which we, the users of the language, have a pretty good general grasp. (If not, then this whole discussion could not be conducted along the lines wv are conducting it.) So the moral appears to be: It is not only basic (or 'lexical') expressions that are up for aspectual classification (leaving open what how the class of basic or lexical expressions is supposed to be determined for this purpose) but also certain expressions that can be obtained from these by combining them with other expressions into larger, compound expressions.

But this then raises two further questions:

(i) When verbs are combined repeatedly with other constituents into increasingly bigger expressions, at what point are aspectual classificationa no longer applicable? and

(ii) What are the principles that determine the aspectual classification of (classifiable) compound expressions on the basis of the aspectual classification of their classifiable parts and the way that this part is combined with one or more other expressions so as to yield the given compound expression?

One context in which these questions come up are the very tests that Vendler proposes for classifying the simple verbs and 'compound verbs' the classification of which directly concerns him. Take for instance the 'in-an-hour-for-an-hour' test. The point of this test is that the *in*- and the *for*-adverbials go with expressions belonging to some of his categories and not with expressions belonging to the others. But now consider the combinations which are possible, such as *write the letter in an hour* or *run for an hour*. What if any are the categories that these expressions belong to? Vendler doesn't ask this question, but it is one that might well be asked and in fact some discussions in the literature turn on precisely these questions.

All these questions are hard, and in spite of the fact that they have been intensively studied for many decades we are nowhere near definitive answers to them. Part of the problem is that some linguists have approached the problems of aspect from the other end, so to speak, starting with a conception of 'sentence aspect' – of aspectual distinctions between full sentences – and have tried to answer what in the sentences they were looking at could be responsible for the aspectual distinctions they could see. If this is your starting point for questions of aspectual classification, then the question will rather be *how far down* such classifications should be assumed to reach.

If both sides – those who like Vendler approach the problems of aspect classification below and those who approach them from above – are right, then presumably we will find aspectual classification all the way up from the lexicon to the full sentence. But of course it cannot be assumed that the relevant aspectual categories will be the same at all levels: categories that are right for lexical verbs, and for phrases that are somewhat but not very much more complex, may not be the right ones for complete sentences (and, perhaps, for expressions that are sentence-like without yet being complete sentences in the strict sense). We shall have to return to these questions again and again.

4.2 Reichenbach: The Tenses of Verbs

The central insight of this paper is that the correct interpretation of certain tenses requires a *Reference time*, as distinct from both the *Speech time*, at which the sentence containing the given tense is made, and also from the *Event time*, the time when the event (or state) described by the sentence is said to occur.

Reichenbach's first and most salient example of such a tense is the past perfect and his first illustration of the role of the Reference time in the interpretation of the Past perfect is the following passage from the novel 'Of Human Bondage' by Somerset-Maugham.[5]

(50) But Philip ceased to think of her a moment after he **had settled down** in his carriage. He thought only of the future. He had written to Mrs. Otter, the *massière* to whom Hayward **had given** him an introduction, and had in his pocket an invitation to tea on the following day.

Reichenbach notes that the Past perfects in this passage locate the events described by their verbs as lying in the past of some past time that has already been introduced into the text. In this case this is the 'moment' at which he ceased to think of the person referred to by *her*.

This time – the moment at which he ceases to think of 'her' – is in Reichenbach's terminology the *Reference time* for both past perfects in this passage.

Reichenbach went on to postulate a Reference time as part of the semantics of all the English tenses: Each tense conveys two connected bits of information,

[5]Reichenbach also gives a second example, a passage from a text by Macaulay about the reign of Charles II of England. Again, the effect of the occurrences of the Past Perfect is one that can be appreciated only within the setting of the multi-sentence text;

'In 1678 the whole series of events had changed eighteen years of misgovernment **had made** the majority desirous to obtain security for their liberties at any risk. The fury of their returning loyalty **had spent itself** in its first outbreak. In a very few months they **had hanged and half-hanged, quartered and emboweled**, enough to satisfy them. The Roundhead party seemed to be not merely overcome, but too much broken and scattered to ever rally again. Then commenced the reflux of public opinion. The nation began to find out to what a man it had intrusted without conditions all its dearest interests, on what a man it had lavished all its fondest affections.'

one about the relation between the Speech time S and the Reference time R and the other about the relation between R and the Event time E.
The following table summarizes his proposal.

(51)

Traditional Name	English tense form	Relation between R and S	Relation between E and R
Past perfect	*has come*	R < S	E < R
Simple past	*came*	R < S	E = R
(Future of the past)	*would come*	R < S	E > R
Present perfect	*has come*	R = S	E < R
Present	*comes*	R = S	E = R
Future perfect	*will have come*	S < R	E < R
Simple Future	*will come*	S < R	E = R
??	??	S < R	E > R

Some of the details of this table may be up for criticism. There are problems, for example, with what Reichenbach proposes for the English Simple Past and Present Perfect. He notes that there are differences in this regard between English and certain other European languages. For instance, German (Reichenbach's mother tongue) will often use the Present Perfect form where English uses a Simple past, and in some German dialects (especially in the South) the Simple past is hardly found any more. It is tempting, and arguably correct, to hold that at least in the southern German dialects the Present Perfect form has *de facto* adopted the semantic characteristics of the English Simple Past (and of the German Simple past as it is still used in the North of Germany).

What is problematic, however, is the assumption that the English Present Perfect should be analyzed as a compound tense at all. Current views of the matter see the perfect as an *aspectual* operator which is used to transform verbs into 'modified verbs', that are used to describe result states of the events described by the unmodified verbs and should therefore also be classified as belonging to a different aspectual category (that of the state verbs).

On this alternative view Present Perfects are just a variety of the (Simple) Present, to be characterized, if we are following Reichenbach here, as involving location of both R and E at the speech time S. In other words, the 'event' is located at S on this alternative analysis. But that is not inconsistent with the predictions Reichenbach makes – and with what is surely right – viz. that when I say, truly, 'I have written that letter.', then the event of writing the letter must have been in the past of the time of speech; for what counts as the 'described event' on the alternative view is not the writing itself but the state of having written the letter. And if that state holds at S, then the writing must have taken place before S.

If this criticism is correct, then it presumably does not only touch Reichenbach's analysis of the Present Perfect, but his proposals for the other perfect forms as well, including that for the Past Perfect. It might be thought that in this way the entire edifice of his account is called into question, since his analysis of the Past Perfect is the starting point for the account and seems its principal pillar. But actually that is not so. Investigations into the semantics of tense and aspect of the past half century have shown that Reichenbach's insight that Reference times other than the Speech time play a decisive part in the semantics of the tenses of English and other natural languages was right and that his Reference times are crucial to the correct analysis of tense. And their importance can be seen with particular clarity in connection with the Past Perfect: that is a form which undeniably requires that the Reference time be situated in the past of S, and thus that R must be distinct from S; on the other hand the Event time E is in the past of R, so R cannot be identified with E.

The remaining question is then what more should be said about this past Reference time. Reichenbach says that in the case of a Past Perfect the described event is located at a time (his 'Event time') in the past of the Reference time. The alternative analysis just mentioned says that the result state of that event holds at the Reference time with the consequence that the event itself must have been in the past of the Reference time. But arguably these are just two variants of one and the same underlying idea. Crucial, and common to the two of them, is that the Reference time must be in the past of S.

The objection that the perfect should be analysed as an aspectual operator and not as part of the tenses as such, is of a piece with a similar critique with regard to the various progressive tense forms of English. There is a wide consensus within semantics today that the progressive too should be treated as an aspectual operator, which applies to verbs or verb complexes and transforms them into new 'verbs' or verb complexes that belong to a different aspectual category; and such applications must take place before the resulting verb or verb complex is combined with its tense. We will come back to this issue – of how aspect is to be distinguished from tense – at length later on.

Reichenbach himself provides some concrete evidence for his analysis of the Past perfect and the Simple past by analysing a number of examples that involve applications of some of the entries in the table in (51). Two of these are given below in (52).

(52) a. I had mailed the letter when John came and told me the news.

b. He was healthier when I saw him than he is now.

The analyses Reichenbach gives of these multi-clause sentences pertains only to the temporal relations between the events described by their clauses. There is a good deal more to say about these sentences. Some of it has to do with the roles played by their different clauses – such as the role played by a *when*-clause that is used to locate the event of its main clause. That, as it turns out, is a quite complex set of issues, and much more complex than Reichenbach makes them out to be. A good part of the later half of the course will be concerned with such matters.

But the little that Reichenbach's analysis of these sentences says about them seems right nonetheless. (53) gives those analyses in much the same diagrammatic forms in which he presents them himself.

(53) a.

1st clause ('I had mailed ..')	E_1 (mail the letter)	R_1	S
2nd clause ('John came')		R_2 E_2 (John come)	S
3d clause ('John told ..')		R_3 E_3 (John tell me..)	S

b.

1st clause ('He was healthier')	R_1 E_1 (he be healthier)	S
2nd clause ('I saw him')	R_2 E_2 (I see him)	S
3d clause ('He is now')		S R_3 E_3 (He be now)

As far as the temporal relations between the events described in the different clauses of these two sentences are concerned these diagrams seem quite right. But from the perspective of current theoretical linguistics there is an obvious question to ask: How does one arrive at these diagrams starting from the forms of the sentences they represent. This is the general question of 'formal semantics':

(54) How can the meaning of a sentence be computed from its form?

It is a question about which Reichenbach has nothing to say in the excerpt from his book his *Elements of Symbolic Logic*that we have been reading, and in fact it is one that came into proper focus only two decades after the appearance of that book.

It is nonetheless clear from what Reichenbach says and from what he is implying that part of what has to be done to compute the kinds of interpretations that are shown in these diagrams is to determine past Reference times (for those tenses occurring in the represented sentences whose interpretation involves an R in the past of S), is to identify these Reference times in other parts of the given sentence, or in the antecedent text or discourse. As we have seen, Reichenbach discusses this dimension of tense interpretation in some detail only for the Past Perfect. And the first example he gives (the quotation from Somerset-Maugham) shows how the interpretation of a Past Perfect in a new sentence may require us to screen the preceding sentence (or sentences?) for a suitable time that can serve as its Reference time. As we will see, one of the important functions of the Past Perfect and some other tenses is to locate the events or states described in their clauses temporally within the time structure that has already been put in place through interpretation of the preceding discourse or text, including other sentences than the one in which the given tense occurs. (The examples (52.a,b) are both single sentences, so that all tenses that are represented in the diagram belong to the same sentence. But as Reichenbach's first example and his discussion of it make plain, whatever the principles of Reference time recovery may be, their application is evidently not limited to single sentences, which can be interpreted as expressing their complete content all by themselves.)

Reichenbach's lesson about the role of the Reference time in the interpretation of the Past Perfect seems clear enough. But as soon as we accept it there is another question that we cannot help asking when we look at the diagrams in (53) and compare these with the table in (51). Four of the six tenses in the two sentences (52.a,b) are Simple Pasts. According to the table in (51) these too involve past Reference times. But how are these times established as part of the interpretation of the Simple Pasts in (52.a,b)?

Here is a first intuition as to what an answer to this question should involve. Three of the Simple Pasts in the two sentences occur in *when*-clauses and one in a main clause. And that difference appears to be important. The function of a *when*-clause typically seems to be that of providing a temporal location for the event described by the clause to which the *when*-clause is attached. If that is right, then presumably the interpretation of the Simple Past of a *when*-clause should not rely on a Reference time that comes from some other clause: *when*-clauses establish their own Reference times – that is part of their point!

For the Simple Past in the main clause of (52.b) the correct story looks to be a different one. If what was just surmised about the function of *when*-clauses is right, then the interpretation of the main clause tense of (52.b) should depend on the temporal information provided by the *when*-clause; and that could arguably take the form of setting the Reference time for the main clause tense equal to the Reference time of the *when*-clause. (Or to its Event time: in the present instance, where Reference time and Event time of the *when*-clause coincide, the outcome is the same, but the choice between the two options may ultimately prove to be of theoretical importance.)

At this point our discussion of these issues can only be tentative. But it is, I think, already becoming visible that the complexity in tense interpretation that Reichenbach appears to have been the first to have seen clearly for what it was, and which he made explicit through his introduction of the concept of Reference time, has much more to it than can be found in the few pages that he ever devoted to the subject in print.

To conclude these comments on Reichenbach, there is a terminological convention that needs to be made at some point in any extended discussion of tense and aspect. Vendler distinguishes between event verbs (the accomplishment, achievement and activity verbs) and state verbs. As the way I put things in this last sentence implies verbs of the former three categories are used to describe events, whereas those of the fourth category are used to describe states. As we proceed, we will see that the distinction between events and states is all-important and that state-descriptions play a much bigger part in the linguistic realization of tense and aspect than Vendler's verb classification may well suggest. But Reichenbach's discussion of tense interpretation on the one hand and his examples on the other suggest that as far as tense interpretation goes, event descriptions and state descriptions function in quite similar ways. Compare for instance (52.a) and (52.b). The clauses of the former describe events, whereas the clauses of the latter describe states. But one of the points that Reichenbach's treatment of the tenses of these two sentences can be taken to make – whether this was part of his intentions I do not know, but for the present argument that doesn't matter – is that to the interpretation of tenses the state-event distinction *is* irrelevant. We will later see that this isn't quite true. But what remains true is that in many contexts the distinction is irrelevant. And in situations where that is so, it is convenient to have a single word that covers both events and states, so that one doesn't have to resort all the time to formulations like 'events and/or states'.

Unfortunately there just doesn't seem to be a good word with this meaning in English. So there was nothing for it but to coin one. The word now generally used for this purpose in linguistics is *eventuality*. (The term was first proposed by Emmon Bach, some of whose work will be discussed later.) The choice of 'eventuality' is perhaps not optimal. But no one has been able to come up with something better and by now the use of 'eventuality' to mean 'event or state' has become so much a part of the field that it would probably be counterproductive to change the term even if someone could come up with one that has a more suitable ring to it. So, 'eventuality' it shall be!

4.3 Bach: The Algebra of Events

Like Vendler's paper, Bach's paper is about Aspect. But as we will see, it takes a quite different line.

The central point of the paper is to draw attention to an important analogy between the 'nominal domain' and the 'verbal domain' – or, of you prefer, between the semantics of nouns and noun phrases on the one hand and that of verbs and verb phrases on the other. To give a first, but telling impression of this analogy, here are Bach's own examples.

(55) a. i. Much mud was in evidence.
ii. * Much dog was in evidence.

b. i. John slept a lot last night.
ii. * John found a unicorn a lot last night.

c. i. Many dogs were in the yard.
ii. * Many muds were on the floor.

d. i. John fell asleep at last three times during the night.
ii. * John slept at least three times during the night.

We start with a critical exposition of Link's account of the semantic differences and relations between singular count nouns, plural count nouns and mass nouns. Although this will lead us somewhat away from the central topic of this course, the matter is of considerable interest in its own right, and it is also important to have a good grasp of it in order to appreciate what Bach has to say about the corresponding aspectual distinction.

4.3.1 Link's Semantics for Count Nouns and Mass Nouns

A. The syntax of nouns and noun phrases

English and other European languages have two kinds of nouns, *count nouns* and *mass nouns*. Examples:

Count Nouns: *cat*, *flower*, *pebble*, *table*, *thing*, *part*, *universe*

Mass Nouns: *water*, *air*, *milk*, *butter*, *bread*, *sand*, *furniture*

Main syntactic differences between count nouns and mass nouns:

1. Count nouns can be pluralized: *cats*, *flowers*, *pebbles*, *tables*, *things*, *parts*, *universes*

   Mass nouns cannot: # *waters*, * *airs*, * *milks*, * *butters*, * *breads*

2. Mass nouns can occur on their own as argument phrases, e.g. as grammatical subjects or direct objects or arguments of prepositions, as in (56).

(56) a. Milk has become expensive.
 b. We have to buy butter.
 c. You can't live without air; and you can't live without water.
 d. Bread is on the table.

Count nouns cannot be used in this way, though their plurals can:

(57) a. * Cat was prowling the neighborhood.
 b. √ Cats were prowling the neighborhood.
 c. * I like flower.
 d. √ I like flowers.
 e. * Universe is very big.
 f. √ Universes are very big.

3. Count nouns can be used with the indefinite article *a* and with *some*. Mass nouns cannot:

$\surd$ *a cat/a flower/a part/a universe*; * *a milk/an air/a bread.*

3. The positive adjectives *much* and the comparative form *less* go with mass nouns but not with count nouns. *many* and *few/fewer/fewest* go with (plurals of) count nouns:

(58) a. $\surd$ Much water/butter/bread/sand/furniture
  b. * Many water/butter/bread/sand/furniture
  c. $\surd$ Many cats/flowers/parts/universes
  d. * Much cats/flowers/parts/universes

N.B. There are quite a few nouns that can be used straightforwardly and felicitously both as count nouns and as mass nouns. Examples are *cheese*, *brick* or *marble*. There are also more general processes for turning count nouns into mass nouns – a famous, if not particularly likeable, example in the literature is the one quoted by Bach: 'There was dog splattered all over the road' – and for turning mass nouns into count nouns. We will come to those below.

Singular count nouns, we saw, cannot occur as argument phrases on their own. In order to make argument phrases out of them they have to be combined, minimally, with an article (*the*, *a*) or some other *Determiner*: *some*, *this*, *that*, *any*, *every*, *no*. It is the complex phrases resulting in this way – phrases like *a cat*, *the table*, *this flower*, *her dog*, *their friend*, *no part*, *every universe* and so on – that are combined, syntactically and semantically, with the predicates whose argument positions they fill.

Note well, however: argument phrases, and this is true in particular for those that involve count nouns, can be a good deal more complex than the examples above. They can involve adjectives, prepositional phrases, prenominal genitives and relative clauses, and in any combination:

(59) *a beautiful cat*; *a huge, loveable longhaired cat*; *every flower in this vase*; *no table that I have ever worked at*; *this mystery-filled universe that we live in*; *the neighbor's cat*; *the biggest cat that anybody ever saw.*

In all examples mentioned in (59) and in the paragraph above it the count noun is the core, so to speak, from which the entire phrase is constructed; and it is also its semantic pivot. We refer to it as the *lexical head of* the argument phrase.

Besides argument phrases that are built from a lexical head noun there are also argument phrases that are very short in that they consist of just a single word. Among these are the ones that consist just of a single mass noun such as *butter* in *There is butter in the fridge* or *Butter is on the table.* But there are others as well, viz. proper names and pronouns: *Jack*, *Obama*, *I*, *you*, *he*, *she*, *it*. All these can occur in argument positions (such as that of a sentence subject) all by themselves. For many theories of the syntactic structure of nominal expressions these one word argument phrases pose special problems. But this is a matter that is of no direct relevance here, so we let it rest.

N.B. Many of the so-called *generative* approaches to the syntax of nominal constructions distinguish between (at least) two *projection levels* above that of the lexical head noun: that of the *N(oun) P(hrase)* and that of the *D(eterminer) P(hrase)*. What we have so far been referring to loosely as 'Noun phrases' and also as 'argument phrases', are really DPs according to this technical terminology. (It is DPs, and only DPs, that can occur in argument positions like that of the grammatical subject of a sentence.) 'NP', in the technical sense we now adopt, is used to refer to an intermediate level in the construction of full DPs, at which the lexical head noun has already been combined with adjectives, genitives, Prepositional Phrases and subordinate clauses, but at which combination with the Determiner has not yet taken place. So the structure of complex DPs is something like this:

(60)

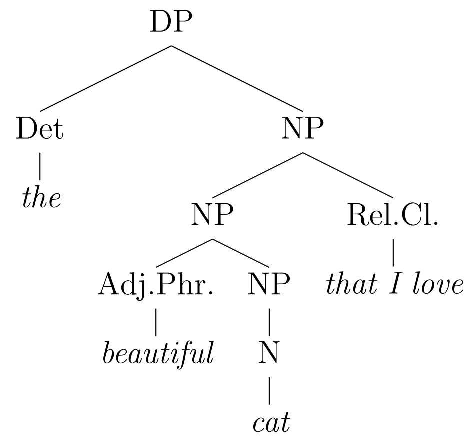

(A number of current generative theories of noun phrase structure assume even more projection levels than these two. Those, however, will not be directly relevant to us.)

B. Semantics of nouns and noun phrases

From a logical and semantic point of view most count nouns function as *1-place predicates*: as expressions that are true of some things and not of others. For instance some things are cats and others are not, and so on.

The standard way to describe this distinction between what a 1-place predicate is true of and what it is not true of (or between what *falls under* the predicate and what doesn't) is in terms of its *extension*:

(61) The *extension* of a 1-place predicate P is the set of all things that P is true of.

Thus the extension of *cat* is the set of all cats, the extension of *flower* is the set of all flowers etc.

Extensions aren't immutable: the extension of *flower* is changing all the time – flowers come and go. So does the extension of *cat* and of the vast majority of count nouns. Also, extensions may vary between possible worlds: There could have been cats that do not exist in our world, and that in fact never did nor will exist here. Thus the extension of *cat* in other possible worlds may be expected to be different from what it is in our world, at this or any other time. To account for these dimensions of variation, semanticists assume that count nouns and other 1-place predicates come with extensions for any combination of a possible world and a time in this world; or, put more formally, that such expressions determine *functions* from possible world-time pairs $< w, t >$ to extensions. Such functions are called *intensions*. Thus the intension of a predicate is a function from worlds and times to extensions – or, in view of what sorts of things the extensions of predicates are, a function from worlds and times to sets of things. In what follows I will ignore intensions, however, and only speak of extensions. (We are making the implicit assumption that a particular world and time have been fixed.)

How does the extension of a count noun that is the lexical head of a DP enter into the semantics of the DP? That depends on the DP's Determiner. Some Determiners – *the*, *this*, *that* – make 'referring phrases' (also: 'referring terms') out of the NPs they get as inputs – phrases that refer to some particular element from the head noun's extension; for illustration see the examples below.

(62) a. The neighbors' cat is always in heat.

b. That Greek student in your class seems very smart.

c. This novel by Eliot is much better than that one.

d. The lecture that he gave to the Aristotelian Society last year was the best I ever heard him give.

The examples all illustrate the same point: the DPs in question pick out one element from the extensions of their lexical heads – a cat, a Greek student, a book, a lecture – and that thing is said to satisfy the predicate – very smart, always in heat, better than something else, the best of its kind – of which the DP is an argument phrase.

Other Determiners do different things with the semantics of the phrases with which they combine. *every* for instance conveys that all members of the extension of the phrase with which it combines satisfy the predicate containing the DP as argument phrase, *some* and *a* convey that some member satisfies it, and *no* that there isn't any member that does. Exactly what the different Determiners do, how they do it, and when their use is appropriate, are among the cenral questions of modern semantics. But what matters for the present discussion is only that count nouns have extensions, and that it is (just) their extensions they contribute to the semantics of those sentences containing them with which we will be concerned in what follows.

Mass nouns too have extensions and here too it is its extension that a mass noun contributes to the semantics of the sentences that contain DPs of which the mass noun is the lexical head. (Mass nouns too can occur as heads of larger DPs, such as *the water in the fridge*, *this butter on the plate here*, *no decent furniture*, besides forming DPs on their own.) But the extensions of mass nouns are different from those of count nouns in that they consist of 'stuff', rather than being collections of things. And 'stuff', in the sense intended here, is something *homogeneous*: when you put two bits of any given kind of stuff together you get again a (somewhat bigger) bit of that same stuff, and when you have a bit of the stuff you can always divide it into smaller bits and what you are then left with are again (smaller) bits of the same stuff.

This, at any rate, seems to be how we conceive of the stuff described by mass nouns. Arguably that is a kind of conceptual idealization. For in most cases we know perfectly well that the stuff described by our mass nouns is *not* infinitely divisible. This is very plain for nouns like *furniture* or *clothing*. Here we can usually see the smallest parts – the individual pieces of furniture or clothing – with our own eyes. But in a way it is also true for the extension of nouns like *salt*, or *sugar* or *sand*. We know that the stuff we are looking at consists of smallest grains and sometimes we can see those smallest bits too, perhaps even with the naked eye. Of course we also know that when we would put those under the sledge hammer we would still be getting just more stuff of the same kind, only pounded to a finer degree of granularity. But we also know that this cannot be carried on ad infinitum. All matter is composed of smallest parts, atoms or molecules. This much physics and chemistry seem to have been established beyond all reasonable doubt and is wisdom that has made its way into the intellectual baggage of every reasonably well educated adult in our culture. But as far as I can see, that doesn't alter anything to the fact that when we use mass nouns, or reflect on their meaning, we hit upon homogeneity as an essential part of our conception of them. So I conclude:

(63) (The conceptual idealization inherent in our use and understanding of

mass nouns)

Normal adult speakers of English (and other languages with a grammaticized mass-count distinction) understand the extensions of mass nouns as *homogeneous* (in the sense of being *additive* and *infinitely divisible*).

(N.B. There are languages without the grammatical distinction between count nouns and mass nouns that we find in English. Many of the languages in the Far East have this property, including Mandarin Chinese, Japanese and Korean, languages that at this point in time are among those whose grammatical properties have been closely investigated using the sophisticated tools and methods of current theoretical linguistics. Such languages, which lack the general count noun-mass noun distinction found in English and similar languages are known as *classifier languages*. In classifier languages, you might say, all nouns, except for a small handful, behave grammatically in essence like our mass nouns. And the handful, the so-called 'classifiers', are used to turn these mass nouns into count nouns. In this regard they are like our nouns *bit*, *heap*, *pile*, *portion* and also, more specialized, *chunk*, *slice*, *draft*, *glass*, *pint*, as they occur in phrases like *bit of trouble*, *portion of ice cream*, *heap of sand*, *pile of sugar*, *slice of bread*, *chunk of meat*, *glass of milk*, *pint of lager*. In fact, these last expressions function in English as 'compound count noun phrases'. In classifier languages *all* count nouns are of such a compound form. What this implies about the conception that speakers of these languages connect with the semantics of nouns, and in particular whether their semantic intuitions differ from ours in relation to nouns that for us are count nouns, is a question that wasn't investigated until quite recently. And it appears to be still too early to draw any firm general conclusions from these investigations.

Count nouns, we noted in passing, can be used in the plural as well as in the singular (whereas mass nouns can only occur in the singular). The plural forms of count nouns are often used to talk about collections of the things that make up the extension of the singular count noun. We find such reference to collections in particular with definite plural noun phrases beginning with *the*, *these*, *those*, as in (64).

(64) a. The students who got an A all are in Grad School.

b. Those students (I've marked them in red) all are in Grad School.

c. These students will all go to Grad School.

The DPs *the students who got an A*, *those students (I've marked them in red)*, *these students* are referring phrases just as their singular counterparts *the student who got an A*, *that student (I've marked her in red)*, *this student*. The difference is that while the latter each refer to a single student, the former must refer to a collection of two or more. One way in which this difference is often expressed is that the singular phrases refer to *atomic* parts of the extension of their head noun, the plural phrases refer to *non-atomic* part of its extension.

(N.B. That the plural DPs in (64) really do *refer* to collections of two or more elements and that they don't do something else with their head nouns' extensions (such as quantifying over them in some way) is shown even more dramatically by the fact that they can occur in argument positions of *collective predicates* such as the subject positions of sentences involving the intransitive verb *gather*. Their singular counterparts cannot fill those positions:

(65) a. $\surd$ The/those/these students gathered in the courtyard.

b. * The/that/this student gathered in the courtyard.

You cannot say of a single student that he or she 'gathered'. And for the same reason 'The students gathered in the courtyard.' cannot be paraphrased as 'Each of the students gathered in the courtyard.' The only correct analysis of this sentence is that 'gathered' is true of the set, or collection, ferreted to by the phrase *the students*.

Singular referring phrases – in particular: DPs with a singular count noun as head – refer to single members of the extensions of their head nouns. But what about plural referring phrases such as *the/those/these students*? One way to answer this question might be this: plural referring phrases also refer to members of the extensions of their head nouns. However, the extension of a plural count noun is different from that of its singular count noun counterpart. The extension of the plural count noun consists of all sets of two or more elements that make up the extension of the singular count noun. (Thus the extension of *students* consists of all sets that consist of two or more students.) In the semantics literature the transition from a singular count noun to the corresponding plural is often indicated by the superscript

$^+$. Thus if P_c is the count noun *cat*, then P_c^+ is the predicate *cats*, and so on. With this notation we can define the above proposal for how the extensions of plural nouns are related to the extensions of the corresponding singular nouns as in (66.a).

(66) (Definition of the *extension* $X^+_{P_c}$ *of* 'P_c^+' and the *extension* $X^*_{P_c}$ *of* 'P_c^*' in terms of the extension X_{P_c} of P_c)

a. $X^+_{P_c} = \{Y : Y \subseteq X_{P_c} \ \& \ |Y| \geq 2\}$

b. $X^*_{P_c} = X^+_{P_c} \cup X_{P_c}$

The literature on the semantics of singular and plural noun phrases also makes use of another, closely related operator that is represented as the superscript *. This operation turns a singular count noun predicate into a predicate that is semantically like the union of the extension itself and its pluralization. (Thus, if 'P_c' is the count noun *cat*, then'P_c^*' can be read as 'cat or cats'.) The definition of $X^*_{P_c}$ is given in (66.b).

C. Mereology vs. Set Theory

One of the important insights of Link(1983) was that it is informative to put the *-extensions of count nouns side by side with the extensions of mass nouns. As argued above, both satisfy additivity – when you add two parts of the extension what you get is again a part of the extension – and they also both satisfy some kind of divisibility principle. The only difference is that for mass nouns divisibility is conceived of as holding unrestrictedly, whereas for count nouns its application is limited. More specifically, those parts of $X^*_{P_c}$ that belong to X_{P_c} (i.e. the *atoms* in $X^*_{P_c}$) cannot be divided into smaller parts of the extension; and larger parts, those that belong to $X^+_{P_c}$, can only be divided in such a way that the portions consist entirely of one or more complete atoms.

But in all other respects the *-extensions of count nouns and the extensions of mass nouns seem to have much the same structure, so much so, Link saw, that it is illuminating to amalgamate them into a single algebraic structure and study that structure. However, to bring out the essential similarity more clearly it is helpful to move from the set-theoretic perspective that is dominant today (especially in mathematics, but also in applications to linguistics and philosophy) to a *mereological* perspective. The central topic of mereology is the relation between parts and wholes: Mereology studies the part-whole relation $\sqsubseteq$, whereas Set Theory studies the membership relation $\in$ between sets and their elements. The two theories target the same general subject matter. But there are important differences between the ways in which they view and treat their subject. The most important one of these is arguably that in Mereology parts and wholes are entities of the same 'level', or the same 'logical type'. Set Theory differs on this point. There is a strong intuition here that a set belongs to a *higher* level than its elements and that for each set there are sets of yet higher levels than it, of which the given set could be an element in its turn. This leads to an extremely complex structure of levels. All of that is absent in Mereology, because the assumption that gives rise to it – sets are of a different level than their elements – isn't made here.

In the present context this difference becomes tangible in connection with our definition of $X^*_{P_c}$. As defined in (66.b), $X^*_{P_c}$ is a set-theoretically mixed bag, consisting of the elements of X_{P_c} on the one hand and sets of them (the elements of $X^+_{P_c}$) on the other. This isn't all that much of a problem formally. But it does suggest something that Link was at pains to counteract: that the denotations of plural DPs are of a different logical type than the denotations of singular DPs. Formally we can restore type identity easily enough, viz. by defining $X^*_{P_c}$ as consisting of sets only, replacing the elements of X_{P_c} by their singleton sets. (A *singleton set* is a set with exactly one element, and when x is that element, then its singleton set is denoted as $\{x\}$.) In fact, this is a ploy that quite a few semanticists working within a set-theoretical framework have adopted. But it doesn't feel quite right either, since there is a strong intuition that what the singular referring DPs really refer to are the objects themselves, *not* their singleton sets.

The point of Mereology is that this awkward choice between two options, neither of which seems to give us all we might want, shouldn't arise in the first place. According to Mereology, the distinction between objects and their singleton sets is a bogus distinction; there should be no temptation to make it. All there is is the part-whole relation. That relation can hold between atomic and non-atomic elements, where the atomic elements are just those to which nothing else stands in the part-whole relation, and the non-atomic ones are all the others. (In mereological treatments of singulars and plurals the atoms are often referred to as 'entities', or 'objects' or 'individuals', and the non-atoms as 'pluralities'.)

From the point of view of Mereology, then, the extension of P_c^* is a structure $< D, \sqsubseteq >$, which, according to standard mathematical terminology, is an *upper semi-lattice* and which, moreover, has the property of being *atomic*. Here are the formal definitions of these notions.

(67) (Definition of 'upper semi-lattice', of 'atom of an upper semi-lattice' and of 'atomic upper semi-lattice')

Let $<D, \sqsubseteq>$ be a relational structure, consisting of the non-empty set D and the binary relation $\sqsubseteq$ on D.

a. $<D, \sqsubseteq>$ is an *upper semi-lattice* iff the following conditions are fulfilled:

(i) For all a in D, $a \sqsubseteq a$ (Reflexivity of $\sqsubseteq$)

(ii) For all a, b, c in D, if $a \sqsubseteq b$ and $b \sqsubseteq c$, then $a \sqsubseteq c$ (Transitivity of $\sqsubseteq$)

(iii) For all a, b in D, if $a \sqsubseteq b$ and $b \sqsubseteq a$, then $a = b$ (Antisymmetry of $\sqsubseteq$)

(iv) For all a, b in D there is a c in D such that
(a) $a \sqsubseteq c$,
(b) $b \sqsubseteq c$ and
(c) if d is an element of D such that $a \sqsubseteq d$ and $b \sqsubseteq d$, then $c \sqsubseteq d$.

b. Let $\mathcal{D} = <D, \sqsubseteq>$ be an upper semi-lattice, a an element of D. a is an *atom of* $\mathcal{D}$ iff for no b in D such that $b \neq a$, $b \sqsubseteq a$.

c. Let $\mathcal{D}$ be as under (b). $\mathcal{D}$ is *atomic* iff for all b, c in D: if it is not the case that $b \sqsubseteq c$, then there is an atom a of D such that $a \sqsubseteq b$ and it is not the case that $a \sqsubseteq c$.

Fact: Consider $X^*_{P_c}$ for any count predicate P_c. Define $\sqsubseteq_{P_c}$ as follows:

for any a, b in $X^*_{P_c}$, $a \sqsubseteq_{P_c} b$ iff either (i) both a and b belong to $X^+_{P_c}$ and $a \subseteq b$ or (ii) a belongs to X_{P_c}, b belongs to $X^+_{P_c}$ and $a \in b$. (in (i) $\subseteq$ is set-theoretic inclusion.)

Then $< X^*_{P_c}, \sqsubseteq_{P_c}>$ is an atomic upper semi-lattice and X_{P_c} is its set of atoms.

Exercise: Show this!

As defined, upper semi-lattices only involve a part-whole relation but no additivity operation. But in a way this is appearance only, for each upper semi-lattice contains an additivity operation implicitly: Within an upper semi-lattice we can define an additivity operator $\cup$ as in (68):

(68) For any a, b, c in D, $a \cup b =$ c iff$_{def}$ $a \sqsubseteq c$, $b \sqsubseteq c$ and for any d in D, if $a \sqsubseteq d$ and $b \sqsubseteq d$, then $c \sqsubseteq d$.

Note that this definition is correct because condition (iv) of Def. (67.a) guarantees that a c that is related to a and b as stated in the right hand side of (68) always exists and is unique. (The uniqueness isn't part of our definition of upper semi-lattices but can be proved from the defining conditions for an upper semi-lattice given in (67.a). Show this!)

In virtue of the possibility of defining $\cup$ as in (68) we could just as well have defined upper semi-lattices as structures $<D, \sqsubseteq, \cup>$, making the relation between $\sqsubseteq$ and $\cup$ that is expressed by the right hand side of (68) one of the defining conditions of such structures. In the sequel we will allow ourselves to speak of upper semi-lattices either as structures of the form $<D, \sqsubseteq>$ or as structures of the form $<D, \sqsubseteq, \cup>$. Either way it will be understood that the explicit or implicit additivity operation is related to $\sqsubseteq$ in the manner of (68).

(N.B. Atomic upper semi-lattices are much like the Boolean algebras defined in Section 3 of these Notes: Every atomic upper semi-lattice can be extended and expanded in a unique way to an atomic Boolean Algebra with the same set of atoms. And conversely we can obtain an atomic upper semi-lattice from any atomic Boolean algebra $<D, \cap, \cup, ^-, 1, 0>$ by ignoring the operators $\cap$ and $^-$ and the reference to the special elements 1 and 0 and by defining

$\sqsubseteq$ by: $a \sqsubseteq b$ iff$_{def}$ $a \cup b = b$.)

The algebras we talked about when we introduced the notion of a Boolean algebra and the algebraic structures that are at issue now also have another feature in common (or, more exactly, there is a close similarity between the Boolean algebras we discussed earlier and the upper semi-lattices we would get if we were to follow the set-theoretic approach in which individuals are replaced by their singleton sets): In both instances, the domain D of the algebra consists of sets of points drawn from some antecedently given set X. In the case of the Boolean algebras discussed earlier these 'points' were possible worlds, X being the totality W of all possible worlds (and the relevant sets were propositions). In the case before us the points are the atomic members of the extension of a given count predicate P_c (and in which X is thus the extension X_{P_c} of that predicate), and the sets are the pluralities formed from those atoms together with the singleton sets of the atoms. Note well, however, that in general the domains of algebraic structures such as (semi)-lattices or Boolean algebras need not consist of sets. Many of the most interesting and important applications of the theory of lattices and Boolean algebras is to structures that are not of such a set-theoretic form.)

D. Mereology of Count and Mass

So much then for the *-extensions of count predicates. From a mereological point of view they are atomic upper semi-lattices and that is how we will treat them here.

Link observed that the extensions of mass nouns can also be thought of as upper semi-lattices, but the difference with the *-extensions of count nouns is that the upper semi-lattices of mass nouns are not atomic. More precisely, Link proposes that we can think of the extension of a mass noun P_m as consisting of all possible portions of stuff of the kind described by P_m. These portions can stand in part-whole relations just as the elements of the *-extension $X^*_{P_c}$ of a count noun, and here too that relation is strictly speaking the only structure-inducing relation we need.

In fact, it follows from what we said earlier about the homogeneity of 'stuff' that these extensions are the opposite of atomic in a strong sense: they do not have any atoms at all. The following definition makes this formally explicit.

(69) An upper semi-lattice $< D, \sqsubseteq >$ is *atomless* iff no element a of D is an atom of $< D, \sqsubseteq >$.

Before we move to the next part of Link's proposal, let us reflect a little more on what the extension of a mass predicate P_m must be like if it is to consist of all portions of the stuff described by P_m. As noted earlier, if this collection of portions is to satisfy the additivity and infinite divisibility constraints – that is, if it is to be closed under these operations – then it will have to be large: with each actual portion, such as the lump of butter on my plate, say, there must be infinitely many smaller portions that could be obtained from it by division; and on the other hand additivity will, when taken in its strict generality, entail the existence of disconnected portions with the oddest distribution in space. But even if we just consider the lump of butter in front of me and try to contemplate all the portions of butter that could be obtained from it by division the mind soon boggles. Let us focus on just one way of cutting the lump: by making a strictly vertical cut somewhere. How many such ways of cutting, and corresponding smaller portions of butter, are there? Well, that depends on various factors and also on various assumptions we are prepared to make. One assumption we could make is that there is a certain area of the plate with which the butter is in contact. Let us consider only cuts that intersect that part of the boundary of the lump that is in contact with this area. How many (strictly vertical) cuts are there which satisfy this condition? In answering that question we must rely on some conception of the structure of space – of what horizontal areas and straight lines intersecting them are like: Given some finite part of a horizontal surface, how many different straight lines are there that intersect it? Arguably there isn't just one such conception. But mathematics and physics point strongly in one particular direction: That space is dense and continuous (in the same sense as we said this is normally assumed for physical time, but in three dimensions rather than one). On this assumption there will be uncountably many cuts that can be made in the manner described, and thus uncountably many different portions of butter that can be obtained in just this way.

That is a lot of butter portions. And yet it is only a tiny fraction, somewhat informally speaking, of the totality of actual and possible portions of butter that the universe must be assumed to contain. (And then think of portions of hydrogen, which doesn't exist just on Earth – something that for all we know may well be the case for butter – but that can be found pretty much anywhere in our universe.) You might object that all this is just based on an abstract mathematical conception of space the physical reality of which is questionable. But note that in the present context this kind of objection doesn't seem to cut much ice. For as we saw, the homogeneity assumption about the extensions of mass nouns is an idealization in the first place. It is hard to see why assumptions about the geometry of space like the one we have just made and the conclusions about numbers of possible cuts we have drawn from that shouldn't be legitimate parts of that conceptual package, especially since it is unclear what other packages are available.

These scant reflections do not settle the matter of course. But at least they enable us to say this much: Given the idealized conception of the extensions of mass nouns as closed under additivity and division that Link advocates, an elaboration of that idealization according to which there are uncountably many portions of the stuff described by most or all mass nouns is perhaps the most natural one we can hope for.

If this is the elaboration of Link's homogeneity assumption that we adopt, then we are faced with a conclusion that may seem surprising: The extensions of mase nouns will always be immeasurably bigger than the *-extensions of count nouns. More specifically, it is not unreasonable that the extension of any count noun is finite. Even of protons, neutrinos and other subatomic particles there can, as received physical theory has it, only be finitely many in our universe. And there are of course far fewer cats or tables than there are protons or neutrinos. But if the extension of any count noun is finite, then so is its *-extension. (It will be much bigger, but its seize will be finite no less.) So there is, on the present conception of count noun and mass noun extensions, no comparison between, say, the number of cats in this world and the number of portions of the stuff that cats are made of.

These considerations are of no direct consequence to the points that Link wants to make about the relation between things and stuff. But they are worth keeping in the back of our mind when contemplating the next move that Link makes. This next move can be seen as the decisive point in and of Link's development. It consists in amalgamating the extensions of mass nouns and the *-extensions of count nouns into a single algebraic structure.

More accurately, Link's idea was to build a single algebraic structure out of the extensions of many count nouns and many mass nouns all at once. Exactly which, or how many, neither he nor Bach ever explicitly discuss. It is even possible to take Link as wanting to present an algebraic model in which all the things and all the stuff in the world are put together, so that the extensions and *-extensions of all mass nouns and all count nouns are subsets of the domain of this structure. I will argue below that this simply cannot be done. The only coherent ways of making sense of the amalgamations Link proposes involve a careful choice of the count and mass nouns whose extensions are to be jointly represented in the structure. The choice has to be made with care because on the one hand it cannot be too big on pain of incoherence while on the other the choice cannot be too parsimonious or the point of the amalgamation would be lost. In particular one would, in order to do justice to Link's intentions, want to include in conjunction with each count noun one or more mass nouns that can be used to describe the stuff of which things falling under the count noun can be made up.

I will return to this issue below. But for now let us assume that there is some set $\mathcal{C}$ of count nouns and some set $\mathcal{M}$ of mass nouns such that the amalgamated structure to be defined below is made up of the extensions of the predicates in $\mathcal{M}$ and the *-extensions of the predicates in $\mathcal{C}$.

E. Amalgamating the mereological structures of different predicates

How do we define the amalgamation of the extensions and *-extensions of the predicates in $\mathcal{M}$ and $\mathcal{C}$? It seems there is a range of options here, all of which would serve Link's purposes, stretching from the parsimonious to the maximally prolix. My personal preferences point towards the parsimonious. But the option that seems to correspond most closely to the version Link adopted lies at the opposite end of the spectrum.

As I said, which of these options we go for seems to be immaterial given Link's concerns, and it definitely seems irrelevant to the applications that are mentioned by Bach. But no matter which option one decides on, there is one major conceptual feature that all of them share and that must be put into full view. It is especially important for us, because of its implications for the aspectual structure of the semantics of verbs, which is Bach's motivation for going into the count-mass distinction in the first place.

The best way to broach the issue is to consider one of the applications Bach brings up, that of Terry's ring and the gold it is made of. At any time t within the duration of the existence of Terry's ring there is the portion of gold of which the ring is made. That portion and the ring share certain properties. For instance they occupy, at any such time t, exactly the same region of space. But there are also properties in which they differ. One of these is their age. The ring was presumably made in the course of the last century (or perhaps one or two or even three centuries earlier; since I don't know much about either Terry or his ring, I am just guessing). But, as Bach literally puts it, 'the gold in Terry's ring is much older'. When we look more closely, we see that that formulation actually has a certain ambivalence to it. What is 'the gold that Terry's ring is made of'? Are we right to identify it with the portion of gold that consists of all and only that gold that is part of Terry's ring at the time t under consideration? I suppose that that is what is intended. And if on this interpretation it is true to say that the gold in Terry's ring is surely much older than Terry's ring, then that seems to entail that this *portion* of gold goes way back, and way beyond the point in time when the ring was made.

But what justifies us in saying this? Did this *portion* of gold exist for all this time? After all, it may well be that the gold in Terry's ring was brought together only when the ring was made. Or it could have been taken, at the point when the ring was made, from some larger lump of gold. Well, you might say, we have already made sure that *this* isn't going to be a problem when we assumed that at any one time the portions of gold would include all the possible portions that could be formed, through division and addition, from the portions available in more concrete terms (e.g. as cohesive lumps or nuggets). So at any time at which all the gold that is now in Terry's ring existed somewhere the *portion* of the gold in Terry's ring existed as well. (There is still a question how we determine the age of a portion of gold given that, presumably, the gold atoms that make it up weren't all created at the same time. Arguably that doesn't matter for the issue that Terry's ring and its gold raise for Bach. For it would seem safe to maintain, given what we have assumed about what portions of stuff there are at any one point in time, that the age of a portion of gold isn't less than that of the youngest atom in it; and it also seems safe to hold that by that criterion the portion of gold that makes up Terry's ring must be much older than the ring itself. Another complication, which Bach doesn't mention and I am also setting aside, is that the gold that makes up the ring will in general vary over time. On the one hand some of the gold will get lost through wear; on the other hand some may be added during occasional repairs, for instance when the ring had to be widened because Terry's finger had grown too stout.)

This much about the less problematic and less important one of the two points that the example of Terry's ring puts into focus. The other, more important issue is this: What really *is* the difference between Terry's ring and the portion of gold that it is made up of? What kind of difference is it? Are there really two different things in the world out there? Or is it a difference solely in the eye of the beholder – is it just a matter of the way I choose to look at what is lying in front of me, or of the way in which I choose to describe it, whether it is Terry's ring I am thinking or talking about or the portion of its gold?

This is a hard question. Part of the problem is that what we have been saying about the extensions of mass nouns is shaped by massive amounts of idealization already. The proposal was to treat the extensions of mass nouns as huge numbers of potential as well as actual portions of the stuff any given mass noun describes – portions that for the most part 'exist' only as the results of potential divisions within a continuum (or, perhaps, a collection of disconnected continua). Some of these portions (but it is only a vanishing small subset of them) are 'actual' in that they are bounded by a surface that demarcates where they end and the rest of the world begins, such as the lump of butter on my plate, the surface of which forms the boundary between it and the surrounding air and plate. And some of these actual portions coincide with elements from the extension or *-extension of some count predicate. The piece of gold that makes up Terry's ring is one of those.

But perhaps the idealization involved in this way of treating the extensions of mass nouns is also part of the answer. If it is true that in characterizing these extensions in this way we have captured something that is real and essential to the way that the users of mass nouns *understand* what it is they are talking about when they use those words, then the case for an ontology shaped by human conceptualization has already been made. This leaves it open whether the extensions of count nouns, and perhaps also their *-extensions, belong to an ontology that is there independently of us. Even if that part of the ontology is unaffected by conceptualization, the conceptual component in our thinking of what is described by our mass nouns may nevertheless thought to suffice as an explanation for why it is that two material entities – the ring and the gold it is made of, can occupy exactly the same place at the same time and yet be distinct. I leave the matter for further rumination.

One last point is needed before we can finally put count and mass noun extensions together. So far we have been assuming that count nouns have not only an extension but also, derivatively, a *-extension. Mass nouns, on the other hand, were assumed to have just an extension, ordered by the part-whole relation which tells us when one portion of stuff is included in another. But on Link's and Bach's approach that isn't all there is to the semantics of mass nouns. We can speak, and think, of two lumps of butter, one on my plate and the other in the fridge, And among the expressions we can use to do that I can use a plural noun phrase such as *the lump of butter on my plate and the lump of butter in the fridge*. This expression is closely reminiscent of conjunctive DPs involving count nouns, like *the avocado on my plate and the avocado in the fridge* – so much so that if the second is treated as referring to a plurality (the non-atomic entity made up of the two avocados), then the first should be treated in this way as well. If this suggestion is taken seriously and we swallow the consequences lock, stock and barrel, then we must acknowledge for mass nouns also *-extensions, on top of the extensions they have been assumed to have up to now. This turns the semantics for mass nouns into an even more complex and prolix thing than it was already. But that is the way that, it would seem, Link and Bach wanted to go.

Actually, as far as I can see, Bach, in his extremely terse presentation of Link's proposal, doesn't put things quite right when it comes to this point. This, as I understand it, is what the situation really is: The extension X_{P_m} of a mass predicate P_m is ordered by the part-whole relation between portions of stuff. In order that we do not get this relation mixed up with the part-whole relation that obtains between atomic and non-atomic elements of the upper semi-lattices we discussed when we introduced the *-extensions of count nouns, let us denote the part whole relation between portions of stuff as '$\sqsubseteq_m$'. When we now, as Link and Bach do, move from extension to

*-extension for mass nouns, then we end up with a structure involving two part-whole relations, the relation '$\sqsubseteq_m$' and the relation between atomic and non-atomic entities, which we will continue to denote as '$\sqsubseteq$'. Similarly, we must distinguish between the additivity operator that is specific to stuff and that has the effect of making one portion out of two smaller portions, and the additivity operator that forms non-atomic entities out of atomic and/or non-atomic ones and that is applicable both to the extensions of count nouns and (now also) to those of mass nouns. We keep these two additivity operators distinct by the same notational device we are using for the part-whole relations, marking the first operator with a subscript $_m$ and leaving the second unmarked: '$\cup_m$' and '$\cup$'.

Clearly the relation $\sqsubseteq$ on the *-extension of a mass predicate P_m is very different from the relation $\sqsubseteq_m$ on P_m's extension. ($\sqsubseteq_m$ only holds between members of the extension of P_m, that is, between atoms of its *-extension, whereas the important instances of $\sqsubseteq$ involve non-atomic elements of the *-extension. In fact, the only overlap between $\sqsubseteq_m$ and $\sqsubseteq$ are the pairs $< a, a >$ for a in the extension of P_m.) But there is nevertheless an important connection between the two relations. Suppose for instance that a and b are portions of stuff described by P_m, and thus members of the extension of P_m, and that $a \cup b$ is the plurality made up from the two of them. Then there will also be the portion $a \cup_m b$ belonging to the extension of P_m as well; and this portion obviously stands in the material constitution relation to the plurality $a \cup b$: it consists of all and only the stuff that is part of this non-atomic element. We will return to this connection presently.

Summarizing this last discussion: According to Link and Bach the semantics of mass nouns involves a *-extension as well as an extension. And the *-extensions involve, just as the *-extensions of count nouns, the mereological structure of atoms and their pluralities. But in addition the extensions of mass nouns have their own mereological structure and that structure is mirrored, in a way that we are not yet able to express, the mereological structure of their *-extensions.

The amalgamation, I said, of the extensions and *-extensions of count and mass nouns can, even after all the assumptions we have now made, still take more than one form. I already put on record my own preference for parsimony on this point. In fact, given a set $\mathcal{C}$ of count predicates and a set $\mathcal{M}$ of mass predicates, the most parsimonious amalgamation into a joint structure, which gives us all we need to make the points that Link and Bach want to make, consists in forming, first, the Domain of the new structure as the union of all the *-extensions of the count nouns in $\mathcal{C}$ and the mass nouns in $\mathcal{M}$. Second, we have to determine the relations and operations of the new structure. Here we can let ourselves be guided by the intuition that the relations $\sqsubseteq_{P_c}$ for the different predicates in $\mathcal{C}$ are restrictions, to the *-extensions of these different P's, of one and the same mereological part-whole relation. To illustrate: you and I are both humans, but we are also both animals, since every human is an animal. So the mereological sum consisting of you and me will be both a plurality (= non-atom) belonging to the *-extension of the count noun *human* and a plurality belonging to the *-extension of the count noun *animal*; and the relation $\sqsubseteq_{human}$ which relates both you and me, as atomic members of the *-extension X^*_{human} of *human*, to the non-atomic member of X^*_{human} that consists of you and me coincides, for these and other elements of the *-extension of *human*, with the relation $\sqsubseteq_{animal}$. Or, put in somewhat different terms, as far as you and I are concerned, forming our mereological sum within the extension $X_{human'^*}$ yields the same result as forming our mereological sum within the extension $X_{animal'^*}$.

The upshot of this is that the relations $\sqsubseteq_{P_c}$ for the different predicates P_c in $\mathcal{C}$ are all 'excerpts' from a general part-whole relation between atomic and non-atomic entities making up a very richly populated, largely homogeneous mereological universe: Whenever two entities a and b belong both to the *-extension of predicate P_c and to the extension of P'_c, then it will be the case that $a \sqsubseteq_{P_c} b$ iff $a \sqsubseteq_{P'_c} b$.

This same consideration also applies to the * -extensions of mass predicates: if a and b are atomic or non-atomic portions of matter that belong both to the *-extension of the mass predicates P_m and P'_m, then $a \sqsubseteq_{P_m} b$ iff $a \sqsubseteq_{P'_m} b$.

So the relation $\sqsubseteq$ we want for our amalgamated structure is simply the union of all the relations $\sqsubseteq_{P_c}$ and all the relations $\sqsubseteq_{P_m}$ for the predicates P_c and P_m in the sets $\mathcal{C}$ and $\mathcal{M}$. Note that because the relations $\sqsubseteq_P$ are assumed to return the same results wherever more than one of them is applicable, the atoms of the new structure, as defined in terms of $\sqsubseteq$, are the very same as the atoms from the *-extensions for the individual P's. And since the operations $\cup_P$ are definable from the corresponding relations $\sqsubseteq_P$ the relations $\sqsubseteq_P$, the same is true for those operations.

What we end up with in this way is an amalgamation structure $< D, \sqsubseteq, \sqsubseteq_m >$, where D, $\sqsubseteq$ and $\sqsubseteq_m$ are as explained above. To do justice to what I take to be Link's and Bach's intentions we have to make one further assumption relating to this amalgamation, which doesn't follow from the coincidence assumption as we have stated it. This is the assumption that when the extensions of two predicates P and P' overlap, then any element in the overlap that is an atomic member of the extension of one of these predicates must also be an atomic member of the extension of the other. (On e consequence of tho is that if $P_c \in \mathcal{C}$ and $P_m \in \mathcal{M}$, their extensions cannot have any elements in common.) On this assumption D has a subset A consisting of all the atoms in the sense of $\sqsubseteq$ (see Definition (67.b)). Some of these are atoms from the *-extensions of count predicates in $\mathcal{C}$ but the vast majority of them are portions of the stuff that makes up the extensions of the predicates in $\mathcal{M}$. Note that this way of amalgamating the *-extensions of the predicates in $\mathcal{C}$ and $\mathcal{M}$ is parsimonious in the sense that D need not be closed under the operation of mereological addition. For suppose that a belongs to the extension of P and that b belongs to the extension of P' but that there is no predicate P'' in either $\mathcal{C}$ or $\mathcal{M}$ the extension of which contains both a and b. Then in D there will be no element $a \cup b$. (So mereological addition will in general be a partial operation on D, which isn't defined for every combination of elements from D.) More informally the point can be formulated as follows: when a and b are entities that are 'fundamentally different' from each other in the sense that there is no predicate in $\mathcal{C}$ or $\mathcal{M}$ that classifies a and b as being of the same kind, then there is no plurality containing these two entities. For an example, suppose that $\mathcal{C}$ contains the predicate *human* and the predicate *chair*, but no predicate (such as 'physical object' or something in this spirit, with an extension that includes both that of human and that of chair). Then D will not contain the plurality consisting of me and the chair I am sitting on.

To repeat, I think that such a parsimonious amalgamation would have done fine for both Link and Bach. But this is not the option that Link chose. In his amalgamation the operation $\cup$ is treated as applicable to all combinations of elements of the Domain of the amalgamation structure. This leads to a domain D' that is larger than the Domain D introduced above. It is obtained from D via global closure under mereological addition. In particular, D' will contain the plurality consisting of me and my chair even of there is no predicate in $\mathcal{C}$ whose extension contains both of us.

The structures $<D', \sqsubseteq, \sqsubseteq_m>$ obtained in this way contain vast quantities of what appears to be just ontological flotsam, which doesn't do any work in linguistically relevant applications of metrology. But since our primary concern here is exegesis of the proposals that (as far as I am able to tell) Link and Bach want to make, it is this, seemingly oversaturated, kind of structure – in which $\cup$ is defined for all combinations of elements of D – that we shall adopt. I will refer to this structure as $\mathrm{EXT}(\mathcal{C}, \mathcal{M})$ (for the *Extension Structure* determined by $\mathcal{C} and \mathcal{M}$).

One reason why Link saw his structures $\mathrm{EXT}(\mathcal{C}, \mathcal{M})$ as a natural choice is that his approach does not consider the relativization to sets of predicates. He just assumed that there is a general distinction between 'things', the elements of the extensions of count predicates, and 'stuff', which makes up the extensions of mass predicates, and seems to have regarded that as enough of a justification for assuming the two general part-whole relations $\sqsubseteq$ and $\sqsubseteq_m$. Central to his conception was the relation between objects – the elements of the extensions of count nouns – and the stuff of which they are constituted, a relation about which more will be said in the next section. This is also the main reason for the dual status that portions of matter have in this account: as instances of the predicate 'portion of matter' they are atoms, but as elements of the extension of some mass predicate they are not. (Formally, this is as we saw the difference between being an atom in the sense of $\sqsubseteq$ but not in the sense of $\sqsubseteq_m$.)

One way in which one might try to reconcile Link's liberal conception of the over-all structure of what is described by count and mass nouns with the notion that such structures should be based on pairs of predicate sets $\mathcal{C}$ and $\mathcal{M}$ would be to assume that $\mathcal{C}$ and $\mathcal{M}$ contain 'port-manteau' predicates like 'stuff' or 'matter' for $\mathcal{M}$ and something like 'thing or matter portion' for $\mathcal{C}$. Putting the *-extensions of 'thing' and 'stuff' together into a single mereological structure will still not give use we something in which mereological addition is defined for all possible combinations of elements from the combined domain: we still have no guarantee that $a \cup b$ exists when a is a thing and b a portion of stuff. To make sure that $\cup$ is a truly total operation, which is also defined for such combinations, we would need some even more general predicate than 'thing' or 'stuff', the extension of which would subsume both the extension of ‘thing’ and that of ‘stuff’. But such a predicate would of course be neither a count nor a mass predicate in the technical sense we have been using these terms, but some kind of hybrid of the two.

However, it is precisely this – including port-manteau predicates like those of the last paragraph in $\mathcal{C}$ and $\mathcal{M}$, or, more generally, including too many predicates in these sets – that I think cannot be done without impunity. The reason is this. As mentioned earlier, science has told us that matter consists of smallest parts – atoms and molecules – and this is wisdom that has become part of the world view of all of us. But that view of matter simply isn’t compatible with the conception of the ‘stuff’ described as infinitely divisible. You cannot have it both ways, no matter how emphatic you may be that the infinite divisibility is ‘of course an idealization’. The idealization is perfectly all right in principle, but only in contexts in which the atomicity of matter is kept under wraps. One way to do that is to limit attention to only a limited number of ‘macroscopic’ count predicates and to avoid ‘microscopic’ predicates such as ‘hydrogen atom’ and the rest. By limiting $\mathcal{C}$ in this way and perhaps imposing corresponding constraints on $\mathcal{M}$ we can extract parts of the over-all ontology for which the infinite divisibility assumption does not lead to explicit conflict. I will refer to such ‘excerpts’ from what appears to be an incoherent and thus illusionary all-inclusive ontology in terms of (wisely chosen) predicate sets $\mathcal{C}$ and $\mathcal{M}$ as *Link extractions*. More specifically, I will refer to the amalgamation structures for which we adopted the name ‘EXT($\mathcal{C}$,$\mathcal{M}$)’ also as ‘Link-extractions’.

F. Adding Constitution

The conflict we spoke of in the last couple of paragraphs takes on a formally explicit character when we add to a Link-extraction the one piece that is still missing from it and that is essential to Link’s account. This is a function h that directly connects the extensions of count nouns with parts of the extensions of mass nouns. The notion of ‘homomorphism’ will be explained momentarily.) h maps the elements of the *-extensions of mass nouns to the portions of stuff of which those elements are constituted. h is a *homomorphism with respect to the operations* $\cup$ *and* $\cup_m$. This means that whenever a and b are elements of the Domain D of our amalgamation structure, then when you add the portion of stuff constituting a and the portion of stuff constituting b you get the portion of stuff that constitutes the plurality $a \cup b$; or, more formally:

(70) $h(a \cup b) = h(a) \cup_m h(b)$

Intuitively h *is* the 'constitution function': it maps each atomic or non-atomic entity onto the portion of stuff that constitutes it. (So we can paraphrase '$h(x)$' as 'the portion of stuff constituting x'.) Link wants this function to be defined not only on the members of the *-extensions of count predicates, but on all elements of D, including the members of the extensions of mass predicates and any pluralities containing such elements. This can be done straightforwardly by specifying h as the identity function on the elements of the extensions of mass predicates and then extending its application by stipulation that (70) also hold for pluralities that contain portions of stuff:

(71) a. If a is a portion of matter (i.e. a belongs to the extension of some predicate in $\mathcal{M}$), then $h(a) = a$;

b. for any a, b in D, $h(a \cup b) =_{def} h(a) \cup_m h(b)$.

Note – this is the promised statement of the close connection between the extensions and *-extensions of mass predicates – that (71.b) tells us something in particular about the connection between the operations $\cup$ and $\cup_m$ within the *-extension of any given mass predicate P_m in $\mathcal{M}$: Applying (71.b) to the case where a and b both belong to the extension P_m – i.e. a and b are portions of the stuff described by P_m – we get that $h(a \cup b)$, the stuff of which the plurality $a \cup b$ of portions a and b, which belongs to the *-extension of P_m, but not to its extension, is equal to the material fusion $a \cup_m b$, which is another portion of stuff in the extension of P_m. (This follows from (71) because in the case we are considering we have that $h(a) = a$ and $h(b) = b$ in view of (71.a).)

This purely formal extension of h to entities that have elements of the extensions of mass nouns as parts doesn't alter the fact that it is the applications of h to atomic and non-atomic elements of the *-extensions of count predicates which real;ly matter. It is here that h does its real work: pick out, for each such atomic or non-atomic entity, the portion of stuff that constitutes it.

If we want to make h a part of the amalgamation structure – something that Link sees as a central point of the same enterprise – and we want at the same time to limit amalgamation to pairs of predicate sets $< \mathcal{C}, \mathcal{M} >$, then we must make sure that the two sets are 'commensurable' in the sense that for each element a in the extension of any predicate P in $\mathcal{C}$ there are one or more predicates in $\mathcal{M}$ which (singly or jointly) cover the stuff that constitutes a. For instance, if *human* is one of the predicates in $\mathcal{C}$, then $\mathcal{M}$ must contain predicates that cover the different kinds of stuff that make up human bodies. ($\mathcal{M}$ might have a single predicate, such as 'human tissue' that can do this job all on its own, but I cannot see any good grounds for assuming that there will always such predicates, which cover all the kinds of stuff that can go into the constitution of the different things that satisfy some given count predicate.)

We can now see in more formal terms why the atomistic conception of matter entails that we cannot build coherent amalgamation structures for predicate sets that contain microscopic predicates like *water molecule* (the predicate that is true of all and only those molecules that consist of one oxygen and two hydrogen atoms). Take some particular water molecule a. $h(a)$ is the portion of stuff that makes up a. Presumably that is a portion of water – what else could it be? So the mass predicate *water*, or some other mass predicate whose extension includes water, must be part of the set $\mathcal{M}$ of mass predicates and the portion of it that makes up a will be part of the extension of this predicate. But this portion cannot be divided into smaller portions of water. So divisibility fails for the extension of this predicate, contradicting the general assumption we made about such extensions. (While there is nothing new in this little argument and it is all in all rather boring and predictable, it does show that restrictions have to be put on the sets of predicates for which a construction of the kind Link envisages is possible at all. And note that it is not only predicates like *water molecule* that have to be excluded. Any general count noun, whose extension is understood as including the extensions of such predicates is equally proscribed. For instance, we cannot admit a predicate such as *thing* if 'thing' is to be understood as including also microscopic things such as atoms and molecules.)

To summarize this exposition we have given of Link's algebraic account of the semantics of count nouns and mass nouns:

For every pair $< \mathcal{C}, \mathcal{M} >$ of sets of count predicates and mass predicates we can construct a Link-extraction EXT($\mathcal{C}$,$\mathcal{M}$) of the form $<D, \sqsubseteq, \sqsubseteq_m>$. But in order that $<D, \sqsubseteq, \sqsubseteq_m>$ can be extended with a constitution function h, the set $\mathcal{M}$ must be attuned to the set $\mathcal{C}$.

Throughout the remainder of our discussion of Bach's paper we will assume that Link extractions come with a constitution function, and that $\mathcal{C}$ and $\mathcal{M}$ have been chosen in such a way as to make this possible. The structures representing the semantics of count and mass nouns that we will consider from now on will thus always be of the form $< D, \sqsubseteq, \sqsubseteq_m, h >$.

A simple example may help to clarify the quite general and abstract considerations of this section. Consider the structure EXT($\{screw\}$,$\{metal\}$) and assume that all screws are made entirely out of metal:

EXT($\{screw\}$,$\{metal\}$) $= < D, \sqsubseteq, \sqsubseteq_m, h >$, where

D is the closure under $\cup$ of $X^*_{screw} \cup X^*_{metal}$ and $\sqsubseteq$, $\sqsubseteq_m$ and h are as defined above.

In this case:

(i) X_{screw} is the set of all screws (in the particular world and at the particular time that we assume to have been we fixed in advance);

(ii) X^*_{screw} is the union of that set with the set of all its pluralities (collections consisting of two or more screws);

(iii) X_{metal} is the set of all portions of metal that are wholly contained within some plurality of screws;

(iv) X^*_{metal} is the union of X_{metal} with the set of all its pluralities (each consisting of two or more portions of metal in X_{metal});

(v) $\sqsubseteq$ is the mereological part-whole relation on D.

(vi) $\sqsubseteq_m$ is the mereological part-whole relation on X_{metal} seen as stuff.

(vii) h is the function that maps each screw to the portion of metal that it is made of and, likewise, maps each plurality Y of screws onto the plurality of portions of metal that the screws included in Y are made of.

(Recall: If the set of portions of metal is identified along the lines discussed earlier for butter, then there are vastly more of those than there are screws or even pluralities of screws in D. Recall also the assumption that X_{metal} is atom-less.)

One way in which we might think of modifying $<D, \sqsubseteq, \sqsubseteq_m, h>$ would be to include in D all the physical atoms belonging to the metal occurring in D. But if we do that, then of course X_{metal} will have to be atomic as opposed to atom-less. For if a is a metal atom (formally: '$MetAt(a)$'), then $h(a)$ will be a portion of metal that cannot be divided into smaller portions: there can be no portions of metal in X_{metal} that are smaller than $h(a)$.

This means that the new structure is atomic through and through (and therefore can be expanded to an a Boolean Algebra by adding appropriate operations $\cap$ and $^-$ and a 1 and a 0; see the definition of Boolean Algebras and the following discussions on pp. 49 ff of these Notes).

The resulting structure will now have the metal atoms as additional atoms in the sense of mereology and differ from $<D, \sqsubseteq, \sqsubseteq_m, h>$ in being atomic rather than having an atom-less part. Such a structure explicitly refutes the idealization according to which 'stuff' is homogeneous, by virtue of being (among other things) infinitely divisible. Furthermore, the new structure is different also in that we now have a further part-whole relation to contend with, viz. that which holds between any screw, or collection of screws, and the physical metal atoms from which it is made up. We must distinguish this relation from the mereological relation $\sqsubseteq$ we already have. But the relation is definable in terms of the notions we have. First, the notion $At_{\sqsubseteq_m}$ of being a physical metal atom can be defined in terms of $\sqsubseteq_m$ as that of being a mereological atom in the sense of this relation; and, second, the new relation – let us call it '$\sqsubseteq_{MetAt}$' – can then be defined by:

(72) $a \sqsubseteq_{MetAt} b$ iff $a \in X_{MetAt}$ & $At_{\sqsubseteq_m}(h(a))$ & $b \in X^*_{screw}$ & $h(a) \sqsubseteq_m h(b)$

Our informal description of the structure above, in which the extension of *metal* goes down to the individual metal atoms, has made use – and in fact, it couldn't but make use – of the count noun *metal atom.* As we have described the structure it is generated by the pair of singleton sets $<\{$*screw*$\},\{$*metal*$\}>$. And if it is the 'true' structure generated by this pair, then of course we can turn it without much change into the structure generated by the pair $<\{$*screw,metal atom* $\},\{$*metal*$\}>$, viz. by assigning to *metal atom* the set of all metal atoms in D and to extend the function h in the obvious way: for each metal atom a $h(a)$ will be that indivisible portion of metal that consists of that atom only, and h can then be further extended to apply to pluralities containing metal atoms as parts in the by now familiar way. Of this structure it is fully plain that the extension it assigns to the mass predicate *metal* is *not* homogeneous in the way we have defined that term; the stuff that individual atoms are made of is no more divisible into smaller parts of the same kind than the physical atoms themselves.

The obvious conclusion from all this is that at least for mass predicates whose extension consists of physical matter homogeneity is an idealization which can be realized only in structures of the form EXT($\mathcal{C}$,$\mathcal{M}$) for carefully selected classes $\mathcal{C}$ and $\mathcal{M}$. As soon as $\mathcal{C}$ contains predicates that contain individual atoms or molecules in their extensions, then *every* physically realistic structure EXT($\mathcal{C}$,$\mathcal{M}$) will flout homogeneousness in the way just discussed. The idealization of homogeneous extensions for mass nouns is possible only when we stay at the macroscopic level, while bracketing out all reference to the microscopic.

But in fact, with some mass nouns the idealization involved in the homogeneity assumption goes much further than this. With *sand* or *gravel* we are able to see the smallest parts with our own eyes; and with a mass noun like *furniture* the idealization is even more extreme. When you start dividing the furniture in the room in the sense of mereology the end is reached pretty quickly, and some or all of the smallest pieces you have reached will as often as not still be quite bulky. A mereological structure in which *furniture* has a homogeneous extension is a far cry from the world as we know it.

There are some obvious general principles, connecting the relations $\sqsubseteq$ and $\sqsubseteq_{MetAt}$ and the function h, that this new structure must satisfy.

First, we can define the function *Atmc* (for 'atomic material constituent'), which maps each screw a onto the plurality of metal atoms contained in it, as follows:

$$Atmc(a) = \cup\{c \in D : MetAt(c) \in \& \ c \sqsubseteq_{MetAt} a\}$$

(Here by $\cup\{c \in D : MetAt(c) \in \& \ c \sqsubseteq_{MetAt} a\}$ we understand the plurality consisting of all the atoms c contained in a. Note that if the screw a contains only finitely many atoms (as no doubt it will), then the fact that our structure is an upper semi-lattice with respect to $\sqsubseteq_{MetAt}$ will automatically guarantee the existence of $Atmc(a)$.)

Here are some simple general principles that can be stated in terms of *Atcm*. Suppose that a, b are screws from D:

(i) $Atcm(a \ \cup \ b) = Atcm(a) \ \cup \ Atcm(b)$;

(ii) $Atcm(a) \neq a$;

(iii) $h(Atcm(a \cup \ b)) = h(Atcm(a) \cup_{MetAt} \ Atcm(b)) = h(Acmt(a)) \cup_m$
$h(Atcm(b)) = \mathrm{h(a)} \cup_m h(b)$.

We have just seen that by adding the count predicate *metal atom* to the set $\mathcal{C}$ of our example gives rise to an additional part-whole relation, $\sqsubseteq_{MetAt}$, between elements of the extensions of different count predicates. But note well that this example is extremely simple. More realistic Link extractions will have substantial sets of count predicates, whose extensions may be connected by a variety of different part-whole relations. Think for instance of what kinds of parts, and parts of parts , and parts of parts, .. it is natural to think of in connection with a Boeing 747, or a modern cruise ship. We will return to the issue of part-whole relations between count noun extensions towards the end of the next section.

A further natural extension of the structures we have been considering up to this point involves space and spatial location. I have decided to keep space out of the structure $\text{EXT}(\mathcal{C}, \mathcal{M})$ and represent it as a separate structure consisting of a set of portions of space (or 'spatial regions') with its own 'mereological relation $\sqsubseteq_{Sp}$ of spatial inclusion. In addition we assume a function *loc* which assigns each element a of D its spatial location $loc(a)$. The following general principles should be obvious:

(i) If a, b belong to the extension of *screw*, then $loc(a) \mid\mid loc(b)$

(ii) (Here $loc(a) \mid\mid loc(b)$ means that $loc(a)$ and $loc(b)$ are non-overlapping regions, i.e. there is no region r such that both $r \sqsubseteq_{Sp} loc(a)$ and $r \sqsubseteq_{Sp} loc(b)$.)

(iii) If $a \sqsubseteq b$, then $loc(a) \sqsubseteq_{Sp} loc(b)$;

(iv) If $a \sqsubseteq_m b$, then $loc(a) \sqsubseteq_{Sp} loc(b)$;

(v) $loc(a) = loc(h(a))$

We noted that many nouns can be used both as count nouns and as mass nouns. In some cases this possibility has been fully conventionalized. An English example mentioned earlier is *cheese*, which is used equally happily as mass noun and as count noun.

But apart from such words that straddle the divide between count nouns and mass nouns there are also general strategies for 'coercing' count nouns into mass nouns and vice versa. The perhaps most prominent operation that transforms count nouns into mass nouns is a very productive mechanism, which is illustrated by examples like these.

(73) a. There is apple in the salad.

b. Today the main dish is rabbit.

This way of transforming a count noun into a mass noun – the way *apple* has been transformed into a mass noun in (73.a) *rabbit* has been transformed into a mass noun in (73.b) – is one that yields mass nouns which describe the stuff from which the individuals described by the count noun are made up. It is referred to in linguistics as the *Universal Grinder*.

It is also possible to coerce mass nouns into count nouns. Familiar examples are *beer* and *wine*, as used in the following sentences:

(74) a. He took another beer from the fridge.

b. This restaurant is best known for its excellent wines.

Note that the ways in which the meaning of the count noun use of *beer* in (74.a) relates to that of the mass noun *beer* and that in which the meaning of the count noun use of *wine* in (74.b) relates to that of the mass noun *wine* are different: *beer* in (74.a) is used to describe a portion of beer (presumably a can or a bottle in this instance). *wine* in (74.b), on the other hand stands for kinds or brands of wine. This is a general feature of count noun coercions of mass nouns: their semantics varies, depending in part on the context, but also on the particular mass noun in question. For instance it is much more difficult (though not impossible) to use *a wine* to refer to a portion of wine (e.g. a glass of wine) than it is for the count noun use of *beer*.

A further point worth noting is that the coercions from count to mass and from mass to count nouns do not involve any overt morphological marking. There is no a priori reason why this should be so. And indeed, there are some cases where the word for an animal is different from the one we use to refer to the meat of that animal. English uses *beef* for the latter purpose, and not *cow*; *veal* and not *calf*; *venison* rather than *deer*. And we also sometimes make use of a more regular, morphologically overt formation option, involving *-meat*, as in *horse-meat*, *wale-meat* and so on. Sometimes more than one option is available (*elk* as well as *elk-meat*, *bison* as well as *bison-meat*), but that seems to be true only for some cases. (You just cannot say (I think) 'We had horse for dinner last night.')

We already mentioned in passing the possibility of turning mass nouns into count nouns through the use of English 'classifiers', as in *a can/bottle of beer*, *a glass of wine*, *a slice of bread*, *a lump of butter*, *a brand of orange juice* and so on. More on this in the next section, which looks more closely at the structures proposed by Link.

Since the Grinder transition from count nouns to mass nouns is so productive and systematic it suggests one natural 'closure' condition on the candidate pairs $<\mathcal{C}, \mathcal{M}>$ for structures $\mathrm{EXT}(\mathcal{C}, \mathcal{M})$: for any count predicate P_c in $\mathcal{C}$,

$\mathcal{M}$ should contain the corresponding mass predicate $Gr(P_c)$ that results from subjecting P_c to the Universal Grinder.

Some further remarks on Link: 'The Logical Analysis of Plurals and Mass Terms: A Lattice-theoretical Approach'

The algebraic structures that Link proposes in (Link (1983)) are of the following form: $\mathcal{B}$ =<<E,$\cup_i, \subseteq_i$>, A, <D,$\cup, \subseteq$>, h >, where A $\subseteq$ E and D $\subseteq$ A. $< E \subseteq_i >$ and $<D, \subseteq>$ are both complete upper semi-lattices and the operations $\cup_i$ and $\cup$ can be defined in terms of $\subseteq_i$ and $\subseteq$ in the familiar way. (Note the change of notation: $\subseteq_i$ corresponds to what we have so far denoted as '$\sqsubseteq$' and $\subseteq$ to what we have been denoting as '$\sqsubseteq_m$'.) A is the totality of all atoms (= the potential referents of singular DPs) and D is the subset consisting of all portions of matter (the potential denotations of singular mass DPs, or 'mass terms' in Link's terminology). E consists of the atoms together with all their sums (the 'i-sums' in Link's terms). Thus the members of A are atoms in the sense of $\subseteq_i$. This is true in particular for the members of D (the 'portions of matter'); these are atoms in the sense of $\subseteq_i$ just as the other elements of A. (But note well: the members of D are *not* atoms in the sense of $\subseteq$). $\subseteq$ orders the members of D in the sense of material constitution; it is $\subseteq$, in other words, which represents that aspect of the structure of D that in the intended instances of Link's account is homogeneous in the sense of closure under mereological summation and unlimited divisibility. h is, as in the structures $EXT(\mathcal{C}, \mathcal{M})$ defined in the last section, the material constitution function. It maps the elements of A onto elements of D and this map is then homomorphically extended to E\A (i.e. to the non-atomic part of E). Thus, as before, if a belongs to A, then $h(a)$ belongs to D; and if a belongs to D, then $h(a) = a$. Furthermore, h is a homomorphism with respect to the relations $\subseteq_i$ and $\subseteq$: for any $a, b \in$ E, if $a \sqsubseteq_i b$, then $h(a) \sqsubseteq h(b)$. (Since sums and i-sums are definable from $\subseteq$ and $\sqsubseteq_i$, it follows that the h-values of any i-sums are the sums of the h-values of the summands of those i-sums.) As can be inferred from all this, the domain of h is all of E, while its range is some subset of D.

I take it to be part of Link's intentions that a primary role of the members of D is to serve as referents of phrases like *the butter on your plate*, *the water in this glass*, *the garbage in all those containers* and so on; and that, by the same token, they also are to provide semantic values for indefinite mass terms like *some butter*, as it occurs in a sentence like *I bought some butter*, and perhaps also for bare mass terms like the one word DP *butter*, as we find it in a sentence like *I also bought butter*. In contrast, referring DPs built from count nouns, like *the neighbor who lives upstairs*, *the magnolia in our garden* and so on are to be found in A.

This much is just Link. But besides the three types of DPs that can be built from mass nouns and that were mentioned in the last sentence, we will now also take a fourth category of DPs into account, which can be built from English mass nouns, and to which I will refer as 'classifier construction DPs'. For these phrases it is not immediately clear where their referents should be located in ontological structure's like Link's. Exploring this question will be useful in that it will allow us to get a better grip on the correlations between the semantics of mass nouns and the semantics of count nouns.

Classifier Constructions, Classifier Constructions and Classifier Nouns

But first, what kinds of constructions are classifier constructions and the DPs that can be formed out of them? The phrase 'classifier construction' is inspired by the phenomenon of classifier languages, which we touched upon briefly earlier in these comments. Recall from that earlier discussion: In classifier languages nouns generally behave like mass nouns, except for a handful of so-called 'classifier nouns', which can be used to turn mass noun phrases into count noun phrases. Such phrases typically have a grammatical form like this:

[Cl 'of' M]],

where Cl is a classifier, M a (mass) noun phrase and 'of' is whatever the language uses to combine classifier and mass noun in the grammatically correct way. European languages such as English are different from classifier languages in that their nominal vocabulary contains as many count nouns as mass nouns (and typically a lot more) and that count nouns and mass nouns can be used to form full noun phrases in roughly similar ways. (Though as we have seen, there are also important differences: English mass noun DPs can be bare (the one-word phrases *butter*, *water*, *trash* are among them), mass noun DPs do not allow for the indefinite article *a* and English mass nouns don't have plurals.) We also noted in passing that English mass noun phrases can be turned into full DPs by way of constructions that are strongly reminiscent of the standard construction of DPs in classifier languages. It is these to constructions that I refer as 'classifier constructions'.

Some examples of English classifier constructions are given in (75.a-d). Classifier construction DPs – DPs that can be formed from classifier constructions like those in (75.a-d) by combining them with a determiner – are given in (75.e-h).

(75) a. bit of butter/water/bread/cheese/ice cream/garbage/furniture
b. quantity of butter/water/bread/cheese/ice cream/garbage/furniture
c. pound of butter/bread/cheese//liter of water//quart of ice cream/ton of garbage
d. lump of butter/glass of water/slice/loaf of bread/hunk of cheese/scoop of ice cream/load of garbage/cartload of furniture
e. a/the/that bit of butter/water/bread/cheese/ice cream/garbage/furniture
f. a/the/that portion of butter/water/bread/cheese/ice cream/garbage/furniture
g. a/the/that pound of butter/bread/cheese//litre of water//quart of ice cream/ton of garbage
h. a/the/that lump of butter/glass of water/slice/loaf of bread/hunk of cheese/scoop of ice cream/load of garbage/cartload of furniture

As these examples suggest, the difference between classifier constructions and classifier construction DPs is like that between a nominal predicate and a nominal referring phrase. Classifier constructions have extensions, like count nouns and mass nouns, whereas classifier construction DPs of the kinds exemplified (75.e-h) are used to to refer to particular entities. This general distinction between referring phrases and other full DPs on the one hand, and the nominal predicates from which these are built on the other should now be familiar and I will not dwell on it any further.

The question that will preoccupy us during the next pages can be asked both in relation to classifier constructions and to referring classifier construction DPs. In relation to referring classifier construction DPs it is: Where in Link's structures can we expect to find the referents of such DPs? In relation to the classifier constructions themselves the question takes the form: Where in Link's ontology do we find their extensions? I will concentrate on the second way of asking the question, and in what follows we will only consider the classifier constructions, ignoring classifier construction DPs week-neigh completely.

So as to be in the right position to tackle the question where in Link's ontological structures the extensions of English classifier constructions should be located, we need to have a closer look at the syntactic structure of these constructions. One difficulty here is that English classifier nouns vary (like, by the way, those of true classifier languages) in what semantic contributions they make in addition to their functional role of turning mass noun phrases into expressions that syntactically behave like count noun phrases. Some classifier nouns, which I will refer to as 'pure classifiers', contribute no such information. Examples are *amount*, *portion*, *quantity*. A portion of butter is just that – some quantity of butter which could have any size, any shape, any properties that a piece or quantity of butter might have. The classifier noun *bit* comes close to being a pure classifier, except that it carries a certain implication that the portions it is used to describe are small for the kind of stuff in question (that which is described by the mass term with which the classifier noun combines). But on the other hand there are also nouns that can be made to play the classifier part in classifier constructions, but that do make quite specific contributions to the semantic of their classifier constructions. An important subcategory of is formed by the open class of *measurement unit nouns*. Measurement unit nouns are nouns like *pound* as in *pound of butter*, *liter* as in *liter of wine*, *yard* as in *yard of taffetas*, *acre* as *acre of arable land* and so forth. And from these classifier constructions we can then form classifier construction DPs such as *one pound of butter*, *a liter/two liters/ three liters/.. of water*, *500 acres of arable land* and so on.

Other non-pure classifier nouns contribute information about shape – examples are *slice*, *cube*, *slither*– or about the container which holds the denoted portion of stuff, like *glass* as in *a glass of wine*, *bowl* as in *a bowl of rice*, *a plateful* as in *a plateful of sauerkraut* and so on.

The syntax and semantics of English classifier constructions can be thought of in two ways. According to the first, which comes close to the informal way we have so far been talking informally about such constructions, classifier nouns function as operators which turn mass noun phrases into count nouns. Semantically, this makes the classifier noun into a functor which maps the extensions of the mass nouns to which it can be applied into the extensions of the classifier constructions that result from applying the classifier noun to them. The second possibility is to construe the classifier noun as the head of the classifier construction, and the combination of *of* and the mass noun that follows it as a modifier of the classifier noun. More precisely, 'of' and the following mass term are to be construed as forming a prepositional phrase, which is adjoined to the classifier noun – much as the prepositional phrase *in the garden* can be adjoined to the noun *tree* to form the complex NP *tree in the garden*. On this second view the classifier noun is a count noun in its own right, with its own extension (and *-extension), and the extension of the classifier construction is obtained from the extension of the classifier noun by the principle that governs combinations of head nouns and prepositional modifiers in general: the extension of the head noun is restricted to those of its elements that also satisfy the modifying adjunct. (Thus just as we obtain the extension of *tree in the garden* from the extension of *tree* by retaining of it those and only those elements which satisfy the condition of being in the garden, so the extension of *slice of butter* is obtained from the extension of *slice* by retaining from it those and only those elements which consist of butter.)

Which of these two analyses of classifier constructions is the 'right' one (or the more appropriate one) is hard to decide, and perhaps the answer should vary depending on what classifier noun is being considered. Arguably the second analysis is more plausible for those classifier nouns that make a substantial semantic contribution of their own, such as *slice*, *lump*, *glass* or *plateful*, while the former might correspond more closely to the grammar that English speakers have internalized for 'pure' classifiers like *quantity*, *bit* or *amount*. A special case are the measure phrases. These have a claim to being a category on their own, with close ties to what are often thought of as quantifying expressions in the realm of count nouns – compare for instance *200 grams of butter* with *half a dozen eggs*. These parallels suggest that for measure phrases the first of our two analyses may well be the more plausible one; but in what follows I will set this intuition aside.

To do full justice to the variety of English classifier nouns it may well be necessary to allow for both of the two construals of classifier constructions that we have just outlined. But if I were to do that I would have to tell the story I want to tell about the semantics of classifier constructions twice over, once for classifier constructions analyzed according to the first construal and once for classifier constructions analyzed according to the second. Since the point of the story would be the same, that would be trying everybody's patience. I will therefore make things easier on author and reader alike by confining myself to just one of the two analyses, that according to which the classifier noun is the head and the following '*of* + mass term' its prepositional modifier.

The Semantics of Classifier Constructions and Classifier Nouns

Let us include a suitable set of classifier nouns into the set $\mathcal{C}$ of the pair $<\mathcal{C},\mathcal{M}>$ – suitable in that for each mass predicate P_m in $\mathcal{M}$ there is at least pure classifier predicate P_{P_m} that can be applied to it to form a classifier construction, and let us assume that $\mathcal{C}$ is closed under the formation of classifier constructions involving classifier nouns form $\mathcal{C}$ and mass predicates from $\mathcal{M}$. This putts us into a position to revise and simplify the assumptions we have so far made about the semantics of mass nouns. Up to now we assumed that mass nouns come, like count nouns, with an extension and a *-extension. This assumption assigns to the elements of the extension of a mass noun a curiously ambivalent status: on the one hand they are non-atomic elements of the mass noun's extension itself – an atom-less structure ordered by $\subseteq$, in which there are presumed to be no atoms whatsoever – and on the other they play the part of atoms in the mass noun's *-extension, an atomic structure that is ordered by $\subseteq_i$. Another point, which we could have high-lighted earlier but so far hasn't been made explicit, is that the only clear cases of terms referring to pluralities of matter portions (that is, to non-atomic members of the *-extensions of mass nouns) are phrases involving classifier constructions, like *the lump of butter on my plate and the lump if butter in the fridge* or *several slices of cake* or *two portions of ice cream.* Since we are now including classifier constructions explicitly among the types of noun phrases that make up the English noun phrase repertoire, this last observation takes on a new importance: If it is only as constituents of classifier constructions that mass nouns can enter into phrases which refer to non-atomic parts of their *-extensions, then that should arguably be taken as indicating that the mass nouns themselves do not have *-extensions at all; the *-extensions involved in these cases are really the *-extensions of the classifier nouns that appear to be indispensable ingredients to phrases that refer to pluralities of potions of stuff. And the general conclusion – this seems a *good* case of Occam's razor – should then be that *-extensions of mass nouns can be dispensed with altogether.

But if we are to get rid of the *-extensions of mass predicates, then something will have to take their place. For as we have just argued, there *are* ways to refer to pluralities of portions of matter and the ontology must be able to account for that. However, if it is only through the use of classifier constructions that such references can be made, then it is also clear what should now take over this task. The crux of the matter is how the extensions of mass nouns are related to the extensions of the classifier constructions in which they can occur. And it is quite clear what the relations is. The matter is best explained with an example. Consider the extensions of the mass noun *butter* and the classifier noun *slice*. The extension of *butter* will contain, among the countless variety of butter portions that make it up, a certain contingent that have the form of slices. These elements are also found in the extension of *slice*, and are distinguished from the other elements of that extension in that they are slices of butter, rather than slices of some other stuff. It is this 'overlap' between the extension of *butter* and that of *slice* which establishes a 1-1 correspondence between a part of the extension of *butter* and the corresponding part of the extension of *slice*

In fact, if the two extensions literally overlapped, then the 1-1 correspondence would be nothing but identity. Whenever two sets A and B, of whatever kind, have a non-empty intersection, then we can describe this state of affairs by saying that A has a part A' and B a part B' which stand in the 1-1 correspondence established by the identity relation: for each element a of A' there is a corresponding element b of B' that stands to a in the 1-1 correspondence and conversely. But this is just a somewhat perverse way of saying that A and B have a non-empty overlap, for the b from B' that 'corresponds' to the element a from A' is nothing other than a itself.

The idea that we can account for the connection between mass noun and count noun extensions in term of literal overlap is attractive for its simplicity But unfortunately it won't do For reasons that will soon be apparent we must keep the extensions of count nouns and mass nouns apart, just as in the earlier ontological structures. That means that the correspondences which we need between extensions of mass nouns and extensions of classifier nouns can't be identity. But Link's structures already provide them: When P_c is a classifier for the mass predicate P_m, then h will map some objects a in the extension of P_c to elements in the extension of P_m – those a which are made up entirely of stuff that is correctly described by P_m. For instance, h will map part of the extension of *slice* – that part which consists of slices of butter – onto elements of the extension of *butter*; each such slice a will be mapped onto the portion of butter of which it is made up. I take it that on the extensions of classifier domains h is always 1-1, but that assumption isn't essential to what follows.

Note that this conception of the relation between the extensions of mass nouns and classifier nouns allows us to stick to the general principle that for all elements a belonging to *-extensions of count nouns, $h(a) \neq a$, while on the extensions of mass nouns it reduces, by fiat, to identity.

Pure classifier nouns are distinguished from non-pure classifiers in that the correspondence established by h involves the *complete* extensions of the mass nouns with which they can be combined into classifier constructions. Take the classifier noun *portion*, the one of which we have made ample use in our informal descriptions of the semantics of mass nouns before we took English classifier constructions explicitly into account. As I put it informally, the extension of a mass noun like *butter* consists of all '(actual and potential) portions of butter'. If that is to be taken literally, then *all* elements of the extension of *butter* belong to the extension of *portion*.

In order that we can do without the *-extensions of mass nouns in a structure generated by a class $\mathcal{C}$ of count nouns and a class $\mathcal{M}$ of mass nouns along the lines just sketched there ought to be for each $P_m \in \mathcal{M}$ at least one classifier noun $P_{cl(P_m)} \in \mathcal{C}$ that can form a classifier construction out of P_m; the work that was done previously by the *-extension of P_m so far is now done by the *-extension of $P_{cl(P_m)}$.

So let us assume that $\mathcal{M}$ and $\mathcal{C}$ are related in this way (i.e. that for every P_m

$\in \mathcal{M}$ there is at least one classifier noun $P_{cl(P_m)} \in \mathcal{C}$). Then we can associate with the pair $< \mathcal{C}, \mathcal{M} >$ a structure EXT'($\mathcal{C}$, $\mathcal{M} >$) in which the members of $\mathcal{M}$ have extensions but no *-extensions, while the members of $\mathcal{C}$ come with *-extensions as well as extensions.

Furthermore, when the *-extensions of the members of $\mathcal{M}$ are no longer part of the structure, there no longer is any need to distinguish between the relations $\subseteq$ and $\subseteq_i$ (nor, of course, between $\cup$ and $\cup_i$, which are definable in terms of $\subseteq$ and $\subseteq_i$). Thus we can now assume that there is just one fundamental part-whole relation $\subseteq$, which on the one hand does the wreak of partially ordering the extensions of mass nouns and on the other that of partial ordering the *-extensions of count nouns.

Note that, formally speaking, $\subseteq$ imposes not only a partial ordering on the extensions of mass nouns and the *-extensions of count nouns, but also on the extensions of count nouns. It is just that the order it imposes on the extension of a typical count noun such as *person* or *cat* is a trivial one, which only holds between any two members of the extension (if and) only if they are the one and the same: on such extensions $\subseteq$ is just the identity relation.

Would that things could be this simple and pristine. But the complexities we were facing before classifying constructions were brought into the game cannot be got rid of quite so easily. Not all count nouns have extensions on which $\subseteq$ is trivial in the sense in which it is trivial on the extension of a noun like *cat*. And among the prime examples of count nouns for which this is not so are the very nouns that are at the hub of the simplification just proposed, viz. the classifier nouns. The extension of a classifier noun is, intuitively speaking, ordered by the part-whole relation in the same way as the extensions of those mass nouns to which it is applicable and with which its extension overlaps. And there are other count nouns as well – count nouns that do not function s classifier nouns – whose extensions seem to display a non-trivial part-whole relation. An example is the noun *part*. Consider the parts of a car. The engine of a car surely qualifies as one of its parts. But so can the pistons, although the pistons also qualify as parts of the engine. And the pistons in their turn are composed of smaller parts yet, and those too qualify as parts of the car, and ...

So it looks as if we are saddled with more than one part-whole relation after all, but now in relation to certain count nouns rather than to mass nouns. We have traded our earlier relation $\sqsubseteq_m$ on the extensions of mass nouns for a relation, as yet unnamed, on the extensions of classifier nouns (and certain other count nouns as well) which, just like $\sqsubseteq_m$ had to be distinguished from $\sqsubseteq$, must be distinguished from Link's relation $\subseteq_i$ because we want the members of count noun extensions to be atoms in relation to $\subseteq$ while they are clearly not atoms in relation to the new (as yet unnamed) relation.

We seem to have lost with the right hand what we thought we had just gained with the left. But there has been some gain. For there is a way of recuperating the non-trivial part-whole relations on the extensions of classifier nouns (and other count nouns such as *part*) from the part-whole orderings of the corresponding portions of matter. In fact, we have already seen one example of how such relations can be recovered in the guise of Definition (72) of the relation $a \sqsubseteq_{metAt} b$. I repeat the definition.

(72) $a \sqsubseteq_{metAt} b$ iff $a \in X_{MetAt}$ & $At_{\sqsubseteq_m}(h(a))$ & $b \in X^*_{screw}$ & $h(a) \sqsubseteq_m h(b)$

Two objects a and b stand in the part-whole relation that we are after just when, and because, the portions of matter that constitute them stand in the material part-whole relation $\subseteq$. Let us call the new part-whole relation '$\subseteq_{const}$' (for 'constitution-based part-whole'). Following the lead of (72) $\subseteq_{const}$ can be defined in terms of (76) as follows:

(76) For any objects a, b in A, let $a \subseteq_{const} b$ iff $h(a) \subseteq h(b)$

With Definition (76) the simplicity that we thought we had gained and then lost again, has been regained once more. And this time it is for keeps. We have only one part-whole relation $\subseteq_{const}$ which imposes on the *-extensions of count nouns the structure of upper semi-lattices that are atomic and on the extensions of mass nouns that of upper semi-lattices that are atom-less. In terms of this relation we can define an additional relation $\subseteq_{const}$ between elements of count noun extensions. This second relation imposes different types of ordering structure on the extensions of various count nouns, from the trivial (but most common) cases such as *cat* where it is simply identity to the cases of pure classifier nouns where it copies the homogeneous order imposed by $\subseteq$ on the extensions of the mass nouns to which the classifier applies.

Identity Criteria and Counting Criteria

The observations about the concepts 'part' and 'car part' point to an important distinction which too often is not given the prominence it needs and deserves. This is the distinction between *identity principles* and *counting principles*. An example showing that counting principles aren't the same as identity principles is given by the concept*part* and its various sub-concepts such as *car part*, *helicopter part* and so on. We focus on *car part*. The extension of *car part* is – like the extension of any other count noun –a reflection of the identity criteria of the concept: car parts a and b are distinct if and only if they are distinct elements of the extension of *car part*. But that information isn't enough if you want to count the parts of a given car. Our conception of counting things comes with a strong intuition that when a and b fall under a given count noun concept and $a \subseteq_{const} b$, then one should not count them both. (Put in more formal terms: We have a firm intuition that any proper counting of what falls under a given concept should involve all and only the elements that together form an *anti-chain within* the extension of the concept *with respect to* the partial ordering $\subseteq$. (In general, an *anti-chain* of a partially ordered set $<X,\subseteq>$ is a subset Z of X which intersects each maximal $\subseteq$-chain in X in exactly one point, in other words, anti-chain and maximal chain have a single element in common. A *chain* of X is a subset of X that is linearly ordered by $\sqsubseteq$, and a *maximal chain* of X is a chain of X that cannot be extended to a larger chain of X.)

How do we count parts – e.g. the parts of the given car, or the parts of its engine – in practice? Well, we have to rely on some additional guidance, which goes beyond our understanding of how individual parts of cars or engines can be identified. Context can provide such guidance in a variety of different ways. Here is one example of how it may do that. Suppose you have been asked to count the parts of a car engine that has been taken apart and that its parts are spread out on the work bench. In this situation it is natural to understand the task of counting the parts of the engine as that of counting the separate pieces on the bench. If one of those pieces is the car's carburetor – it could have been taken apart further, into its components, but it hasn't – then the carburetor should be counted as one part, and not as the many smaller parts into which it could have been decomposed further if one had decided to do so. Other contexts will provide the relevant counting criteria via different clues. But some clue about what to do about situations in which a and b a re distinct candidates for being counted while $a \subseteq b$ must be given. Otherwise we wouldn't know what to do.

This counting problem can of course arise only when the extension of the given concept is non-trivially ordered by $\subseteq$. If the order is trivial (as it is in the case of *cat* or *person*), then counting will always be unequivocal, since for any two elements a and b that are candidates for being counted, if $a \subseteq b$, then $a = b$, so we can count 'them' only once. Arguably this is the 'normal' case, and it is presumably because we have come to see it as the normal case that identity principles and counting principles are so often perceived and treated as if they was no difference.

Summary of our Survey of Nominal Ontology

Let us state once more in full what our reflections on Link's proposals for the structure of the ontology for count and mass nouns have led us to. Our adjustment of those proposals has crystallized into ontologies with the following formal structure.

Let $\mathcal{C}$ be a set of count predicates and $\mathcal{M}$ a set of mass predicates and assume that for each predicate P_m in $\mathcal{M}$ there is at least one 'classifier predicate' P_c in $\mathcal{C}$ which 'covers' P_m. (In intuitive terms this means that P_c can be used to combine with P_m in a classifier construction. The formal implication of 'cover' will emerge below.) We refer to the subset of $\mathcal{C}$ that consists of its classifier predicates as '$\mathcal{CL}$'.

A *Link-Extraction* EXT($\mathcal{C}$,$\mathcal{M}$) *based on* $\mathcal{C}$ and $\mathcal{M}$ is a 4-tuple $< D, A, \subseteq, h >$ such that

(i) D is a non-empty set (the 'universe' of EXT($\mathcal{C}$,$\mathcal{M}$));

(ii) A is a subset of D;

(iii) For each predicate P_c in $\mathcal{C}$ A includes the *-extension *-Ext(P_c) of P_c);

(iv) For each predicate P_m in $\mathcal{M}$ $D \backslash A$ includes the extension Ext(P_m) of P_m;

(v) $\subseteq$ is a partial ordering of D;

(vi) $\subseteq$ imposes the structure of an atomic upper semi-lattice on A; for each P_c in $\mathcal{C}$ *-Ext(P_c) is an atomic sub-semi-lattice of this semi-lattice;

(vii) $\subseteq$ imposes the structure of an atom-less upper semi-lattice on $D \setminus A$; for each P_m in $\mathcal{M}$ Ext(P_m) is an atom-less sub-semi-lattice of this semi-lattice;

(viii) h is a $\subseteq$-preserving map from D into $D \setminus A$;

(ix) h is the identity function on $D \backslash A$ (that is, if $b \in D \backslash A$, then $h(b) = b$);

(x) h is disjoint form the identity function on A (if $a \in A$, then $h(a) \neq a$);

(xi) For each predicate P_m in $\mathcal{M}$ there is a predicate P_{cl} in $\mathcal{CL}$ such that for each b in Ext(P_m) there is an a in Ext(P_{cl}) such that $h(a) = b$.

With the help of $\subseteq$ on $D \backslash A$ we can define a constitution relation $\subseteq_{const}$ on A. $\subseteq_{const}$ gives the non-trivial part-whole relation of the extensions of classifier predicates in $\mathcal{CL}$ and certain other count predicates such as *part*.

One of the things that ontologies of the kind defined above ought to be good for is to account for the extension shifts that come with the transformations from mass to count and from count to mass predicates. In Link Extractions it is h that plays one of its central roles in they connection. For instance, with the Universal Grinder transformation of a count noun P_c into the corresponding mass noun $M(P_c)$ comes the change from the *-extension of P_c into the extension $\{m : (\exists a \in$ *-Ext$(P_c))$ m = h(a)$\}$. And with the conversion of a mass noun P_m into a count predicate with the help of a classifier predicate P_{cl} comes a change of the extension Ext(P_m) of P_m into the extension of the classifier construction and this extension is that part of the extension of P_{cl} which consists of all those a whose h-images are in the extension of P_m, i.e. the set $\{a : a \in \text{Ext}(P_{cl})\ \&\ h(a) \in \text{Ext}(P_m)\}$. For the semantics of the less systematic operations that turn mass nouns like *beer*, *wine* or *water* into count nouns, denoting contextually fixed portions of beer, kinds of wine or bodies of water, the ontology offers no ready-made formulas for the corresponding extension transformations. That is in one part because the 'language of Link-Extractions' (with its relation '$\subseteq$' and its functor 'h') lacks the expressive power to do this. But in another, more decisive part it is because as we have defined Link-Extractions the relevant ontological categories, such as kinds, aren't included in them. To adequately describe the semantics of such switching operations between count nouns and mass nouns we need ontologies that are more comprehensive as well as more richly structured.

Before we turn to the eventuality ontologies that are the main target of Bach's paper here is one final reflection on the nominal ontologies which served Bach as the blueprint for his eventuality structures. The intuitive considerations that led Link to the ontologies that, in our final formalization, take the form of Link-Extractions EXT($\mathcal{C}$,$\mathcal{M}$) are reasonably straightforward in connection with count nouns describing concrete entities, which are made out of matter. But it is probably true for all languages that have count nouns, and it is certainly true for a language like English, that a large part of them is not about concrete things. Abstract nouns are all over the place. Here is a handful, but nobody should have trouble expanding this little list, and in various directions: *shape*, *tone*, *symbol*, *virtue*, *contract*, *statement*, *number*, *wave*. What could the 'stuff' be that the things in the extensions of these words are made up from? The rhetorical flavor of this question is reflected by the fact that the Universal Grinder has no purchase on these nouns. (Sentences like 'There was contract all over the lawyer's desk' and 'The blackboard was smeared with number' sound very strained, if not outright ungrammatical.) This is another reason for limiting Link-ontologies to (mutually attuned) pairs of count predicate and mass predicate sets and to leave it to the individual metaphysicians how far they are prepared to push Link's approach (i.e. how comprehensive the sets $\mathcal{C}$, $\mathcal{M}$ are for which they consider Link-Extractions to be well-defined).

A category that needs special consideration when it comes to the comparison between nominal and verbal ontology, are nouns for events, processes and states. English has no end of these – *party*, *race*, *adventure*, *trip*, *run* (the noun), *war*, *peace*, *admiration*, *flow* (the noun), and so on. Some such nouns are ambiguous between clear event interpretations and interpretations as descriptions of entities that may be associated with events but that we would not want to classify as events themselves. Some examples are: *exam*, which can be used to refer to events that consist in someone (or -ones or -thing or -things) being examined but also to refer to the set of questions that were or will be administered on the occasion when an event took or will take place in which the knowledge or intellectual ability of some person or people will be tested; or *stuffing*, which can be used to refer to the process of stuffing the goose but also to the stuff that goes into it on such an occasion; or *question*, which we use to the act of making an inquiry but also to the content of the inquiry. And then there are the nouns about which we may be uncertain wether what they describe is an eventuality or not. What is a wave? What is a wind? What is a passion?

As these last examples indicate, it is not easy to demarcate the set of nouns that describe eventualities. But what matters most for the present discussion is that there are nouns (and many of them) of which it is beyond dispute that they do describe eventualities. Moreover, some languages have general mechanisms for turning verbs into nouns that describe the very same eventualities as the verbs from which they are derived. English is especially productive in this regard. First, the formation of gerunds, by suffixing *-ing* to the stem of the verb, seems to work more or less without exception. And in addition, English is also very liberal in its use of the bare stems of verbs as corresponding nouns. Thus *walk* can be used as a noun, to refer to what you do when engaged in the activity described by the verb, *sleep* as a noun for what you do when you sleep and so on. This means that in English there are no eventualities that can be described by some verb but not by any noun. (A further noteworthy feature of gerund formation in English is that many gerunds can be used both as mass and as count nouns. Thus we can say: 'Again there were several shootings this week.' as well as 'Shooting is dangerous.' On the whole the mass noun use of gerunds seems to be more generally available and should probably be regarded as their primary use. When a verb has both a bare stem noun and a gerund, such as the verb *walk* for which we have the nouns *walk* and *walking*, then as a rule the bare stem noun is used to describe the events that can eb described with the verb, whereas the gerund is used to describe the corresponding processes, see the next section.)

To the extent that Bach is right and events are made up of processes in much the same way that concrete objects are made up from stuff, there would be no problem about including event and process nouns among the sets $\mathcal{C}$ and $\mathcal{M}$ that generate viable Link-Extractions. it is just that for such pairs of count and mass-predicate sets, the extensions of some of the clout nouns and the corresponding mass nouns are rather different from the extensions of count nouns for physical objects and mass nouns for their constituting matter. However, for the purpose of the central point that Bach wants to make about the structural parallels between nominal and verbal ontology a focus on Link-Extractions EXT($\mathcal{C}$,$\mathcal{M}$) for which $\mathcal{C}$ and $\mathcal{M}$ contain event and process nouns is not helpful. The intuitive analogy that Bach wants to draw attention to is between the event-process relationship on the one hand and the relationship between objects and the stuff constituting them on the other: when you reflect on the matter the former relationship is structurally very much like the latter. So for the purpose of Bach's central point it is more helpful to focus on Link-Extractions for pairs $<\mathcal{C},\mathcal{M}>$ in which the sets $\mathcal{C}$ and $\mathcal{M}$ are restricted to predicates to which Link's original intuitions are most directly applicable.

4.3.2 Bach's Analogy between the Count-Mass and the 'telic'-'non-telic' Distinction

Bach proposes that the distinction which is central to Link's analysis of count nouns and mass nouns – the extensions of count nouns are conceived as 'discrete', those of mass nouns as 'continuous'; or, in the technical vocabulary we have been using: the distinction between mass nouns as 'homogeneous' and count nouns as 'non-homogeneous' or 'anti-homogeneous' – has its mirror image in the verbal domain; here it is the distinction between 'events' on the one hand and states and processes on the other.

(I have put ‘events’ in scare quotes here, since Bach uses this term in a different and more restrictive way than it has been in these Notes ever since the discussion of Vendler’s ‘Verbs and Times’. In particular, Vendler includes among what he refers to as ‘event verbs’ his activity verbs, which Bach, in his classification, categorizes as ‘process’ verbs.)

In fact, Bach’s ‘events’ vs. states-&-processes opposition can also be seen as a contrast between two types of verbs. The division that he talks about closely corresponds to the distinction between Vendler’s telic vs. non-telic aspect categories (accomplishments and achievements vs. activities and states). However, as we will see below, the relation between Bach’s and Vendler’s aspectual distinctions is more complicated; they do not necessarily apply to the same things.

Perhaps the most central, and also the hardest question about aspect:

(77) What is aspect *about*: What do aspectual notions apply to, to the eventualities described by verbs and verb phrases, or to the verbs and verb phrases themselves?

Let us keep this question firmly in mind. We will return to it more than once.

Here is the list of examples Bach gives in the section where he introduces his algebraic treatment of aspect.

(78) a. John kiss Mary: atomic event;
b. Mary stumble and Mary twist her ankle: plural event;
c. Mary stumble: atomic event;
d. People discover the hidden cove: plural event;
e. Sam build a cabin: atomic event;
f. Sam pound in a nail: atomic event;
g. Jones poison the populace: atomic event;
h. Jones poor poison into the water main: atomic event.

All the eventualities described by the sentences in this list are of the 'event' sort (in Bach's terminology). But there is a difference that opposes (78.b) and (78.d) to the other items in (78). (78.b) and (78.d) describe *pluralities of events* rather than single events. In the terminology we have been using in our discussion of the semantics of count nouns: What (78.b) and (78.d) describe are non-atomic entities in the domain of events. The other examples in (78) describe atomic events. So, exploiting one analogy with Link's account of ontological structure in the nominal domain, (78.b) and (78.d) describe elements from the $^{+}$-extensions (and thus of the *-extensions but not of the extensions) of the event predicates they represent, whereas the other examples in (78) describe elements of the extensions of the given event predicates.

As Bach observes, there are certain English expressions that we understand as qualifications of event pluralities. One of these is the adverbial *three times*. *three times* expects the description of an event plurality with which it can combine, and when it does not get such a description the result is infelicitous. Illustrating example, also from Bach's paper, are given in (79).

(79) a. John fell asleep three times during the night.

b. ? John slept three times during the night.

c. ? John was sleeping three times during the night

Bach notes that (79.b) is harder to interpret than (79.a) and the same is true for (79.c). The reason is that the phrase *fall asleep* qualifies as an event description, so that (79.a) can be interpreted as saying that during the night there were three events of the kind described. The corresponding phrases in (79.b) and (79.c), *sleep* and *be sleeping*, do not qualify as event descriptions and that accounts for why (79.b) and (79.c) are strange.

The examples in (78) also put into view another point that is reminiscent of what we have observed for the nominal domain. It is illustrated by the pair (78.e,f). Pounding in a nail can be one of the many events that go into building a cabin. In fact, it can also be part of the event of nailing down the rafters, which in its turn is part of putting the roof on, which in its turn is part of building the cabin. For Bach all these events are atomic, but at the same time they stand to each other in part-whole relations. The case here seems closely similar to what we have been saying about the noun *part* in our discussion of Link: objects. i.e. the extensions of count nouns, can stand to each other in part-whole relations and these relations can be quite complex. The relations between all the different events that go into the building of a cabin provide an example of this in the verbal domain.

Interlude: Event Identity

A second point is illustrated by the pair (78.g,h). This point has been very prominent in one part of the literature on events, viz. in the philosophy of action. The philosophy of action is a part of the theory of events because actions are kinds of events, if admittedly a quite special kind which has close connections with the different branches of practical and norm-directed philosophy, such as ethics and political philosophy. It is especially within the theory of action that relations between event descriptions like those in (78.g,h) are of great importance. The point of (78.g,h) is that the two sentences can, in suitable but perfectly natural contexts, be used to describe what is arguably the same event, but in ways that are so different from each other that it need not be obvious that they are being used as descriptions of the same event. In particular, the agent of the event described in (78.h) may be unaware that the event it describes could also be truthfully described in the words of (78.g): he may have poured poison into the water main for some quite different purpose (to get rid of bed fish, Bach suggests) and will be likely to be horrified what his action led to, which its description in (78.g) so plainly reveals.

This issue too has its counterpart in the nominal domain, although it has played a far less prominent part there. But a simple example will suffice to show that it does turn up there as well. Suppose you have just picked up a tiny screw from the floor and that this object can be correctly described not only as a 'screw' but also as a ' part of Mary's extremely precious mechanical wristwatch'. Knowing that the second description applies, and not only the first, may make a dramatic difference to what you will do with this little screw, somewhat like the dramatic difference between events that can be described by (78.h), but not by (78.g) and those that can be described by (78.h) as well as (78.g).

Neither Link nor Bach have much to say about these modal and epistemic aspects of thing and event descriptions, and in these comments we won't pursue this topic either. But especially in the verbal domain the modal/epistemic dimension is of great importance, if only because event description is a crucial ingredient to planning: when we make plans about what to do, the actions we contemplate and the events we expect or consider possible are available to us *only* as the possible instances of our conceptualizations or descriptions of them. But as we will see below, there is also an important modal dimension to the structure of events, processes and states that is not tied in this direct and irrevocable way to event, process and state descriptions.

One last remark on the subject of not transparently equivalent event descriptions like those in (78.g) and (78.h). Some philosophers (Kim, Goldman) hold the view that different event descriptions always describe distinct events (unless they are equivalent for 'analytic' reasons and their equivalence should be obvious to anybody who knows the language in which the descriptions are stated). On this view the agent of the events described in (78.g) and (78.h) has in fact performed two events, that described in (78.g) and that described in (78.h); and, by the same token he has performed, on that same occasion, untold numbers of other events, such as the one described as 'John pour poison into a water pipe that runs of the Main underneath Broad Street', and so on. An early, forceful and persuasive voice against this was Davidson's. Davidson argued that events are individuated in terms of their temporal and, when they have one, spatial location and by their roles as causes and effects in a network of causal relations that gives our world its causal coherence. Even now it cannot be said that Davidson's view is accepted universally, but if my estimates are correct, then a substantial majority of today's philosophers side with him and I doubt that there are any linguists who have given the matter any thought and would side with Kim and Goldman rather than with Davidson.

That Bach is also on Davidson's side is not only indicated by some of the things he says, but by the very analogy between nominal and verbal ontology that is the crux of his paper. This is so because what some philosophers consider a point of debate in relation to descriptions of events, has never been seriously considered an issue in relation to the descriptions of things. Consider for instance the screw that John just picked up from the floor and that is now resting in the palm of his right hand. It can be described both as the 'screw that John just picked up from the floor' and as the 'screw that is missing from Mary's wristwatch'. These descriptions are not transparently equivalent and in fact it is quite possible that John does't realize this; he doesn't know that what he is holding is missing from Mary's watch. But no one in his right mind would take this observation as a possible reason for thinking that there are really two screws in the palm of John's hand, or – as one would then be forced to accept as well – some indefinite, very large number of screws, each of them corresponding to a different description of the screw he is holding. (It is unlikely that John could hold that many screws in the palm of his hand (even if they were all this small), and if he could he would probably collapse under the weight.

In short, anyone who takes the analogy between the nominal and the verbal as seriously as Bach should see the relation between events and their descriptions as obeying the same general logical principles that we take to govern the relation between objects and their descriptions; and that will place him, as far as the verbal domain is concerned, squarely within the Davidsonian camp.

Bach's Algebraic Structures

Let us, before we address some of the differences between nominal and verbal ontology, present the structures of events, states and processes that Bach proposes and that are meant to bring out the close parallels between them. In doing this we have to improvise somewhat, since Bach doesn't go into much detail about his algebraic formalization. However, given the quite detailed proposals we have been making in our exposition of Link we may as well state the structures Bach has in mind in the format we have been led to adopt in the course of our explorations of Link's ontology. Once more we assume that our structures are relative to a pair of sets of predicates – a set of 'event' predicates $\mathcal{EV}$ and a set of state and process predicates $\mathcal{PR}$. And once again these two sets should be attuned to each other, in the same sort of way that we argued earlier $\mathcal{C}$ and $\mathcal{M}$ must be in tune. (We will return to this point presently.) Bach adds to his algebraic event structure the relations $\circ\!<$ of complete precedence and O of overlap between eventualities, taking these temporal relations to be part of the essence of event structure. There is no direct counterpart to this in the nominal domain; the closest analogy is with the location of material entities within the structure of space.

Here is my minor adaptation, for a pair of sets of predicates $< \mathcal{EV}, \mathcal{PR} >$, of the structure that Bach proposes:

(80) $<D, E, \sqsubseteq_e, \sqsubseteq_p, h, \circ<, O>$

$<D, \sqsubseteq_e, \sqsubseteq_p, h>$ must satisfy the same structural constraints that we articulated in connection with Link's mereological ontologies for sets of count and mass predicates. In particular, h now maps each 'event' (in Bach's sense) onto the process without which the event could not have occurred. D is the totality of all events, processes, states and their pluralities, E the set of events and their pluralities and $D \setminus E$ the set of processes and states (and their pluralities, if any). (I will turn to the relation between states and processes towards the very end of these comments on Bach's paper.) $\sqsubseteq_e$ and $\sqsubseteq_p$ correspond to the part-whole relation between elements of the *-extensions of count nouns and the part-whole relation between elements of the extensions of mass nouns, respectively and h maps D onto $D \setminus E$, the identity function on $D \setminus E$ and (of necessity) disjoint from the identity function on E. We may ask, as we did for Link-Extractions, if it isn't possible to reduce these two primitive part-whole relations to a single one. The possibility for that would now rest on (i) the question whether what may seem pluralities of processes and states aren't on closer inspection always pluralities of corresponding events, and (ii) if there is a general operator, and one realized in our language, that turns process or state descriptions into event descriptions which can always be used to refer to the apparent pluralities of processes or states described by those process or state descriptions. Some of the examples Bach mentions indicate that process descriptions do not combine happily with what appear to be eventuality counters such as *three times* (see (79.b,c)) and that suggests to things: (a) that pluralities of bits of state or process do not exist after all, and (b) that state and process descriptions cannot be coerced without further ado into descriptions of corresponding events. (If they could, then (79.b,c) wouldn't have got the question marks they have (and it seems rightly so) in Bach's paper, even if it had been true that there are no pluralities of bits of state or process.) The question whether there are general linguistic ways of turning state or process descriptions into descriptions of corresponding event (i.e. events which h turns into the states or processes described by the input descriptions) does not seem to have an obvious answer. I do not think there is any 'functional' event verb that takes process and state descriptions as inputs (e.g. in the form of an infinitival or gerundive clausal or other verb projection; a verb V, say, that allows us to form 'V running' or 'V being sick' in the same way that we can form 'start running' or 'start being sick', but where 'V running' would mean the same thing as 'had a run' and 'V being sick' the same as 'had a bout of illness'). In fact, as indicated by the final observation in the parenthetical part of the last sentence, it may well be that the closest that we can come to a linguistic realization of such a transformer of process and state descriptions into event descriptions is via nominalizations in conjunction with portmanteau verbs like *have* or *go through*, as in *have a run* or *go through a bout of illness*.

How generally such 'verbal classifier constructions' are available in English or other languages is a question we will not pursue any further here. But it is at this point clear enough, I think, what such a general transformation of process and state descriptions into event descriptions should accomplish to feel justified in assuming an abstract operator that does precisely this, whether or not we have linguistic mens to express all possible applications of this operator. If we assume that there is such an operator, then reduction of $\sqsubseteq_e$ and $\sqsubseteq_p$ to a single primitive relation is possible just as we argued for this possibility in the nominal domain. But rather than make this assumption and change definition (80) accordingly, let us stay with eventuality structures in the form in which (80) specifies them. For what follows that formulation will do just as well.

Let us return to more homely ground: What are the intuitions that support Bach's structures as defined in (80)? The basic intuition is easiest to explain in connection with the events described by accomplishment phrases like 'write a letter' and the activities without which such events could not have come about. Each event of writing a letter involves a sequence of things happening, the movements of the pen on paper or of the writer's fingers across the keyboard, or the flow of speech from the manager into the dictation machine from which her secretary will produce a written copy. Arguably there are various other actions that will be part of letter writings in each of these cases, so in practice it will be hard to determine precisely what should be considered part of the entire process that is involved in any particular letter writing event. But the intuition that there must be various things going on in order for there to be a completed event of letter writing seems (I think) plausible enough, and so (I think) is the intuition that the event should not be identified with the totality of those things.

How to justify the transition from such cases, for which these intuitions seem quite firm, to the general relation between what is described by event descriptions and what is described by process descriptions in general is another matter. This is obviously a big step, the implications of which are not easy to fathom. There is a lot here that needs careful reflection and analysis. But in these comments we follow Bach in his assumption that such a generalization is possible and justifiable. There are however a number of general questions that arise in connection with eventuality structures and that have either no immediate analogue in connection with Link-Extractions or else a possibly different answer, and some of these we should at least mention even if we are in no position to answer them. One question is whether with each event (in Bach's sense) there is associated a process that stands to it in the same sort of relation in which the stuff constituting a material object stands to the object that it constitutes. Another one – related to, but definitely distinct from this one – is the ontological nature and structure of events and processes. Here the parallel Bach perceives between Link's nominal ontology and his own proposal for the verbal domain make his convictions plain enough: Event predicates like that in (79.a) have a 'discrete' semantics – more formally, they have extensions that consist of atoms in the sense of the relevant part-whole relation and *-extensions extending these – and state and process predicates like those in (79.b) and (79.c) have a 'continuous', or 'homogeneous', semantics, which is formally reflected by atom-less extensions, with a part-whole ordering in relation to which each element is a non-atom. This general intuition is confirmed well enough by our intuitions about event, process and state descriptions, as we have seen and discussed them extensively in earlier parts of these Notes: if e is a particular event exemplifying an event description like *Mary write a letter*, then no proper part of e will qualify as an instance of *Mary write a letter*; but a proper part of an instance e of a process description like *Mary run* or of an instance s of a state description like *John be asleep* will still be something describable as an instance of Mary running or of John being asleep. In other words, the extensions of process and state descriptions are structured by a principle of divisibility that mirrors the divisibility of the extensions of mass nouns, whereas the extensions of event descriptions lack this property, just as the extensions of count nouns.

But what about closure under additivity, which in Link's ontology is another aspect of the homogeneous character of the extensions of mass nouns? Here there appears to be some difference between the nominal and the verbal domain. For mass nouns we assumed that their extensions are closed under the formation of mereological sums: for any two portions in the extension X of a given mass noun X contains their sum as well. The sum may be present in X only as a virtual entity, even when its summands are present more concretely, as currently well-circumscribed bits of the stuff (such as the lump of butter on my plate and the lump of butter in the fridge), and it may even be the spatially distributed material sum opt portions that are themselves only virtual, e.g as the potential results that could be but haven't actually been made.

Admittedly this is all highly abstract. But what can we say about the analogous assumptions regarding the extensions of process and state predicates? Suppose that Mary did some running this afternoon and some other bit of running yesterday morning? Does it make sense to say that there is also a 'temporally distributed' process of her running one part of which is temporally located within yesterday morning while the other is located within the afternoon of today? Well, that is not so easy to say. But on the other hand it seems to be what Bach's general ontological set-up commits us to. For it is clearly part of that set-up that if Mary had a run e_1 this afternoon and another run e_2 yesterday morning, then there also exists the plurality $e_1 \cup_e e_2$ formed out of these two running events. But then the 'processual substance' of that plurality – the total process which consists of what was involved in yesterday morning's run on the one hand and in what was involved in this afternoon's run on the other – must also exist as a process; otherwise h could not relate non-atomic events and processes in the systematic way that is part of the way in which Bach's ontology structures are set up.

The Modal Dimensions of Homogeneity

But there is also a side to the additivity of states and processes that has no obvious analogue in the nominal domain. Processes or states that occur in the actual world and that, for whatever reason, come to an end at some point could have gone on for longer and those possible longer states and processes would still be instances of the same process or state descriptions that we use to describe the states and processes that actually occurred. Suppose for instance that s is a state of John's being asleep, which started at some time after lunch and that ended at three o'clock in the afternoon – that was when, for that he would have gone on sleeping for a little while at least. So in this other imagined world there would have been a state of John sleeping that would have ended at some later time than his sleep in the actual world. There is a natural intuition that this longer lasting state in the possible world imagined is not only an instance of the same state description 'John be asleep', but that it is in fact the very state of John being asleep in the actual world, but with the difference that in the imagined world that state has a longer duration. Likewise for the process description 'Mary run'. Suppose e is na actual process instantiating this description. At some point Mary stopped running, but it is (let us assume) reasonable to assume that Mary could have gone on running for some while longer: the world could have continued somewhat differently from the point at which she stopped running from the way it actually did, and in that alternative continuation of the world Mary would have gone on running for a little longer; in other words, there would have been in that alternative continuation of our world a process e' that also answers the description 'Mary run', and that might even be considered the same process as Mary's run e in the actual world, except that in the imagined continuation the running would have spanned a longer interval. For event descriptions, such as 'Mary write a letter', this is not so. If e is a complete letter writing event, then in no continuation of the world from the point where the letter is finished and e reaches its culmination could there be an extension of e which also fits the description 'Mary write a letter', let alone that such a continuation could be thought of as part of the same event as e. Mary could go on with her letter writing after the point where e culminates, but that would then be the writing of a different letter, and if she were to complete that letter, then in the imagined world there would exist a plurality of two letter writings, not a single long event instantiating the description 'Mary write a letter.'

I believe that this modal dimension to our intuitions about how state and process descriptions differ from event descriptions is an important part of our understanding of the difference between them. Ideally, therefore, a formal ontology of states, processes and events should incorporate a modal dimension as well, in the form of worlds that can develop in different directions, giving rise to a kind of complex tree of 'possible histories', with an immensely rich structure of branchings in the direction of the future. But that would complicate the structure considerably and it is something that Bach does not consider. Since these are comments on his paper, we follow him in ignoring the modal dimension here. But the issue merits further exploration.

Modal considerations also have application in the nominal domain, but here they do not seem to be as central as they are in the realm of events, processes and states. However, once ontological structures for the verbal domain are enriched with a modal dimension, then it would be natural to develop a similar extension for the nominal domain, so as to preserve, in the spirit of the general thrust of Bach's paper, as much of the formal analogies between the two domains as we can.

One crucial (and obvious) difference between nominal and verbal ontology is that events, states and processes happen in time and as soon as they happened or been, they are gone irrevocably and become part of the past that is accessible to us only through testimony. There is nothing more you can do to change them; they have become unalterable fixtures of history. The entities that make up nominal ontology are crucially different in this regard. They may have, and pretty much all of them do have, a limited life-span, coming into existence at one time and going out of it again at some later one. But while they exist different things can happen to them at different times; they can have their own limited histories, made up from the succession of what happens to them in the course of their existence. This is one reason why the distinction between objects and the stuff from which they are made appears so firmly grounded. Bach brings up the example of Terry's ring, which must be distinguished from the gold of which it is made, for one thing because the gold goes back to the early days in the development of outr planet, or even beyond those, and thus is much, much older than the ring, as artifact, with its particular shape and purpose. For another example of the same sort, suppose you give me a lump of plasticine with the suggestion that I model it into the representation of something. I first make it into an armadillo, but dissatisfied with my effort I kneed the armadillo back into a shapeless lump and then turn that into a model of the Capitol of the State of Texas. Pleased as I may be with this second creation you insist that I return the plasticine to you and destroy my Capitol, turning it back into a shapeless lump you had handed me, so that it can serve again for the educational purposes to which you put it. The stuff from which my two artistic creations were made, the plasticine, can be expected to be of much more recent date than the gold in Terry's ring. But it goes back to a time before my two creations and it has outlived both of them. So here two the difference in life span between the given quantity of plasticine, armadillo and Capitol is a compelling reason for regarding them as three distinct entities in the world.

For the distinction between events and the processes that on Bach's theory are constitutive of them no such life-span-based arguments are available. Event and process are by their very nature coeval. It is for this reason that in justifying the kind of ontology that Bach proposes a greater emphasis must be put on our understanding of the semantics of event, process and state *descriptions*. More or less all of the *evidence* we have for the assumptions we make about the structure of verbal ontology is based on that understanding. Elsewhere Bach has coined the term 'natural language metaphysics' to describe the task of deriving and justifying ontological structures like those he proposes for events, processes and states on the basis of an analysis of the conceptual implication of the logical forms of the languages we speak and the uses we make of those forms. In relation to verbal ontology the term seems especially apt, as so much of what we can say about this ontology can be motivated only by careful reflection on how that ontology manifest itself in the linguistic forms that we humans use to talk about it.

In the light of these considerations Bach is right to include temporal relations explicitly as components of the formal event ontologies he proposes (see (80)). So far we didn't say anything about this component. There isn't a great deal that needs to be said about them, but it is high time to say it. And what little needs to be said we can say quite concisely in the light of what has been articulated earlier in these Notes in considerable detail:

(81) The temporal relations $\circ<$ and O of any structure given in (80) are binary relations on the set D which satisfy the postulates for event structures given on pp. 27 -30 of these Notes.

In the light of what we have been saying when discussing Vendler this would seem to imply that achievements (the events described by achievement verbs and verb phrases) are mapped onto minimal, or atomic, processes, which would contradict the assumption that processes are infinitely divisible – just as we saw there to be a contradiction between the infinite divisibility of the extensions of mass nouns and the assumption that matter is composed of atoms and molecules as their smallest parts.

Aspectual Distinctions: What is it that they apply to?

There are obvious parallels between Bach's bipartite division of eventuality descriptions and the eventualities they describe – into events on the one and processes and states on the other hand – and the aspectual divisions made in Vendler's 'Verbs and Times'. Vendler's classification is fourfold, but as we have seen, the arguably most important division is between accomplishments and achievements on the one side and activities and states on the other. That too is a binary division and when we adjust for the somewhat different terminologies, the two divisions seem closely similar.

But there is nevertheless an important prima facie difference between these two classifications with regard to what they are supposed to apply to. Vendler's classification is in the first instance a classification of verbs, first and foremost lexical verbs, but then also phrases consisting of a verb together with a direct object, or some propositional phrase; but even these compound phrases do not stray very far from the lexical verbs from which they are built. Bach's classification, on the other hand, targets phrases which, syntactically speaking, are situated in the upper reaches of sentence construction. For instance, the examples in (79) are complete sentences and those in (78) are 'sentence radicals', phrases that are complete sentences except for lacking a finite tense. (Bach assumes that putting in the finite tense is the last operation in building the syntactic structure of a complete finite clause.)

One of the hard and still unsolved problems of natural language semantics is how, in English and other languages, the aspectual properties of lexical verbs are connected with the aspectual properties of full clauses. The connections appear to be quite different for different languages, even if we restrict attention to those that have any tense morphology at all. And for English (and other languages of, in particular, Western and Central Europe, such as the Germanic and Romance languages), the problem is a complex and difficult one because there are many different factors that contribute to the aspectual properties of full clauses and sentences and many different points in the course of constructing a complete clause or sentence from the words and morphemes that make it up at which these factors can make their impact, by changing the aspectual settings as they had been determined up to that point. For the most part the aspectual properties of the main verb of a clause play an important, and often a decisive, role in determining the aspectual properties of the clause, but as a a rule they do so through interaction with the many other factors that play their part in shaping the ultimate aspectual profile of the clause as a whole.

In these Notes (and in the course they accompany) we haven't said anything about this very complex set of issues so far, and we won't have much to say about it later on. (An exception is the discussion of the mechanisms of 'coercion' in the paper 'Temporal Ontology and Temporal Reference' by Moens and Steedman on which there are comments later on in these Notes.) But it is important for us to be at least aware that there is a kind of gap between the properties of lexical aspect that are the focus of someone like Vendler and the sentence level aspect which is the primary concern of linguists like Bach.

The aspectual properties of complete clauses and sentences are of particular importance for the interpretation of discourse. This will be the central theme of the two papers by Partee, which are the two next in line for our discussion. As shown in those papers, tenses and other devices that affect aspect and/or temporal reference play a crucial role in the information they carry about how a new sentence in a discourse or text, with its particular tense and, sometimes, some of those other devices, is temporally and aspectually linked to the interpretation of the sentences preceding it. In the second paper Partee develops a theory that derives semantic descriptions of multi-sentence discourses in which there is explicit reference to the different events and states. (Processes in Bach's sense do not play a separate part in her theory, but are treated as a species of ('dynamic') states.) These discourse descriptions are obtained from the eventuality descriptions contributed by the individual sentences, and the theory can therefore be understood as a complex algorithm for integrating the eventuality descriptions provided by individual sentences into these discourse descriptions (whose structural organization is somewhat different from that of the sentences, so that the integration involves modification of the sentential eventuality descriptions as well). But while a formulation of Partee's theory at a level of mere description transformation is possible (and she develops such an account in considerable detail), the resulting discourse descriptions are naturally interpretable – and meant to be interpreted – as descriptions of eventuality complexes of the kind that are found in Bach's ontological structures.

In short, Bach's ontology may be thought of as the 'real world' counterpart to Partee's discourse representations: a discourse representation à la Partee is a *true* description of the world of which it speaks if the events, processes and states it represents can be found in the Bach ontology for that world.

This way of putting things suggests that Bach's ontology is an ontology for the events, states and processes described by complete sentences and clauses and the discourses and texts of which complete sentences and clauses are the building blocks. How complete sentences and clauses come to be the eventuality descriptions they are, via a compositional process that computes their aspectual and other semantic properties from the semantics of the words out of which they are constructed syntactically, is a further problem, and as noted above, it is one that appears to vary from language to language. But whatever the details of these computational mechanisms and the principles that govern them, these are all mechanisms that convert eventuality descriptions of one kind, for the lower nodes in the syntactic tree, into descriptions of another kind, associated with nodes the next level up in the tree. What the ontological correlates are of the eventuality descriptions that result along the way of these compositional computations is not always easy to determine. But there are at least two places where contact must occur. One, we have just seen, is at the end of the compositional process, where the eventuality descriptions of complete sentencers and clauses must be interpretable as applying to the world about which the sentence or discourse is used to say something. But another one is at the 'bottom' – i.e. at the start of the compositional process –, where our knowledge of the semantics of lexical verbs manifests itself among other things in our ability to understand and test Vendlerian classifications.

The evidence that Bach adduces in support of the central analogy of his paper seems to appeal primarily to our knowledge of lexical semantics (see (78) and (79)). (That Bach advocates a three-way distinction between vents, processes and states rather than the four way-distinction of Vender or the five way distinctions of those who include semelfactives as a distinct category, is of no relevance to the general point we are concerned with right now.) But it seems to me that the picture Bach wants to convey in *The Algebra of Events* is based on what we can observe in simple sentences and clauses, in which the relation between sentence aspect and lexical aspect is the most direct we. For it is for those cases, and for those cases only, that it makes sense to apply our lexical intuitions to what look like sentence radicals.

However, if the picture conveyed in Partee's papers is correct, according to which we are only dealing with a binary aspectual distinction at the level of complete sentences and clauses, then there must be some indirectness in the relation between lexical aspect and sentence aspect even for such sentences and clauses, in which the relation is as direct as possible: somewhere a transition has to be made from the three-way or more-than-three-way aspectual distinctions at the lexical level to the two-way distinction at the top. How this transition is made in sentences of different syntactic form is, it turns out, a particularly complex problem for English and many other European languages, in which the aspect of a complete clause or sentence will often depend on a variety of factors and not just on the choice of verb. My understanding of the complex mechanisms that determine sentence aspect in such languages is still fragmentary. But minimally we should retain this much from the discussion above: The ontology of the world which determines whether a given sentence or discourse that makes a claim about a certain part of it is true should be assumed to take the form of Bach's eventuality structures. For someone who accepts the approach developed in Partee's *Nominal and Temporal Anaphora* this boils down to the question whether the representation for the sentence or discourse is true of the real world about which a claim is being made. And in view of Partee's assumption that those representations only distinguish between events and states, it should be only the events and the states of the world's ontology that are directly involved in the question whether the sentence or discourse representation (and, with it, the particular sentence or discourse it represents) is true. But if that is so, then more needs to be said about the relations between states, processes and events in Bach's eventuality ontologies than we have so far done. The final remarks of my present comments on Bach's paper are a first stab at this question.

Natural Language Metaphysics vs. Real Metaphysics

If all we pay attention to is that (i) at the 'bottom' – there where the lexical verbs make their contributions to sentence aspect – the relevant distinctions are threefold or more, that (ii) at the 'top', where sentence aspect has finally been determined, on the basis of the aspectual properties of the verb and, often, much else, the distinction is a binary one, describable as that between events and states, and if (iii) it is further observed that the notions 'state' and 'event' already play their part at the bottom level; then one might be tempted to think that what happens in the cpurse of compositional aspect computation is a classification *reduction*: from a more detailed classification, which includes the categories of 'state' and 'event', to a simple one in which only those two categories survive. For instance, if one assumes, as we have done with only minor adjustments throughout these notes, Vendler's four-fold classification of verbs into accomplishment, achievement, activity and state verbs, the task may seem to be that of explaining how the first three categories are amalgamated into a single category of 'event', while the category 'state' retains its significance throughout the compositional process. Likewise, if one adopts Bach's three-way classification into events, states and processes, there is reduction from a threefold distinction to a binary one, in which the concept 'process' has been discarded, presumably because it is now redundant.

More than enough has been said in these notes to make it obvious that this is a picture that bears little resemblance to the way semantic composition works. It is possible to make use in a meaningful and explanatory way of a binary division between events and states at the sentence level not just because of the more than 2-fold aspectual classification at the bottom only that between the states and the events which is part of that classification has retained its relevance. It is also the case that many of the states that, in the kind of approach we are talking about, *are* relevant at the sentence level are different from the states that can be described by lexical state verbs. (Early on in his paper Bach speaks of 'dynamic states' and to the extent I can tell, it is these new, non-lexical states that he has in mind when using this term.)

That complex phrases at projection levels well above that of the lexical verb should be descriptions of kinds of states that are not described by any lexical verbs need not be a reason for surprise. But the problem we are facing is not just that of a larger and more diversified ontology of states than that needed to account for the semantics of lexical state verbs. The real problem is a different one: What is the *status* of these new states? This problem is an instance of a much more general one, which can be described, using

Bach's own terms, as the relation between 'Real Metaphysics' and 'Natural Language Metaphysics' (Bach (1986)). 'Real Metaphysics' is that branch of philosophy which is concerned with the nature of the world - what kinds of things there are, in what fundamental ways they are related to each other, and in what sense (or senses) different categories of things can be, or *exist*. This is a branch of philosophy that goes back to Greek Antiquity, with Aristotle as its most prominent representative and arguably its founder, and it has remained a central concern throughout the history of philosophy. But, Bach argues, for the semantics of natural languages we need something else – not an inventory of what, to the best of our scientific knowledge, the world ultimately consists of, but what entities are implicitly assumed by the speakers of any given natural language. For the enterprise of identifying the different kinds of these entities and their fundamental relations bach introduced the term 'Natural Language Metaphysics'. The 'dynamic states' that arise as bearers of eventuality descriptions at the sentence level (as for instance in the approach opt Partee in the second of the next two papares we will discuss) should be thought of as belonging to the 'Natural Language Metaphysics of English' rather than to Real Metaphysics.

There is no reason to assume that Real Metaphysics and Natural Language Metaphysics are disjoint, no reason why some of the entities that are needed in Natural Language Metaphysics should not be part of Real Metaphysics as well, no reason why the intuitions about what there is that are enshrined in the ways we speak should not coincide with what there is according the best of our scientific knowledge. In fact, in practice the dividing line between Real and Natural Language Metaphysics is difficult to draw, and where it is drawn, it is typically a reflection of the theoretical convictions or *parti-pris* of those who draw it. The category of events is a good example of this. For those who think of Real Metaphysics as a reflection of the fundamental insights of modern physics, macroscopic events like that of Russia's annexation of the Crimea or my having a shower would not be included among the basic ingredients of our world, while they are indispensable ingredients of Natural Language Metaphysics. But others, who think that Metaphysics should not take the reduction of macroscopic entities to microscopic complexes for granted, may want to see such macroscopic events as part of the world as it exists independently of how we talk about it in the languages we speak, and thus as part of Real Metaphysics. For them such macroscopic events will belong to the overlap of Natural Language Metaphysics and Real Metaphysics.

Let us, for the sake of argument, assume a Real Metaphysics that counts macroscopic events as part of the 'furniture of the language-independent world', and let us assume that the same is true for 'lexical' states like that of Mary being in Paris, of Bill being in pain or of Paris being in France. These events and states, in other words, are assumed to belong to Real Metaphysics as well as Natural Language Metaphysics. What then are we to say about the various kinds of 'dynamic states' that Bach alludes to and that are needed to account for sentence aspect and for its compositional computation from the meanings of the different constituents from which complete sentences are put together? In order to approximate an answer to this question we need to have a look at two things: (i) what is the status of the event-state distinction at the level of complete sentences and discourses and (ii) what are the kinds of compositional constructions that introduce dynamic states in the course of the compositional computation of sentence aspect (as illustrated in the comments on the second paper of Partee later on in these Notes).

on an account like Partee's bring 'dynamic' states into play? the two issues are closely connected and I will discuss them together.

One of the operations that are involved in the semantic composition of many English sentences is, we have seen, the Progressive. And one way to analyze the Progressive is as an operator that transforms descriptions of events into descriptions of states. Consider in this connection the following examples.

(82) a. First John vacuumed the living room. Then he did the dishes. Then he made the beds and finally he dealt with the study.

b. Marie came in through the kitchen shortly before three. John was doing the dishes. She said hallo and went straight to her study.

(83) a. When Alan opened his eyes he saw his wife who was standing by his bed. She smiled.

b. When Alan opened his eyes he saw his wife who was standing by his bed. She was smiling.

The point of the two examples in (82) is this: Both (82.a) and (82.b) can be used as partial descriptions of an episode in the life of Marie and John, in which John spends part of the afternoon dealing with various bits of housework, a succession of tasks of which doing the dishes is one. (82.a) presents what John did as this series of successive events, by means of sentences with event verbs in the simple past. (82.b) describes a part of the episode described in (82.a) in more detail, focussing on the spouse who is coming home and finds John in the process of doing the dishes as she enters the kitchen. Here the doing of the dishes is referred to by a sentence in the Past Progressive. So, on the assumption that progressive sentences describe ('dynamic') states, the very same part of the episode that is described as an event in (82.a) is described as a state in (82.b). But what exactly does that come to? Is there besides the event described in (82.a) – that of John doing the dishes – *in addition* a state of his being in the process of doing the dishes? That appears to be what the analysis of non-progressive and progressive uses of event verbs commits us to. But is this reasonable?

A different but partly similar point is illustrated by (83). Here the two bits of discourse, (83.a) and (83.b), are to be thought of as describing two different possible episodes, one in which Alan's wife smiles at him in reaction to his opening his eyes and seeing her and one in which she is smiling already at the point when he opens his eyes and her smile is what he sees when he sees her. The reason why (83.a) is an adequate description of the first episode and (83.b) an adequate description of the second is that the Simple Past *smiled* in (83.a) is naturally interpreted as indicating that the event described follows the events described by the preceding sentence, whereas the Past Progressive *was smiling* in (83.b) is naturally interpreted as indicating that the smiling was going on while the events of the preceding sentence took place.

Here too the analysis under discussion commits us to the position that a smiling *event* is described in (83.a) and a smiling *state* in (83.b) . But from the perspective of what entities make up the two episodes that (83.a) and (83.b) could be used to describe, this distinction is less than compelling. Whether the smiling occurred after Alan opened his eyes and saw his wife or started before he opened his eyes and saw his wife and continued while he did seems immaterial to what sort of entity the smiling was – if it was an event in the episode described in (83.a), then it should also be an event in the episode described by (83.b); and conversely, if the smiling described in (83.b) is a state, why shouldn't it also be a state in the episode described by (83.a)? Here too, then, it looks like our analysis commits us to a duplication of entities, with states 'reduplicating' events, for which a metaphysical justification would be hard to come by.

In fact, it isn't just in terms of Real Metaphysics that the duplication of events by states is hard to justify; even a justification in terms of Natural Language Metaphysics seems problematic. Isn't this a case where the existence of certain entities is forced upon us by a particular choice of linguistic analysis, rather than a commitment implicit in the language to which that analysis is being applied, which could perhaps be analyzed just as well in some other way that doesn't carry this commitment? To achieve some clarity on the point of this last question, we have to focus more closely on the various constructions for which it appears reasonable to hold that they describe states and of which at least some versions of the approach of Partee's papers claim that they do describe states.

In English and many other languages there are quite a number of such constructions. Besides the English Progressive (and constructions in other languages to a similar effect), there is the Perfect in its different forms (Present Perfect, Past Perfect, Future Perfect, infinitival Perfects), which can be analyzed as yielding descriptions of *result states*, and which is analyzed as a result state operator in the versions I am talking about. Other operators the outputs of which are naturally interpreted as state descriptions are negation and various forms of quantification, including habitual and generic quantification. (This is not meant to be an exhaustive list, but it is plenty for present purposes .) Thus, each of the sentences in (84) is, according to the assumptions in question, to be analyzed as the description of a state.

(84) a. Bill didn't submit his abstract (yesterday).

b. Whenever John makes a telephone call, he first smokes a cigarette.

c. Whales suckle their young.

For each of the constructions exemplified in (84) two questions can be raised: (i) what reasons are there for treating them as state descriptions? (ii) Isn't it possible to account just as well for the linguistically relevant fact without making the assumption that these constructions describe states?

These questions lead jus back to the examples in (82) and (83). It should have become clear from what we have said about those examples that the choice of tense forms – the choice between Simple Past and Past Progressive

– is motivated by the temporal relation in which what is described by Simple Past or Past Progressive is to be understood as standing to the event or events of the preceding sentence. By itself that may not seem a good reason for treating Past Progressive sentences as referring to states while the corresponding Simple Past sentences are construed as referring to events. But there is a further motivation for this assumption, viz. that the temporal relation between what is described by the Past Progressive sentences in (82) and (83) and the event or events of the preceding sentence is that of temporal inclusion and not that of temporal succession, as we find with the Simple Past sentences, and that in this respect the Past Progressive sentences behave in just the same way as sentences describing lexical states. (To verify this, substitute for instance *his wife/Marie was tired* for the Past Progressive sentences in (82.b) and (83.b).) This proves to be a general pattern: All the constructions that on the approach at issue play the part of state descriptions are understood as temporally related to eventualities of the antecedent discourse in the same way as sentences that act as descriptions of lexical states. Treating these constructions as state descriptions thus captures a certain generalization: In past tense discourse – but the generalization holds even more generally than that – the eventualities introduced by state descriptions are to be understood as including or overlapping the eventualities introduced by the relevant preceding sentence or clauses. In contrast, those eventualities that are introduced by sentences that are treated as describing events are normally understood as standing to earlier eventualities in a different temporal relation, often that of temporal succession.

By themselves this consideration might be less than persuasive. For couldn't one make do just as well with an account in which the constructions in question are treated as predicates of *intervals of time*? That would lose us some uniformity in that separate principles would now be needed for on the one hand the discourse level temporal relations involving states and on the other principles for the relations involving the intervals described by these constructions. But that might arguably be a price well worth paying if it gains us ontological parsimony and avoids the prolixity of states duplicating events for apparently no theory-independent reason and that many might therefore want to so without if at all possible.

But there is a further consideration in favor of acknowledging non-lexical states as part of Naturqal Language Metaphysics and it is one that I am inclined to see as decisive. It is evident that what is expressed by any of the constructions that the type of analysis whose merits and demerits we are discussing treats as descriptions of states can enter into causal relations, both in the role of cause and in t5hat of effect. An example of a non-lexical state description playing the part of cause is given in (85.a). The subject phrase of this sentence seems just as good a cause phrase as the vent-describing subject phrase of (85.b).

(85) a. That John smoked a cigarette whenever he made a telephone call was a cause of his untimely death.

b. That John made that trip to Afghanistan was the cause of his untimely death.

c. That John had angina pectoris was the cause of his untimely death.

Temporal intervals are not the kinds of entities that enter into causal relations. So if the *that*-clause of (85.a) were construed as describing a temporal interval, no coherent semantic account of (85.a) would be possible. But states do enter into causal relations, something that is plainly true of lexical states, as shown in (85.c). In short, sentences involving the constructions in question must be regarded as describing some kinds of entities that can enter into causal relations, and in the light of our observations about the temporal interpretations of such constructions in discourse, states seem the right candidates.

To some even this argument may not seem conclusive. But it is hard to better. What we would want ideally is a demonstration that non-lexical states of the kinds we have been discussing have cognitive reality as ingredients to the compositional processes involved in language interpretation by actual speakers. But there is no way of demonstrating this so long as we cannot look directly into the workings of the human mind and as we all know that is something that we cannot do – knot now and perhaps never, tr at any rate not for a long, long time. Short of a demonstration of this kind, the best we can hope for is one that an approach which commits itself to the existence of non-lexical states is superior to any account that tries to make do without such a commitment. The arguments i have sketched in favor of non-lexical states can be seen as a first attempt in the direction of such a demonstration. But a conclusive demonstration would require a much firmer grip on the full range of possible accounts of the semantics of tense and aspect, and of natural language semantics more generally, of which an account of tense and aspects would have to be an integral part; and this too is not available now any more than it was at the time when Bach's paper was written. So a demonstration of this second kind is also well beyond our current possibilities.

One last point. We began this section with the observation that at the start of the compositional process eventuality descriptions (the lexical verbs) are classified into three, four or even more semantically relevant classes and that at the end of this process a binary distinction seems to be all that is needed. In our cursory perusal of what happens in the process that connects the classifications at the bottom with those at the top we then found that in order to make sense of it we need more kinds of entities (more types of states) rather than fewer aspectual distinctions. But what *about* the initial distinctions at the bottom? What part do they play in the compositional process if all but one of them are suppressed somewhere along the way to the top? It isn't possible to address this question properly at this point of our explorations. But much of the answer will become visible in our discussion of the paper by Moens and Steedman, in which it is shown how different aspect-changing operators that eventuality descriptions may meet in the course of sentence composition are sensitive to Aktionsart distinctions like those of Vendler. Only when the initial representation of the lexical verb has passed through all such Aktionsart-sensitive operations that the given sentence has placed along the path it has to traverse on the way to the representation of the complete clause can these other aspectual distinctions be discarded, with the event-state distinction being the one that survives throughout.

4.4 Partee: Some structural Analogies between Tenses and Pronouns in English

The two papers by Partee that are the topic of this subsection of the notes focus like the paper by Bach on an analogy between the nominal and the verbal domain. But this time the analogy is not between nominal and verbal predicates, but between 'referential' phrases. More precisely, it is an analogy between phrases that have always been regarded as referential in some sense, viz. pronouns, and phrases that until had been considered anything but referential, viz. tenses, but for which Partee suggests, on the strength of the striking parallels between pronouns and tenses to which she draws attention, that they should be treated as referential elements as well.

Before we turn to the parallels first a word on 'referential'. One of the standard notions within semantics and the philosophy of language is that of a 'referring term'. A *referring term* is an expression that is used to some particular person or thing. Perhaps the most obvious examples are proper names, such *Washington* or *Washington,D.C* or *Washington Irving*. But demonstrative phrases – *this last day of our holiday*, *that hat you are wearing* – and definite descriptions – *the woman who lives next door*, *the chair in my bedroom*, *the halfway point between Dakar and Timbuktu* – are always or typically used as referring phrases in this strict sense as well. Other examples are the first and second person pronouns *I* and *you*. *I* always refers to the one who counts as the person who is making the utterance of which it is part and *you* (when used as a grammatically singular phrase) always refers to the one the speaker or author targets as addressee. Third person pronouns, finally, can also be usd as referring terms in this sense. An example is the *he* in Partee's example given here as (86).

(86) He shouldn't be in here.

But for third person pronouns, use as referring terms in this narrow sense is only one of several uses and it doesn't seem to be the most common one. Other uses are those in (87).

(87) a. The woman in the house next door, she almost ran over me.

b. Did I tell you about the woman in the house next door?. She almost ran over me.

c. Every boy in my class fancies a girl who can't stand him.

d. If one of the arrows hits the target, it's mine.

As terminology has it, all third person pronouns in (87) are cases of *anaphoric* pronouns: they all get their interpretation from some other noun phrase in the sentence or discourse, their *anaphoric antecedent*. But between these different cases of pronominal anaphoricity there are further important differences. The first two, the occurrences of *she* in (87.a) and (87.b), can still be classified as referring terms, since thy do refer to some particular referent, viz. the woman in the house next door. But the other two pronouns – the *him* of (87.c) and the *it* of (87.d) – are not referring terms in the sense defined. They do not refer to one particular thing, but rather take on varying 'referents', as a function of some set of things that their antecedent runs through, or that their antecedent can choose a referent from. (87.c) is a clear example of the first case: *every boy in my class* talks about each individual in the set it indicates (the set of boys in my class); and *him* will pick up whichever boy from that set we focus on. (Keep in mind that the sentence makes a statement about all boys in the set!) *it* in (87.d) is an instance of the second case. Whichever arrow is (the) one that hits the target, that will be the referent of *it*.

These descriptions of how the pronouns in (87.c) and *it* of (87.d) work are somewhat clumsy and not very precise. In fact, to arrive at truly satisfactory accounts of the semantic contributions that these pronouns make, isn't all that easy and it has taken logic and semantics an inordinate amount of time to reach the point where correct analyses of these pronoun uses became possible. And with regard to the *it* of (87.d) the debate over how best to account isn't over even today.

For the role of *him* in (87.c), on the other hand, we have had for some time what appears to be a fully adequate account. But that account crucially involves an account of the semantics of *quantificational* phrases like *every boy* and, in some form or other, that of a *bound variable*: *him* in (87.c) is analyzed as playing the part of an occurrence of a variable that is bound by the quantifier phrase *every boy*.

It is also possible, though not strictly necessary and therefore less compelling, to treat occurrences of pronouns that function as referring terms (as in (87.a) and (87.b)) as instances of *free* variable occurrences. It is because of this additional possibility that Partee sees (third person) pronouns as 'variables', whose occurrences can be either bound, as in (87.c) and (87.d), or free, as in (87.a) and (87.b). On the basis of the analogies between tense and pronouns she points out in the paper, she then concludes that tenses should also be analyzed as 'variables', rather than as sentential operators in the manner of tense logic.

The starting point of this introduction to 'Some structural Analogies between Tenses and Pronouns in English' was the term 'referential'. It ought to have become more or less clear by now how that term is intended: A referential expression is one that can be used to do some or all of the things that the different pronouns in (87) do – in Partee's terminology: an expression that can be used to do the job of free and/or bound variables. Referring terms, as that term was just explained, are thus one type of referring terms - those that only allow for uses that correspond to free variables. But third person pronouns, qua *expressions*, are not referring terms, since they allow for bound variable uses as well; but they can be used as referring terms.

Personally I am wary of identifying the third person pronouns of English and other natural languages with 'variables'. For when you look closely at the ways in which pronouns work, you find that those mechanisms are quite different from the regimes that govern the use of variables in mathematics and formal logic. What is true as that pronouns often do the work of bound variables in predicate logic; but how they succeed in doing that is another matter, and it os there that the differences between pronouns and the variables of formal logic become salient. Hence my preference for the term 'referential expression'. We will return to this issue when discussing Partee's second paper.
We will see that while the proposal that tenses are variables points in the right direction, it cannot be more than a first step. Partee's second paper is a first revision and further elaboration of this proposal, and we will see that even that won't be the last word. But first the evidence as she presents it in 'Some structural Analogies between Tenses and Pronouns in English'.

To fully appreciate the importance of the parallels to which Partee draws our attention in this paper, it is necessary to try and put ourselves in the fraym of mind that was dominant when the paper was written. This was the tim when Tense Logic was still thought of as providing us with a useful and truthful model of how the tenses work in languages like English. That may seem hardly credible today, for as we saw during our brief review of Tens Logic, as a model of the workings of natural language tenses Tense Logic has pretty much everything wrong with it. Still, there are some things about it that are more or less right, or at least halfway plausible; and those were enough to blind the community into the faith that here was the way forward. Partee's paper was a lucid and persuasive statement to the effect that this simply couldn't be right.

The first and still most famous example she gives to show that Tens Logic doesn't get certain things right is that in (88).

(88) I didn't turn off the stove.

We are asked to imagine this sentence as uttered by one of a couple to the other as they are just turning onto the expressway on the way to the place where they will be spending their holiday. It will be quite clear in such a setting what time the speaker is talking about: it is the time when they are closing up the house before getting into the car and driving off. It seems natural to describe the role of the past tense in this utterance, with this intended meaning, as used 'deictically', much like the pronoun *he* in (86). The two cases – of the pronoun in (86) and the tense in (86.2) – aren't perhaps completely parallel in that there is some intrinsic vagueness in precisely what time is being talked about in (86.2). But the similarity os close enough; and it is a useful exercise to see exactly why Tense Logic (more specifically: the (P,F)-calculus, with its operator P that is meant to symbolize the simple past tense) is incapable of representing the intended content of this utterance.

We won't give a real, conclusive proof that the intuitive truth conditions of (86.2) aren't matched by any formula of the (P,F)-calculus. (Such proofs are usually hard even when what is being proved seems obvious, or seems obvious after a few tries. But we will look at what seem to be the two most plausible candidate formulas for the logical form of (86.2) and see that they don't work, and why they do not work.)

Let q stand for the 'generic' proposition expressed by 'I turn off the stove' (the proposition that is true at all and only those times at which the speaker S does turn off the stove). Then there are two main candidates for the formalization of (86.2):

(i) $P\neg q$
(ii) $\neg Pq$

But neither of these will do. (i) says that there was some time in the past when Sdidn't turn off the stove. Surely that is true: there must have been countless times at which S didn't turn off the stove; in fact, the vats majority of times must have been such times; for any tim at which on dos turn off one's stove there is normally going to be a much longer period when one will not do so. So (i) would have been true even if what S is expressing by uttering (86.2) on the given occasion would have been false. (Perhaps she did turn off the stove after all.)

(ii) isn't much good either. In all likelihood S will have turned off the stove many times in the past, whether or not she did on the present critical occasion. So the proposition expressed by (ii), that it is not the case that there was a time in the past at which S turned off the stove will be false irrespective of what she did no this occasion. So again the difference between turning off the stove on this occasion and not turning it off is not captured.

It is intuitively clear what the problem is: The operator P of tense logic is *'existential'*: Pq says, at t, that there is *some time in the past of t* at which q is true. And as Partee observes, that is precisely not what what (86.2) is trying to convey. (86.2) is not about arbitrary times in the past but about some particular time. The only way to get a formula of the (P,F)-calculus that doesn't get the truth conditions of (86.2) demonstrably wrong is to add a conjunct which identifies the time that (86.2) is about. That is, we could formalize (86.2) as (iii):

$P(r \mathbin{\&} \neg q)$,

where r stands for something like 'we are closing up the house before going on this particular trip'. But how could (iii) be the logical form for (86.2) in general? (Note that this will work only if we make sure that r picks out the intended time uniquely. If r is also true of previous occasions when S was closing up the house before going on a trip, then (iii) won't work for the same reason as (i).) Where are the propositions r supposed to come from?. ((iii) would have been fine if S had made the time she was talking about explicit, as in (89).

(89) I didn't turn off the stove, when we were leaving the house just now.)

The conclusion that this first comparison between tenses and pronouns suggests is that what is wrong with the (P,F)-calculus is the existential nature of the meaning of its operators. But is this always wrong or only in certain cases, such as (86.2)? We will return to this question before concluding the discussion of this first Partee paper.

Next, tenses are like pronouns in being often anaphoric. Recall that the pronoun *she* in (90.a) and (90.b) (which are repetitions of the earlier (87.a,b)) is anaphoric to the DP —*Sheila* and that in these cases the effect of this is that the pronoun and its anaphoric antecedent *Sheila* corefer in a straightforward sense: *she* too refers to the person to whom the speaker of (90.a) or (90.b) is using the name *Sheila* is used to refer. And note once more that the pronoun's antecedent is part of the same sentence in (90.a), but belongs to a different sentence in (90.b).

(90) a. The woman in the house was backing out of her driveway just now and she almost ran over me.

b. Did I tell you about the woman in the house next door? She almost ran over me.

c. Sheila had a party last Friday and Sam got drunk.

d. Sheila had a party last Friday. Sam got terribly drunk.

e. When Susan walked in, Peter left.

f. When Susan left, she took the keys.

g. When Susan left, it was raining.

h. When we were closing up the house, I forgot to turn off the stove.

The tenses of the main clauses in (90.c-h) are anaphoric in much the same sense as the pronouns in (90.a,b). The analogy is particularly close between (90.c) and (90.a) – the anaphoric constituent in the second conjunct of a conjunctions and its antecedent in the first conjunct – and (90.d) and (90.b) – anaphoric constituent in the second of a pair of consecutive sentences and its antecedent in the first sentence. And just as in (90.a) and (90.b) the pronoun *she* picks up the referent of its antecedent *the woman in the house next door*, so the second occurrence of the simple past tense in (90.c) and (90.d) picks up the time of the event described by the first past tense clause or sentence and allows the time of the second clause/sentence to be identified with it. The sentences in (90.e-h) are like the examples in (90.c,d) in that they too contain occurrences of the past tense, viz. those in their main clauses, that are interpreted as anaphorically related to the times or events of some other clause (here the *when*-clause).

While the parallels that Partee draws attention to are undeniable (and the point of bringing them to attention crucially important at the time when she did), we should nevertheless be alert to the fact that there also is a potential difference between the tense and the pronoun case. The relationship between pronoun and antecedent in (90.a) and (90.b) is plainly that of coreference: the pronoun refers to the same entity as its antecedent (the neighbour who almost ran over the speaker). But do we also have coreference in the tense examples (90.c-h)? Well, you might say, isn't that just how you just described the case: 'the 'anaphoric' tense picks up the time of the event of the previous sentence or clause and identifies the time of the event of its own clause or sentence with it'.?

Yes, indeed. That is how I described the case, and perhaps it really is the right way to describe it; and if it is, then the parallel between the pronouns in (90.a,b) and the tenses in (90.c-h). But we shouldn't take it for granted that this is the correct description. What we are ultimately interested in – it is the only thing in these cases that we can use our speakers' intuitions to test – are the temporal relations between the events that the clauses or sentences are used to describe. Now, of course, these events can't be identical: the event of Susan walking in isn't the same event as that of Peter leaving; walkings in aren't the same events as walkings out. But is might still be the case that two event sentences are 'temporally coreferential' in the sense that the events they describe are (or were) simultaneous. But that almost certainly isn't so for all the examples concerning tense in (90). For instance, the most likely scenario to be described by (90.e) is one in which Peter leaves *upon* Susan walking in: the leaving happens *after* the entering. And in (90.c,d) the event of Sam getting drunk is naturally understood as happening during the event of Susan's party. But that still need not be simultaneity in the strict sense in which two events are simultaneous if and only if their durations fully coincide. For it may have taken Sam only a small part of the time the party lasted to reach his state of inebriation. And furthermore, another scenario is possible too, one in which Sam wasn't at the party, but is Susan's former boy friend whom she ditched not long ago and who has learned about the party and has had to conclude that he has not been invited. He has to get through the evening somehow and feels there aren't many option apart from getting drunk. If this was the situation, then Sam may well have been drunk already before the party started. Or, Sam being a slow drinker with a lot of alcohol tolerance, he may have got truly drunk not until well after the parts has ended.

An interesting case, in the context of the present discussion is (90.f). What is the temporal relation between the event of leaving and the event of taking the keys? When exactly *was it* that she took the keys? Is there any way of answering that second question that can help us to settle the first? All that (90.f) seems to convey by way of information is that once Susan was out of the house she had the keys with her, while they were in the house until she left (assuming that it was the house that she was leaving).

An apparently clear case among the examples in (90) is (90.g). This sentence seems to have just one possible interpretation – that according the event (or better perhaps: the process, or state of affairs) of it raining was going on at the time when Susan left: the rain-process (or state of affairs or event) was going on when Susan left. In other words, the leaving event was temporally included in the raining event. *when*-sentences of this sort, in which the *when*-clause describes an event and the main clause a state or process, may be argued to be special in that their main clause is always understood as presenting a certain ongoing condition that obtained at the time when the *when*-clause event occurred. But when we discuss Partee's second paper we will see that even for such sentences the issue we are discussing is not that simple.

(90.h), the last example of (90), is much like (90.g). Here too the main clause can be seen as describing an ongoing condition, which held throughout the duration of the *when*-clause event. In this case the condition consists in that an event of a certain kind, that of the speaker turning off the stove, did not occur over a certain period of time – in this case a period of time that includes the *when*-clause event of leaving the house. This observation applies even more straightforwardly to the direct *when*-clause variant of Partee's original example, viz. 'When we were closing up the house, I didn't turn off the stove.', whose main clause just denies the occurrence of a turning-off-the-stove event over a certain period of time. (90.h) is like this last sentence in that the verb *forget* with infinitival *to*-complement entails that an eventuality of the type described by that complement didn't take place (or didn't occur over a period of time indicated by the sentence of which *forget* is the main verb. In the case at hand: the period of time that it took to leave for the given holiday trip.)

(N.B. The aspectual side of negation is intriguing in that negated simple past tense event verb clauses behave differently from their past progressive variants. To see this, consider once more our very first pair of examples about Alan and his wife, now extended with the negations of their second sentences.

(91) a. Alan opened his eyes and saw his wife. She smiled.

b. Alan opened his eyes and saw his wife. She was smiling.

c. Alan opened his eyes and saw his wife. She didn't smile.

d. Alan opened his eyes and saw his wife. She wasn't smiling.

(91.c) is like (91.a) in that in both cases the second sentence is about what happens (or doesn't happen) in response to Alan opening his eyes. And (91.d) similarly resembles (91.b): In either of them the second sentence describes a condition that is already obtaining at the time when Alan opens his eyes. We can explain the semantics of the negated sentences in (91.c) and (91.d) by assuming that the negation has scope over the tense and, in the case of (91.d), the progressive.)

The next set of examples is to show that tenses can also resemble pronouns in that they can behave as *quantificationally bound variables*.

(92) a. Every boy in my class fancies a girl who cannot stand him.

b. Every student talked to a student in front of him.

c. Every classroom in this building has a beamer mounted on its ceiling.

d. If one of the arrows hits the target, it's mine.

e. When you are eating Chinese food, you're always hungry an hour later.

f. Whenever Susan comes in, John immediately leaves.

In (92.a,b) the pronoun *him* is bound by the quantifying phrases *every boy in my class* and *every student*, respectively, and in (92.c) the pronoun *its* stands in such a relation to the subject phrase *every classroom.* What it means to say that in each of these cases the pronoun is bound by the subject phrase is easy to explain to someone who knows something about predicate logic, but not so easy to someone who doesn't. (For the one who knows about predicate logic it will suffice to observe that, for instance, (93) is the symbolization of (92.b) and that the constituent in this formula that corresponds to the pronoun *him* is the penultimate occurrence of the variable x, which is bound by the quantifier $\forall$ that corresponds to the determiner *every* of the grammatical subject of (92.b).

(93) $(\forall x)(student(x) \rightarrow (\exists y)(student(y) \,\&\, infrontof(y,x) \,\&\, talkedto(x,y)))$

For those unfamiliar with predicate logic the following hint will have to do: *him* in (92.b) acts as a variable bound by the quantifying noun phrase *every student* in the sense that when you want to verify whether (92.b) is true you will have to determine for each of the students in questions whether there was another student in front of *that* student that the first student talked to: the person referred to by*him* will for each of these verifications, refer to the particular student that that verification applies to.)

(92.e,f) demonstrate that something like quantificational binding is also possible in the temporal domain. The matter is perhaps clearest for (92.f) in which the conjunction *whenever* functions as a quantifier over times much like *every* functions as a quantifier over 'individuals' (where an 'individual is anything that can be described with the help of a count noun). More precisely, the entire *whenever*-clause plays the sam kind of role as the subject noun phrases *every boy in my class*, *every student*, *every classroom* do in the sentences (92.a-c). Concentrating on (92.f) and (92.b): Just as *every student* in (92.b) expresses a quantification to the effect that everything that satisfies the noun of the phrase (i.e. *student*) satisfies the specification given by the verb phrase (*talked to a student in front of him*), so does the 'quantificational adverbial' *whenever you are eating Chinese food* expresses that each time at which Susan walks in is one such that at a time immediately after it John leaves.

(92.e), in which 'all'-quantification is expressed by the 'quantificational' adverb *always*, illustrates the same point, although the details of its grammatical construction are somewhat different. Exactly how adverbs like *always* work has been a point of ongoing debate between linguists since the seventies. But the net effect of *always* in a sentence like (92.e) is intuitively clear: In every situation in which someone eats Chines food they are hungry an hour after having done so. Generalizing: When *always* occurs in the main claus of a compound sentence consisting of a *when*-clause and a main clause, then the meaning of the compound sentence is that every situation of the kind described by the *when*-clause satisfies the specification given by the main clause.

While there can be no question that (92.e) and (92.f) do have these quantificational meanings, and that their analysis must involve 'bound variables over times' in some way, we should nevertheless pause at this point and ask ourselves whether it is really their *tenses* that act as these variables or that are responsible for bringing them into the picture. In fact, how the different syntactic and morphological elements of sentences like (92.e) and (92.f) conspire to produce the quantificational meanings they have is a puzzle on all pieces of which there isn't complete agreement even today. So Partee's suggestion that such examples show tenses to be capable of playing the part of quantificationally bound variables should be treated with caution. But note well that the arguments are further ammunition against the principle that all temporal relations expressed in natural language should be analyzable in terms of operators from Tense Logic: precisely the doubt that it is the tenses that are responsible for the universal temporal quantifications expressed by these sentences, the argument that all such temporal relations must be analyzable in terms of *tense* operators, is weakened further.

A key position among the examples in (92) is held by (92.c). This sentence is an example of a type of construction that became a central focus of linguistic debate in the second half of the seventies and is usually referred to as the 'donkey sentence' phenomenon, or the phenomenon of 'donkey pronouns' (after sentences like 'If a farmer owns a donkey, he beats it', which Geach (*Reference and Generality*, 1962) had discovered in medieval philosophy and of which he realized the importance for modern logical theory).
If one of the arrows hits the target, it's mine. (92.c) seems to be saying that for any of the arrows that the phrase *the arrows* is referring to here it is the case that if it hits the target it is the speaker's. So here too we have an apparent effect of universal quantification, quantification over arrows. What is curious about such sentences, and what has been responsible for the endless debates to which they have led, is that indefinite noun phrases, which normally express *existential quantification*, contribute in these sentence what looks like a universal quantifier. What seem needed here is an explanation of how this is possible and where it is that indefinites make what seem to be universal and where what seem to make existential contributions.

The reason why the donkey sentence phenomenon is important in the present discussion is that sentences like those in (92.e,f), which as we saw express universal quantification over times, arguably involve something like indefiniteness as well. Consider the slight variant (94) of (92.f).

(94) When Susan is in town, John hides in his cabin.

This sentence is also naturally interpreted as involving universal quantification, viz. as saying something to the effect that on every occasion when Susan is in town John hides in his cabin. But what exactly dos the *when*-clause of this sentence contribute? One plausible suggestion is that what it says is that there is an event (or state) of Susan being in town; and the main clause of the sentence then says that at that time John hides in his cabin. That is, the relationship between the existential information contributed by the *when*-clause and the way in which the main clause exploits that information is strongly reminiscent of the relation between *one of the arrows* and *it* in (92.c). In the second paper of Partee's that we are discussing it is the parallel between sentences like (92.c) on the one hand and sentences like (94) on the other that shoulders a considerable part of the burden of explanation.

Sentences in which, as Partee puts it in this paper, tenses play the role of quantificationally bound variables ranging over time are not restricted to cases of universal quantification. This is shown by the examples in (95).

(95) a. No one could tell what he was being tested for.
- b. Most classrooms in this building have a beamer mounted on their ceiling.
- c. John never answers when I call his home.
- d. John never talks when he is eating.
- e. John never changes his mind when he has made a decision.
- f. John never drives when he has been drinking.
- g. Most of the time, if I write John a letter, he answers within a week.
- h. Mostly, if a man commits perjury, he has to continue committing perjury.

4.5 Partee: Nominal and Temporal Anaphora

This paper revisits the observations that Partee made in her Partee (1973). Its first aim is to find a new formulation of the analogies between tenses and pronouns which makes use of the novel perspective on reference and anaphora that had been made available by the advent of Dynamic Semantics, in the form of File Change Semantics (Heim (1982,1988)) and Discourse Representation Theory (Kamp (1981*b*)). The version of Dynamic Semantics Partee makes use of in this paper is that if DRT.

Before applying the ideas of Dynamic Semantics to the analysis of tense Partee briefly reviews what Dynamic Semantics has to say about pronouns. Since DRT will play a part in this class later on as well, I will cover this same ground, but I will show in addition how DRT's semantic representations – its so-called *Discourse Representation Structures*, or 'DRS's'–, are built from syntactic sentence structures by applying the procedure to a few examples.

The 'Urexample' of DRT, and the example that Partee starts with, is the two sentence discourse in (96) – a 'mini-text' consisting of two sentences the second of which contains pronouns that point back to noun phrases that are part of the first: *he* in the second sentence points back to *Pedro* in the first sentence – that is, *Pedro* is the *anaphoric antecedent* of *he* – and likewise *a donkey* is the anaphoric antecedent of *it*. It is the anaphoric connection between the indefinite noun phrase *a donkey* and the pronoun *it* in the following sentence that had been a source of trouble for the formal approach to the semantics of English and other human languages that dominated semantics one until the advent of Dynamic Semantics, that of Montague Grammar. (N.B. Montague Grammar is still the dominant approach to natural language semantics today; ways have been adopted to get around the to get around the '*a donkey* - *it*' problem that make it possible to retain the general principles on which Montague Grammar is built.)

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

The DRT approach to this and other multi-sentence texts (i.. to bits of text that consist of more than one sentence) is to first build a semantic representation – that is, a DRS – for the first sentence, then use this DRS to build a DRS for the second sentence, which gets merged with the DRS for the first sentence as it gets constructed, then use this DRS in constructing a DRS for the third sentence (in case there is one) and so on.

(97.b) shows the DRS for the first sentence of (96), which is repeated here as (97.a). Once this DRS has been obtained, it can serve to construct the semantic representation for the second sentence. According to the method used here this construction takes the form of extending the DRS for the first sentence with the semantic contributions that the second sentence makes to the discourse. (The method for constructing DRSs for sentences and texts shown below is the one assumed in the work that Partee was relying on; nowadays DRSs are constructed in somewhat different ways, but the differences do not matter for what follows.) The DRS in (97.d) is the extension of the DRS in (97.b) which incorporates the semantic contributed by the second sentence, repeated in (97.c).

(97) a. Pedro owns a donkey.

b.

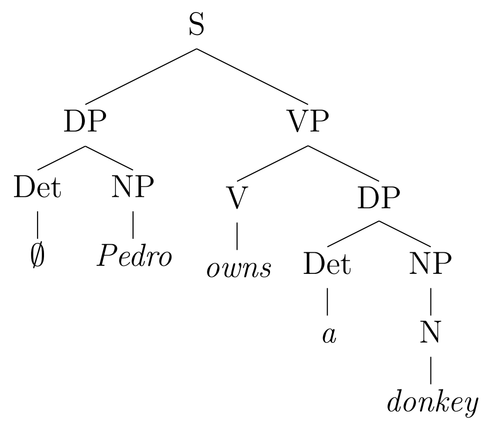

c.

$u \quad v$
$u = Pedro \quad donkey(v)$
$u\ owns\ v$

d. He beats it.

e.

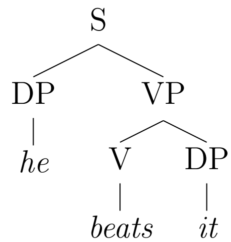

f.

$$\begin{array}{|c|}\hline u \quad v \quad w \quad z \\ \hline u = Pedro \quad donkey(v) \\ u \ owns \ v \\ w = u \quad z = v \\ w \ beats \ z \\ \hline \end{array}$$

According to the DRS construction method Partee assumes DRSs are constructed 'top-down' from *syntactic structures* of the sentences that make up the given discourse. Here DRT, in all forms of formal semantics relies on the insights of Noam Chomsky, who saw early on that the best way in which the grammar of a human language such as English can be described – or their'syntax', to use the term current among linguists to refer to grammar in the specific sense Chomsky had in mind – should be seen as consisting of general principles that enable speakers of the language to put together grammatical sentences from words and that enables their hearers/readers to 'unscramble' the strings of words that reach them as put together according to the principles that the speaker/author has applied in putting together those strings.

Exactly what the principles of a syntax English and other human languages are like has been hotly debated ever since Chomsky put his conception of grammar on the table. But the general conception of a grammar as consisting of principles that enable speakers to build syntactically well-formed sentences from words has remained unchallenged throughout those debates, so that today there is a very large (and, as fas as I can see, dominant) section of the linguistics community that takes this conception for granted.

Part of this conception is that the syntactic structure of a string of words that is built correctly according to the principles of the syntax of the language can be represented in the form of a 'tree', which shows smaller expressions are stepwise integrated into larger ones with the complete sentence as the final outcome, shown at the 'top' of the tree.

On the whole there has been a greater degree of convergence among syntacticians with regard to the structure of such syntactic trees for the different grammatical sentences than there has been with regard to the exact formulation of the principles that underly the construction of these trees. Here we will only show the trees. As a matter of fact, the trees we will be simplified versions of the trees that most syntacticians would now assume for the sentences in question – simplifications that have been chosen to free the method of DRS construction from having to deal with syntactic subtleties that are irrelevant to the DRS construction task at hand.

One important general feature of generative grammars is that their rules make use of *grammatical categories*, such as N(oun), V(erb), V(erb)P(hrase), N(oun)P(hrase), D(eterminer)P(hrase), S(entence), R(elative)C(lause) and so on. Each of these labels defines a syntactically relevant class of expressions – syntactically relevant in that the rules themselves are sensitive to them. For instance, a given grammar may have as one of its rules that any (expression of the type) DP may be combined with a(n expression of type) VP to yield a(n expression of type) S. This rule can only be applied to pairs consisting of a DP and a VP – in that sense it is sensitive to the categories DP and VP. But at the same time rule also plays its part in the definition of categories. For it tells us something about what goes into the category S: Any expression obtained by concatenating a DP and a VP belongs to the category; or, what comes to the same thing, if an expression can be analyzed as consisting of a DP followed by a VP, then it must be an S. And this same dual relationship to the categories of the given grammar – sensitivity to them and contributing to their definition – also applies to the other rules of the grammar. More precisely, it is the system constituted by all the rules of a given generative grammar that defines the different categories, even while the rules in their turn are sensitive to the categories they jointly define. (Those familiar with Recursion Theory will recognize from this description that a generative grammar is a definition by simultaneous recursion of the family of its grammatical categories.)

Some of the grammatical categories that are found in generative grammars were familiar from traditional grammar long before the concept of Generative Grammar made its entry on the scene. And there are quite a few others that originate in particular generative grammars but that have gained a certain autonomy: not only are they found in one generative grammar after another, they have also taken on a life of their own in that other methods have been developed to determine which expressions belong to them (in particular within computational linguistics), so that the question: 'What is the generative grammar for language L?' becomes: 'What set of rules generate this set of (independently determined) categories?'

In fact, it isn't just the grammatical categories that have acquired a certain independence from particular generative grammars If they didn't have such independence by virtue of antedating the Generative Grammar enterprise. The same is true for the structural representations, in tree-like form, of certain forms of sentences that are found in English and other languages. Often factually and even conceptually distinct generative grammars for a human language L will converge on the trees they assign to many of the sentences of L, even if the details of the ways they generate those sentences may be quite different. The syntactic trees on display for th most part belong to this structural common ground.

We return to our example. The tree in (98.a) is, once more, the syntactic tree we assume for (97.a) and which was already displayed in (97.b). In (98.a) this tree is placed into the *condition set* of the DRS we are going to construct for it – the condition set occupies the lower tier of the outer rectangle in (98.a). This outer rectangle will serve as the frame within which the DRS construction which will demarcate the DRS that will result from this construction, viz. the one shown in (97.c). The successive construction steps are shown in (98.b-d). The step leading to (98.b) deals with the highest composition operation that can be applied to the tree in (98.a) – that step which builds the string labelled 'S', viz. the string *Pedro owns a donkey*, from its 'subject' NP part (*Pedro*) and its VP part (*owns a donkey*). It introduces a *discourse referent* u for the subject phrase *Pedro* which is introduced into the upper tier of the DRS (its so-called *universe*) and that also is inserted for the subject noun phrase *Pedro* in the syntactic tree. The role of u is to represent the individual denoted by the subject noun phrase, which in this case, where the subject phrase consists just of the name 'Pedro', is the person named *Pedro* that the speaker of (96) uses 'Pedro' to refer to. This information is coded by adding the condition '$u = Pedro$' to the condition set of the DRS (that is in the same lower tier which also contains the syntactic structure that is being decomposed).

The next construction step, the results of which are shown in (98.c), deals with the one remaining syntactic construction operation that can still be performed on the syntactic tree of (98.b), that which builds the VP string from the verb and the direct object DP *a donkey*. Again a new discourse referent, v, is introduced into the universe of the DRS and substituted for the DP in the syntactic structure – v stands for some donkey but no further specification of it is given; the information is coded by the DRS condition '$donkey(v)$', which is now added to the condition set.

At this point we have reached a structure which cannot be reduced any further and which expresses that the relation expressed by the verb – the relation of ownership – holds between the individuals represented by u and v. We can express that relationship in the simpler and more familiar form '$owns(u, v)$'. Rewriting the relationship on this way turns (98.c) into its more convenient alternative (98.d).

(98) a.

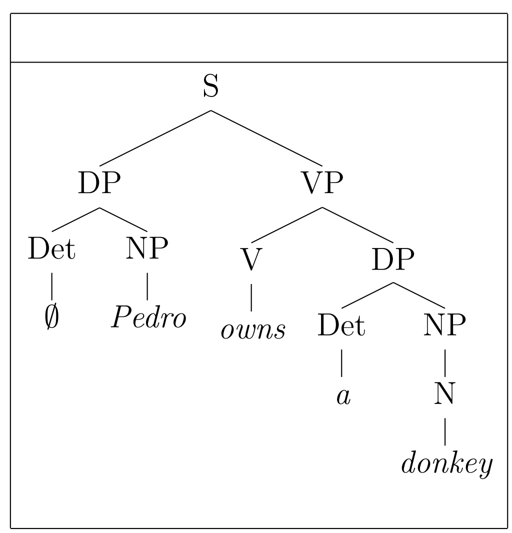

b.

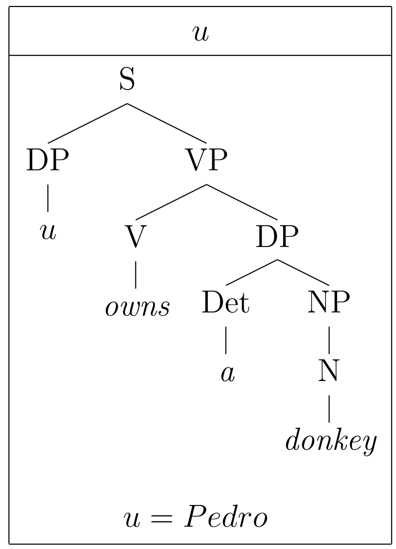

c.

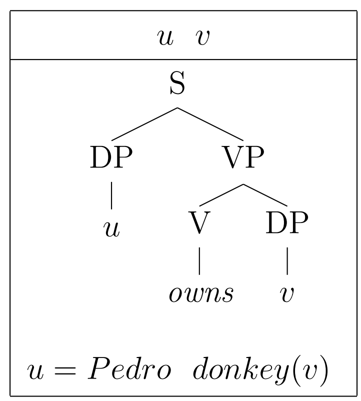

d.

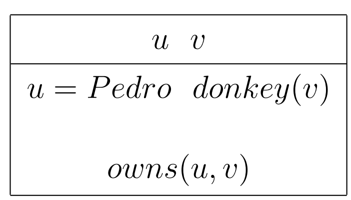

In order to extend the DRS in (98.d) to a DRS for the two sentence discourse in (96) we add the syntactic tree for the second sentence to the condition set of this DRS and then apply reduction operations to this tree that are similar to those that we just applied to the syntactic tree for the first sentence. The result of the tree insertion is shown in (99.a), the result of the first operation in (99.b), that of the second operation in (99.c) and (99.d) simplifies (99.c) in the same way that (98.d) simplifies (98.c).

The one difference between the processing of the syntactic tree in (99.a) and that in (98.a) is that the two DPs of (99.a) are pronouns which have to be interpreted as anaphoric to the DPs *Pedro* and *a donkey* in the first sentence. Part of the operations that involve the pronoun DPs of (99.a) is to encode these anaphoric connections. This is done by setting the discourse referents that are introduced for the pronouns as part of these operations equal to discourse referents that were introduced for their anaphoric antecedents. Thus, if w is the discourse referent introduced for *he*, then the anaphoric connection between *he* and *Pedro* is encoded by the equation '$w = u$', which is added to the condition set of the DRS. Likewise for *it* and *a donkey*.

(99) a.

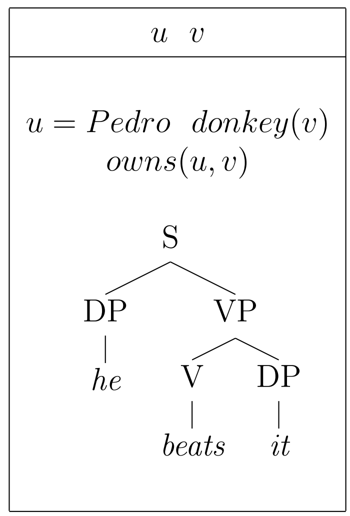

b.

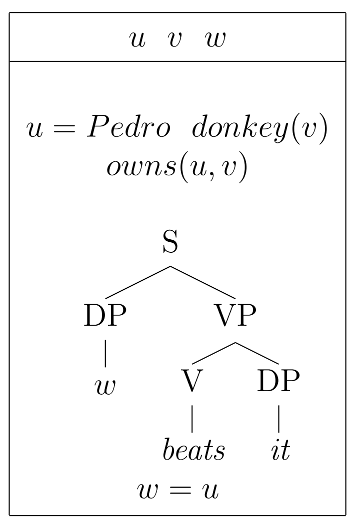

c.

$$\begin{array}{|c|}\hline u \quad v \quad w \quad z \\ \hline \\ u = Pedro \quad donkey(v) \\ owns(u,v) \\ \end{array}$$

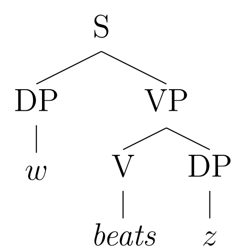

$$\begin{array}{|c|} w = u \quad z = v \\ \hline \end{array}$$

d.

$$\begin{array}{|c|}\hline u \quad v \quad w \quad z \\ \hline \\ u = Pedro \quad donkey(v) \\ owns(u,v) \\ \\ w = u \quad z = v \\ beats(w,z) \\ \hline \end{array}$$

So far so good. But what we need is a method of DRS construction that also can account for the temporal and aspectual dimensions of sentence meaning. For the sentences in (96) these matters are comparatively simple: both are present tense sentences which each assert that a certain state holds at the utterance time. But precisely because tense and aspect are simple in this case, (96) is good starting point for integrating the processing of temporal and aspectual information with the kinds of processing steps that we saw illustrated in (98) and (99).

However, in order to be able to do this we must first enrich the syntactic structures for the sentences in (96) (and likewise for other sentences), so that the structure tells us when and where to deal with the temporal and aspectual information carried by the sentence. For now I will adopt a limited solution of this problem. (We will return to this matter again in our discussion of the paper by Moens and Steedman.) This solution is familiar from generative approaches to syntax in that it places the addition of tense morphology as coming just before the relevant sentence constituent constructed from the verb is combined with the grammatical subject. The node marked 'T' carries a tense 'feature', which is overtly expressed by the tense morpheme. (We distinguish three possible values for the tense feature: past, present and future. All further distinctions, having to do with progressive and perfect, will be dealt with as matters of aspect.) For the first sentence of (96) this proposal yields the tree in (100).

(100)

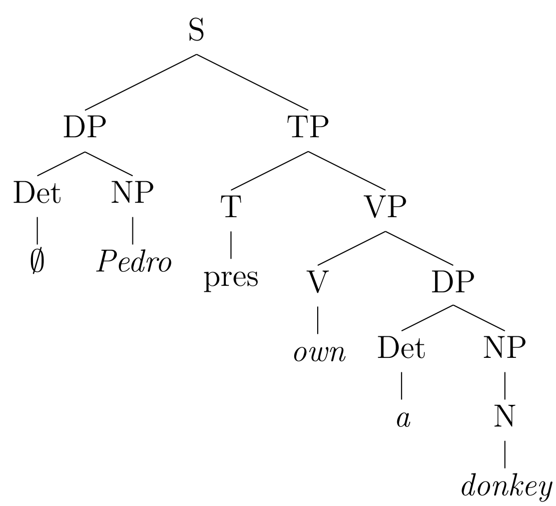

Top down semantic processing of (100) now involves a step, in between those that deal with the subject and direct object arguments, which deals with the information provided by the tense feature value ‘pres’. This step deals with the temporal location – at the utterance time – of the eventuality described by the VP constituent. We assume that ‘pres’ stands for the standard use of the present tense, according to which the eventuality must be a state which temporally includes the utterance time n. We encode this information using the predicate ‘State’, which says of an eventuality that it is a state, and the temporal inclusion relation $\subseteq$. That is, we introduce a new discourse referent ev_1 (where the use of the letters ‘ev’ is to indicate that the entity represented by ev_1 is of the sort ‘eventuality’), introduce this discourse referent into the universe of the DRS and add to the condition set the conditions ‘$State(ev_1)$’ and ‘$n \subseteq ev_1$. In addition we substitute ev_1 for the feature value ‘pres’ under the ‘T’-node, to indicate that this is the eventuality described by the VP. Thus, after the first two processing steps – that involving the subject DP *Pedro* and the tense feature – have been applied to (100), the representation reached is that in (101.a).

(101) a.

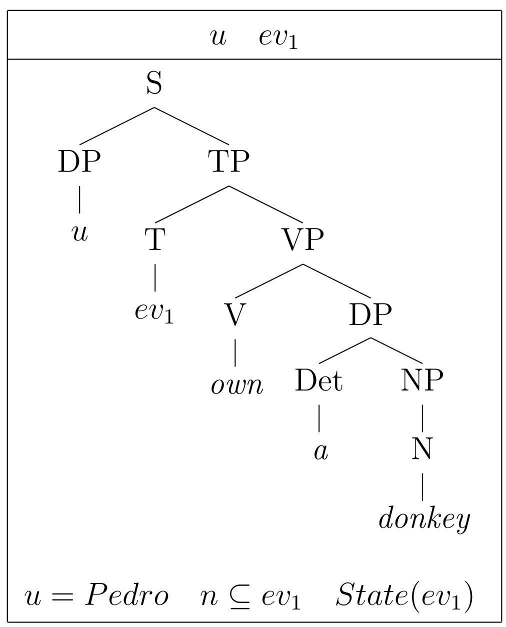

b. 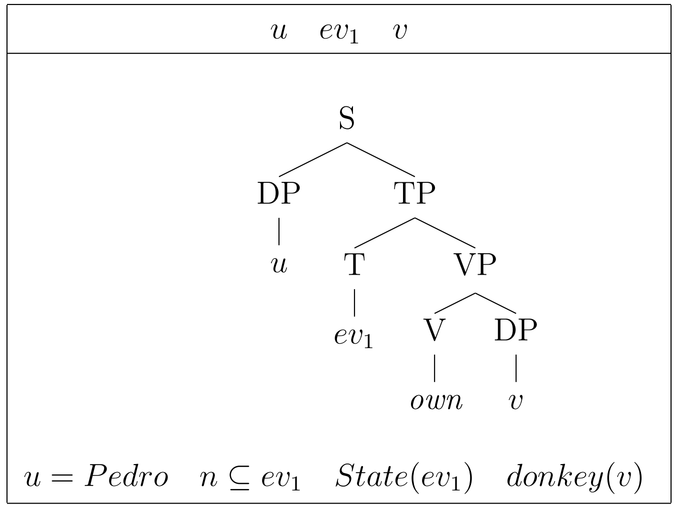

c.
$$\begin{array}{|c|}\hline u \quad ev_1 \quad v \\ \hline \\ \\ u = Pedro \quad n \subseteq ev_1 \quad State(ev_1) \quad donkey(v) \\ ev_1 : own(u, v) \\ \hline \end{array}$$

d.
$$\begin{array}{|c|}\hline u \quad s_1 \quad v \\ \hline \\ \\ u = Pedro \quad n \subseteq s_1 \quad donkey(v) \\ s_1 : own(u, v) \\ \hline \end{array}$$

The next step deals with the direct object DP in the same way as before, with the result shown in (101.b). To simplify this representation we must now treat the verb *own* as a 3-place predicate, with the arguments ev_1, u and v. The simpler notation we use for this predicational relationship is: '$ev_1 : own(u, v)$'. (101.c) gives the result of this simplification. And a further notational simplification is possible as well, which consists of replacing the eventuality discourse referent ev_1 by the discourse referent s_1, in which the symbol 's' indicates that it ranges exclusively over states. This renders the sortal restriction condition involving the predicate 'State' redundant, so this condition can now be dropped. This further simplification is shown in (101.d)

Note that once we reach, after the third step, the level of the lexical verb, we have additional information about the aspectual category of the verb. In the present case, where the verb is the stative verb *own*, this information is consistent with the condition '$State(ev_1)$' that was added to the DRS as part of processing the feature 'pres'; so things are as they should be. Had, on the other hand, the verb been an event verb (such as for example *eat*), then there would have been a clash between this condition and the lexical information that ev_1 must be an event. We will return shortly to this point in connection with the verb *beat* in the second sentence of (96).

In fact, this is the only new complication we encounter when the representation construction for (96) of which (101) shows the first half by processing the second sentence along the same lines. (Otherwise the new representation construction differs from the old one in (96"') the same way as that in (101) differs from the one in (98).) The problem presented by *beats* manifests itself when the construction reaches the end node (or 'leaf') labeled 'V' which anchors the occurrence of the lexical verb *beat*. Let us assume that the lexicon – that part of the grammar which contains all the words of the language, with their relevant properties – this verb is listed as an event verb – a verb that serves to describe events, and not states. Then the representation construction reveals an inconsistency at this point between this lexical specification and the already adopted condition '$State(ev_2)$' for the eventuality described by *beat*. This doesn't mean, however, that the construction aborts at this point. Rather, the presence of '$State(ev_2)$' is an encouragement to 'coërce' the given occurrence of *beat* into a state verb (i.e. to reinterpret it in such a way that it comes to qualify as a state describer). In this case coërcion is possible and it takes the form of reinterpreting *beat* as a *habitual*, that is as a description of a habit or practice to beat the direct object. Habits and practices are kinds of stats that the subject can be in and that can be felicitously asserted to hold at the utterance time.

That coërcion must take place as part of the interpretation of occurrences of event verbs in sentences like the second sentence of (96) has long been realised, as has the fact that the result of such coërcions is typically a habitual or dispositional reading. But exactly what habituals and dispositions are and how the habitual or dispositional reinterpretations of event verbs are related to the meanings of those verbs as 'unadulterated' event describers, are questions that remain without a satisfactory answer. Here I do no more than indicate the problem by adding a subscript $_{hab}$ to the occurrence of *beat* in the final DRS.

The most relevant stages in the DRS construction for the second sentence of (96) are shown in (102). (102.a) gives the starting position for this stretch of semantic processing of (96): the DRS from (101) extended with the new syntactic tree for the second sentence. (102.b) is the result of applying the tense processing rule, (102.c) the representation at the point where the processor is confronted with the conflict between '$State(ev_2)$' and the lexical semantics of *beat*, and (102.d) is the final representation for (96), which results from the resolution of this conflict through 'habitual coërcion'.

(102) a.

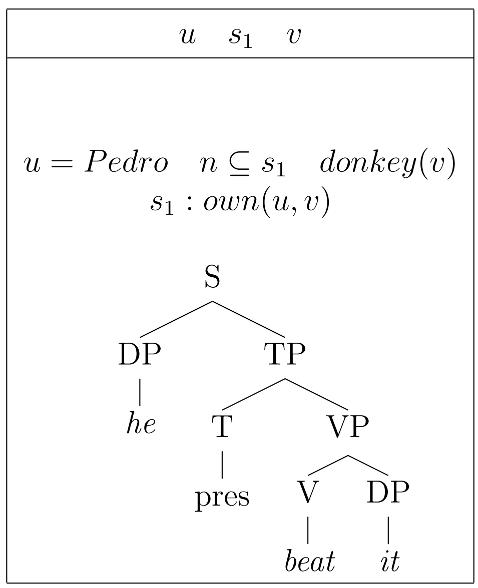

b.

$$u \quad s_1 \quad v \quad w \quad ev_2$$

$$u = Pedro \quad n \subseteq s_1 \quad donkey(v)$$
$$s_1 : own(u, v)$$

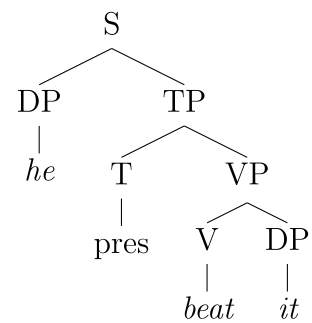

$$w = u \quad State(ev_2) \quad n \subseteq ev_2$$

c.

$$u \quad s_1 \quad v \quad w \quad ev_2 \quad z$$

$$u = Pedro \quad n \subseteq s_1 \quad donkey(v)$$
$$s_1 : own(u, v)$$

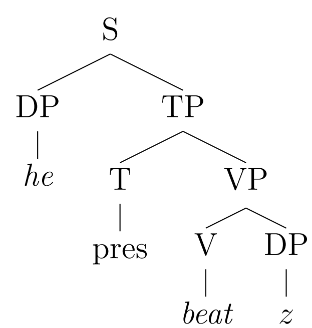

$$w = u \quad State(ev_2) \quad n \subseteq ev_2 \quad z = v$$

d.

$$\begin{array}{|c|}
\hline
u \quad s_1 \quad v \quad w \quad s_2 \quad z \\
\hline
\\
u = Pedro \quad n \subseteq s_1 \quad donkey(v) \\
s_1 : own(u, v) \\
\\
w = u \quad n \subseteq s_2 \quad z = v \\
s_2 : beat_{hab}(w, z) \\
\hline
\end{array}$$

N.B. The processing of tense in this example is exceptionally and misleadingly simple. There are two occurrences of the simple present tense in this example, each of which has been interpreted as claiming that a certain state holds at the utterance time n (just as two occurrences of the pronoun I in an utterance by a speaker S will be interpreted as referring both to S). The net effect of this is that the two sentences are interpreted as speaking about the same time, but that result is obtained in this case by interpreting their tenses independently from each other. As far as temporal reference is concerned (96), is thus *not* an example of the analogy between tenses and (3rd person) pronouns that Partee is interested in. In fact, the present tense is, as just noted, more like the 1st person than like the 3rd person pronoun. The close analogy with 3rd person pronouns that Partee is after in both her 1973 and her 1984 paper is rather with past tenses (and to some extent also with the future tense). We will return to this when showing the representation construction of some of the later examples from 'Nominal and Temporal Anaphora'.

Many logical operations – conditionals, negation, quantification, disjunction prominently among them – give rise to irreducible complex conditions in the representing DRSs. The one example we consider here is that of an *if.., then..*-conditional. Conditionals are represented in DRT by DRS conditions of the form '$K_1 \Rightarrow K_2$', where K_1 and K_2 are DRSs. An example is the DRS in (103.c) for the English conditional in (103.a). (103.c) is the DRS given for this conditional on Partee's paper, but in the somewhat different format for such DRSs that became more widely used not long after her paper was published. (103.b) is a syntactic tree for that yields this representation given the right construction operation for turning the English *if .., then ..* into an $\Rightarrow$ condition.

(103) a. If Pedro owns a donkey, he beats it.

b.

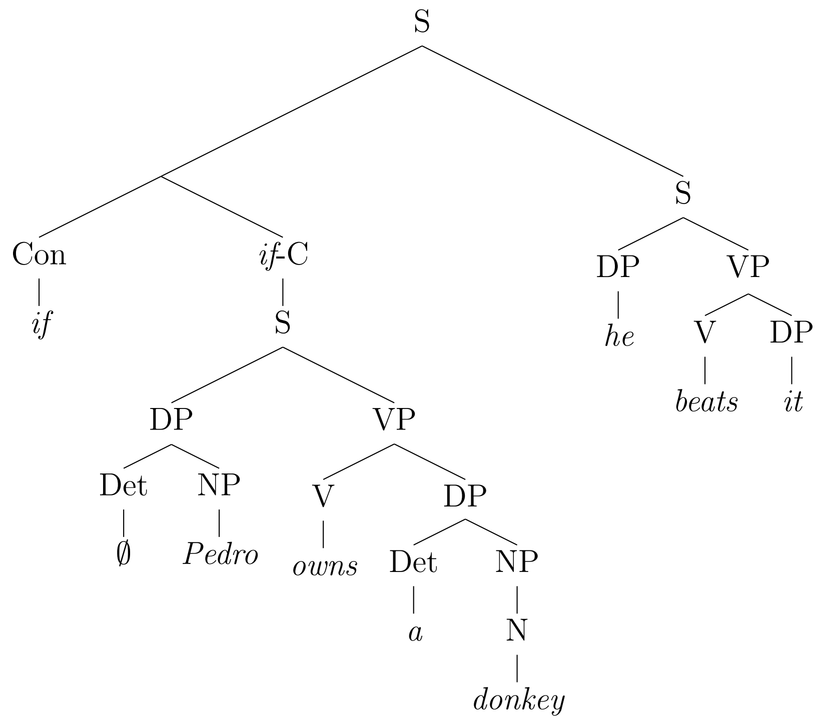

c.

$$\boxed{\begin{array}{|c|}\hline u \quad v \\ \hline u = \textit{Pedro} \quad \textit{donkey}(v) \\ u \ \textit{owns}\ v \\ \hline\end{array} \Rightarrow \begin{array}{|c|}\hline w \quad z \\ \hline u = \quad w = u \quad z = v \\ w \ \textit{beats}\ z \\ \hline\end{array}}$$

d.

$$\begin{array}{|c|}\hline u \\ \hline u = \textit{Pedro} \\ \begin{array}{|c|}\hline v \\ \hline \textit{donkey}(v) \\ u \ \textit{owns}\ v \\ \hline\end{array} \Rightarrow \begin{array}{|c|}\hline w \quad z \\ \hline u = \quad w = u \quad z = v \\ w \ \textit{beats}\ z \\ \hline\end{array} \\ \hline\end{array}$$

But in this case too we will eventually want to know more about the temporal structure of *if*- and main clause. So what we want for a start is a syntactic structure in which the tenses are separate constituents, as in (104).

(104)

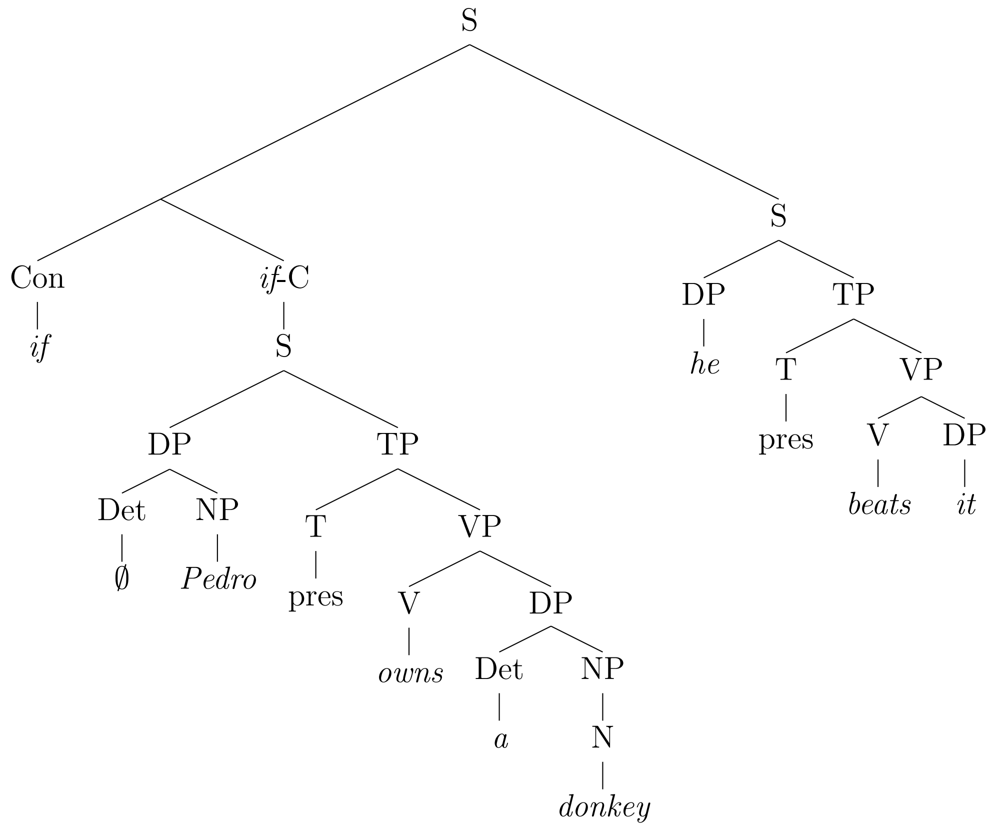

The first step in the DRS construction for (104) separates the *if*-clause adjunct from the main clause, establishing a ⇒-condition with as its left (or *antecedent*) DRS one that contains the syntactic structure of the *if*-clause S and as right (or *consequent*) DRS one containing the syntactic structure of the S of the main clause, see (105.a). (105.a) gives the complete DRS, which is like (103.c), except that it also represents the described states.

(105) a.

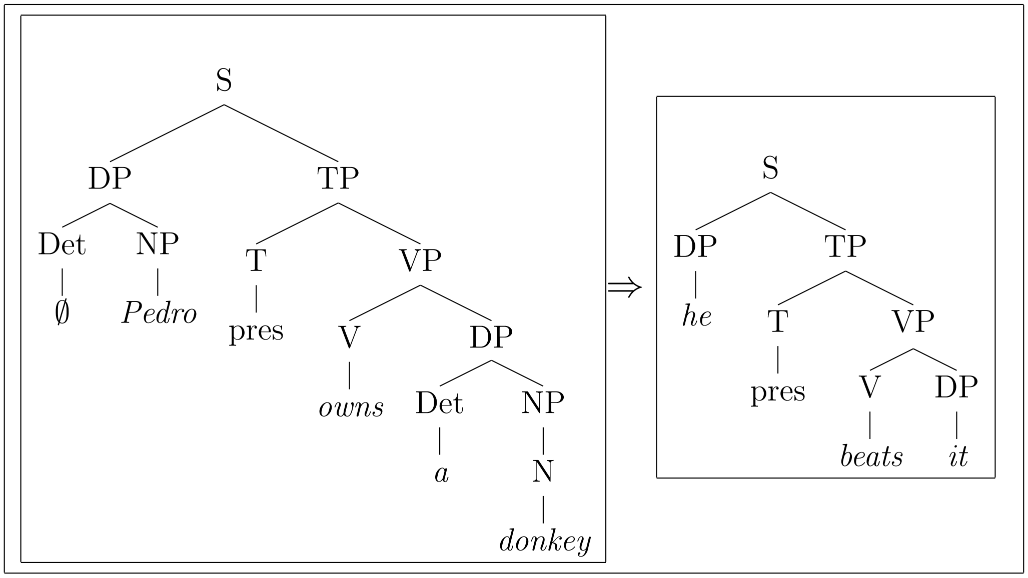

b.

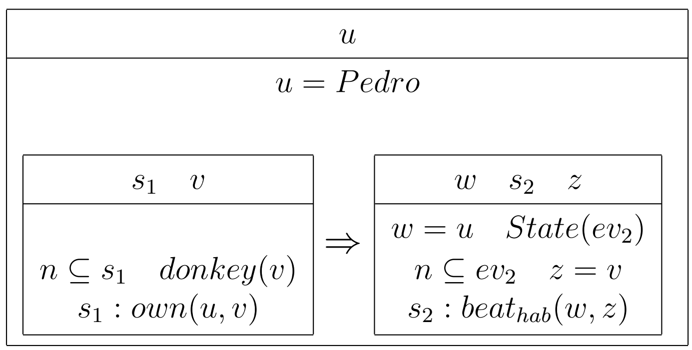

Quantifying noun phrases, such as the subject DP of (106.a), also give rise to complex DRS conditions. Universal quantifiers like *every farmer* can be represented by mans of ⇒-conditions – this is because the semantics for such conditions makes the discourse referents in the universe of the antecedent DRS play a universally quantified role (whereas those in the universe of the consequent DRS play an existential role) – or by a so-called *duplex condition*, which is the option shown below. Duplex conditions are introduced by the processing steps that deal with the DPs containing the quantifying determiners (in the way that the DP *every farmer* contains the determiner *every*). (106.b) gives the syntactic tree for (106.a) laid out according to the same principles as those in (96.4), (101.a) and (104.b). (106.c-e) show relevant construction stages.

(106) a. Every farmer who owns a donkey beats it.

b.

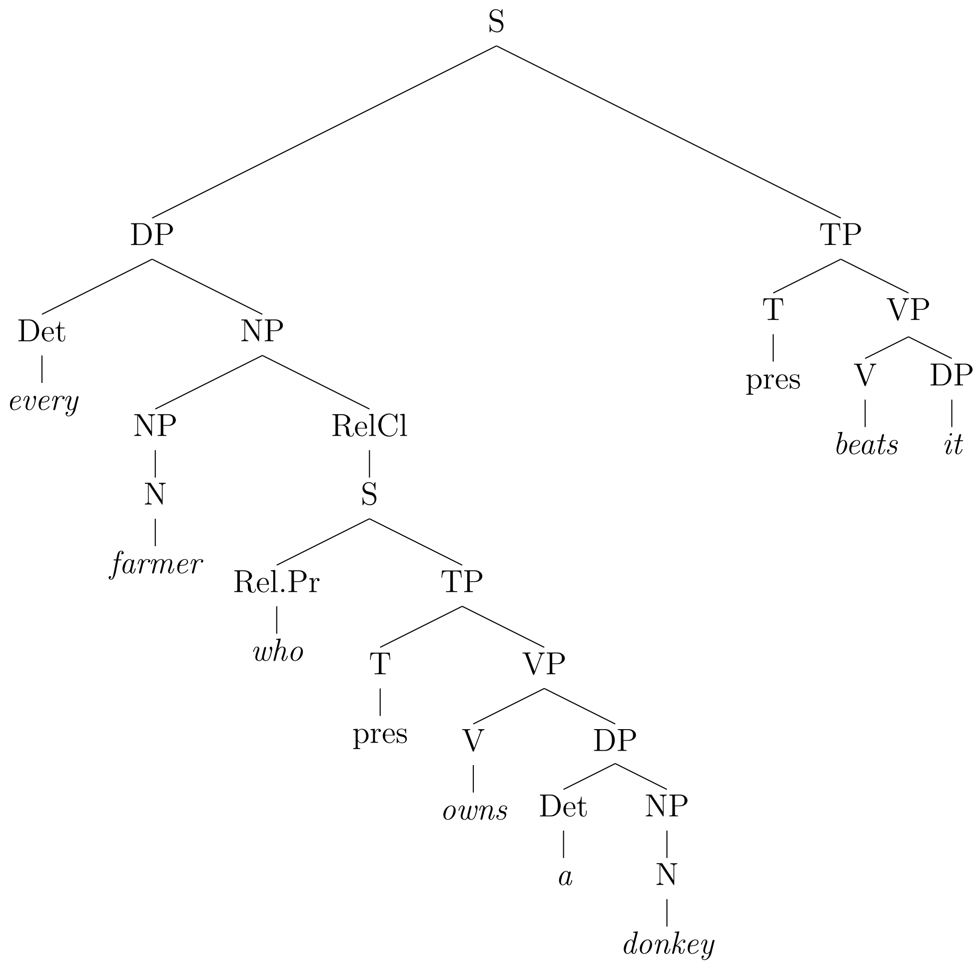

c.

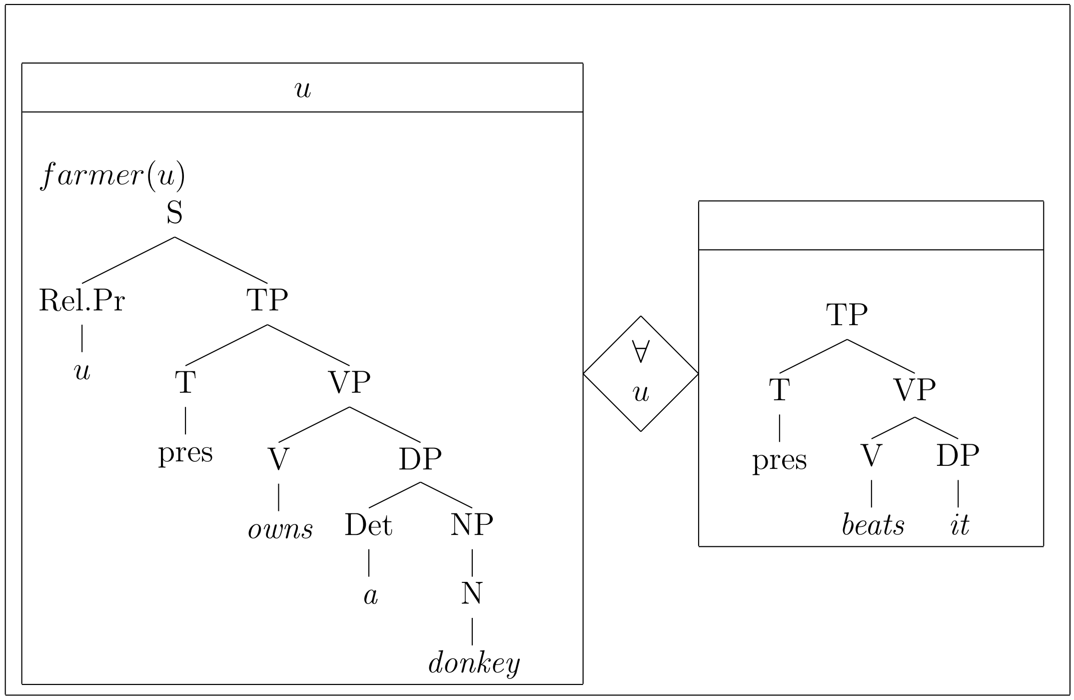
u
farmer(u)
S
Rel.Pr
u
TP
T
pres
VP
V
owns
DP
Det
a
NP
N
donkey
∀
u
TP
T
pres
VP
V
beats
DP
it

d.

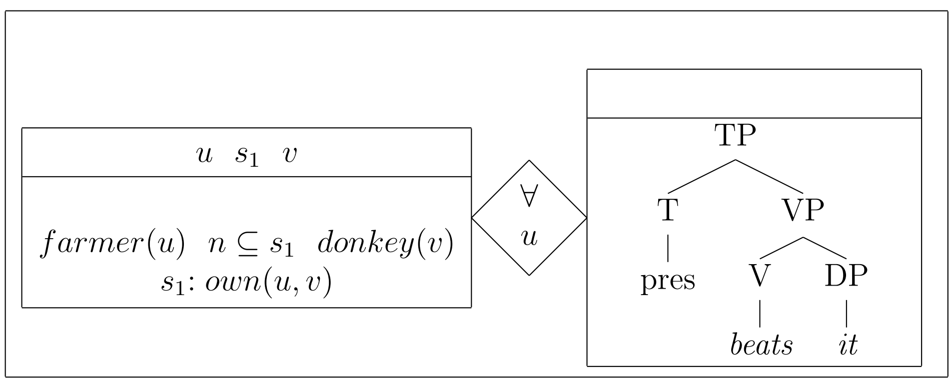
$u \quad s_1 \quad v$
$farmer(u) \quad n \subseteq s_1 \quad donkey(v)$
s_1: $own(u, v)$
∀
u
TP
T
pres
VP
V
beats
DP
it

e.

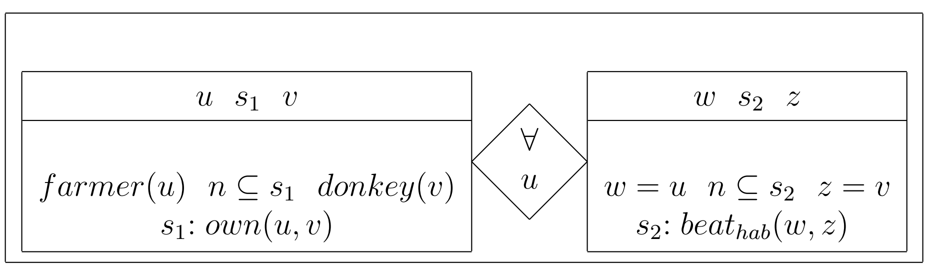
$u \quad s_1 \quad v$
$farmer(u) \quad n \subseteq s_1 \quad donkey(v)$
s_1: $own(u, v)$
∀
u
$w \quad s_2 \quad z$
$w = u \quad n \subseteq s_2 \quad z = v$
s_2: $beat_{hab}(w, z)$

This much by way of preparation to what Partee (1984) has to say specifically about tense. Her focus is on past tense narrative, a type of prose in which the sentences of which the narrative is made up are all in the past tense and where event sentences typically have the effect of ‘pushing the story forward’. (It is important in the longer run, and also in connection with the following papers, to keep in mind that past tense narratives form just one among many different text types. Different text types may involve different principles for computing temporal relations on the basis of the syntactic forms of the sentences involved.)

The first such example that the paper discusses in detail is that in (107.a). This example consists of three event clauses, a stative clause, two event clauses and two stative clauses. All the event clauses advance the time of the story whereas the stative clauses do not; they described conditions obtaining at the time of the last mentioned event.

Partee gives the DRS (107.b) as semantic representation for (107.a). Her discussion of subsequent examples reveal more of the general principles that underly the construction of this DRS; but it will serve our purposes better if we bring these considerations to bear already at this point.

(107) a. John got up, went to the window, and raised the blind.
e_1 e_2 e_3

It was light out. He pulled the blind down and went back to bed.
s_1 e_4 e_5

He wasn't ready to face the day. He was too depressed.
s_2 s_3

b.

e_1 e_2 e_3 e_4 e_5 s_1 s_2 s_3 r_s
$e_1 < e_2 < e_3 < e_4 < e_5$
e_1: John get up
e_2: John go up to the window
e_3: John raise the blind
e_4: John pull the blind down
e_5: John go back to bed
s_1 O e_3
s_2 O e_5
s_3 O e_5
s_1: It be light out
s_2: John not be ready to face the day
s_3: John be too depressed

If we want to relate the construction of the semantic representation for (107.a) and the following examples to the constructions above, the first question we need to settle is that of syntactic form. Here we will make things as easy for ourselves as is compatible with our present purpose. We make the first clause of (107.a) – a simple past tense sentence consisting of a subject DP and an intransitive verb – be our paradigm. Its syntactic structure will be assumed to be that in (108).

(108)

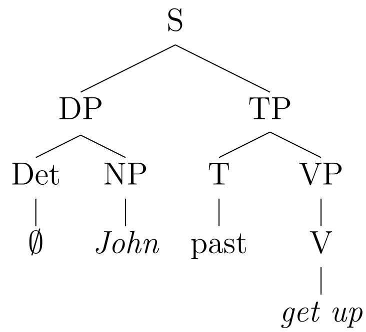

We will assume such syntactic structures for all the clauses in (107.a), treating what follows the subject DP as an ‘intransitive verb’ that is not analyzed further into smaller components. In the present context this is wholly unproblematic but for one case: the subject DP *it* of *It was light out* is what is called a ‘dummy subject’; it is present only for syntactic reasons, i.e. to provide a grammatical subject where there is no semantic argument to fill that position. The construction step dealing with such dummy subjects does not involve the introduction of a new discourse referent to take the place of the subject DP. In fact all that happens at this point is that the DP is excised from the syntactic structure, as an indication that it has been dealt with. Thus, if we were to construct a representation just for this sentence, then the first two representation stages would be those in (109.b,c). The remaining steps will then convert (109.c) into a state specification of the form ‘s : *lightout*’, located as holding at some time in the past of n. (Details of the processing of the past tense follow below; see in particular (112.)

(109) a. It was light out.

b.

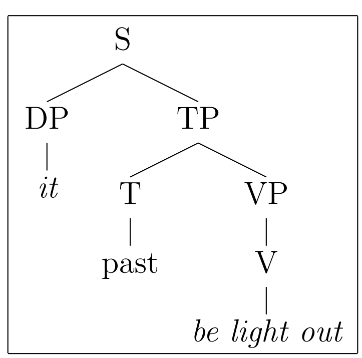

c. 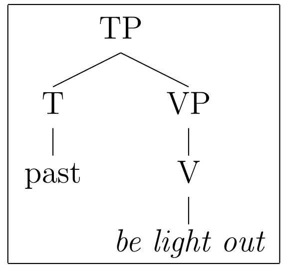

There is one other matter that we will ignore in what follows. Some of the clauses in (107.a) are complete sentences standing all by themselves ((109.a) is an example), while others are conjuncts of longer sentences, separated from the other clause or clauses of the sentence by *and* or by a comma. We will ignore these distinctions and treat all clauses as separate sentences that are separated from their predecessors and/or successors by full stops. That is of course quite artificial, since what w get when we replace all the *and*'s and commas in (107.a) by full stops is only marginally acceptable English prose. Exactly why that is so – why (107.a) is so much better than what we get when we leave the *and*'s and commas out – still defies linguistic theory. But on the other hand it appears that sentence conjuncts require the same kind of semantic processing as clauses standing on their own (at least when the clauses in question are simple past tense clauses like these). So the simplification does not affect what should be said about the semantics of (107.a). Lt us assume, moreover, that each of the non-initial sentence conjuncts in (107.a) are replaced by full sentences in which the subject DP is *he*. So the second clause is now *He went to the window* and so on.

With these syntactic preliminaries out of the way we can turn to the substance of Partee's proposal. This proposal, for much of which Partee gives credit to Hinrichs, is based on the general principle that every past tense event described by a past tense event sentence comes with a *reference point* which is situated directly after the end of the event and which serves to locate the eventuality that the text or discourse mentions next. It is in this way that Hinrichs and Partee account for the temporal progression that event sentences typically produce in past tense narration.

On the one hand, then, event clauses in bits of past tense narrative like (107.a) *introduce* reference points (together with the events they introduce). But on the other hand there also must be a way in which past tense event clauses must *look for* such reference points for their temporal location. For it is only by virtue of the next one of a pair of event clauses 'finding' the reference point introduced by the first and locating its event at that point that narrative progression gets established.

It is here that Hinrichs and Partee rely on the two-dimensional theory of tense proposed by Reichenbach. Reichenbach classifies the simple past as a tense that situates its reference time in the past from the utterance time and locates the event time as coincident with the reference time. Moreover, determining this reference time – this is also quite clear from the examples Reichenbach discussed – is part of the task of interpreting tensed sentences. It is still unclear what the rules for determining reference times are in general, as they are selected by and apply to different text and discourse types and different rhetorical tropes within the same type. Hinrichs and Partee present no more than a very partial answer to this general question. But even if the text and discourse passage to which their rules do apply are only a handful among many, their implicit position that these cases are in some important sense paradigmatic is arguably right.

For texts like (107.a) their search principle is simple enough: go for the reference point introduced by the last (simple past) event sentence.

In addition there is a general difference – it is independent of any issues of text- or discourse type – about how the eventuality described by the clause is temporally related to the reference time. Here events and states differ. A state temporally includes the reference time, an event is temporally included within the reference time.

The upshot of it all is that quite a number of operations need to be executed once the tense operation has become available. In particular, when the tense is 'past', then (i) a new eventuality discourse referent must be introduced to represent the described eventuality; (ii) a past reference point must be identified; (iii) the right temporal relation must be established between described eventuality and reference point. One of the questions that an explicit formulation of the interpretation of such must decide one way or another is at which point of the DRS construction these different operations are to be carried out. We adopt the following stipulations: (i) and (ii) are performed at the point of dealing with the information provided by T (i.e. the 'decomposition' of

TP into T and VP); and the temporal relation between the new eventuality discourse referent ev and the reference point discourse referent rp is put in place provisionally via the 'ambiguous' condition '$(Event(ev)\ \&\ ev \subseteq rp) \vee \cdot\ (State(ev)\ \&\ rp \subseteq ev)$'. The 'alternative reading' sign '$\dot{\vee}$' indicates that a choice has to be made between the alternatives flanking it before the representation can be considered complete. The choice between the two alternatives can be made only at the point when enough is known about the aspectual properties of the eventuality description to know whether what is described is an event or a state. For the clauses that make up (107.a) this will be the case only when there is access to the lexical properties of the verb. Thus disambiguating '$(Event(ev)\ \&\ ev \subseteq rp) \dot{\vee} (State(ev)\ \&\ rp \subseteq ev)$' will be the final step.

The last notational stipulation we need before it is possible to highlight the construction of the DRS for (107.a) by displaying a selection of stages has to do with how we encode the result of identifying the reference point. We will do this in the form of an *assignment condition* of the form '$rp := r$, where r is the time discourse referent that has been identified as reference time.

As we have seen, according to Hinrichs and Partee a further operation is needed in the case of event clauses: that of introducing a reference point (directly) after the event introduced by the clause. This operation can be executed once it is known that the clause is event- rather than state-describing; and in the cases at hand that is known only at the end. So this operation can and will only be performed simultaneously with the resolution of the temporal relation between ev and rp.

Partee only considers the case where the reference point needed to interpret the past tense of clause S_i is identified with that of the last event sentence preceding it. But that of course cannot work for the first event sentence S_1, as it is the first clause of the discourse and thus not preceded by any other, by which an antecedent reference point could have been established. In this case, Partee says, a reference point simply has to be assumed (or 'accommodated' as one often says about such assumptions to which the interpreter is forced by the canons of interpretation). In the construction below we assume that the reference time r_0 is introduced into the DRS at the point of the first processing of a past tense.

At last everything is in place for the DRS construction for (107.a). (110.a) gives the starting position for the DRS construction for the first clause, (110.b) the result of applying the operations triggered by the past tense and (110.b) the result of recognizing the eventuality described by the clause to be an event and (110.d) the simplification that results from using a discourse referent restricted to events (rather than to events and states). Moreover, since at this point the current reference point representing term rp has done its current work, we eliminate it from the representation – by dropping the condition $rp := r_0$ and replacing all other occurrences of rp by r_0. In this way rp can be used afresh in dealing with the next clause.

(110) a.

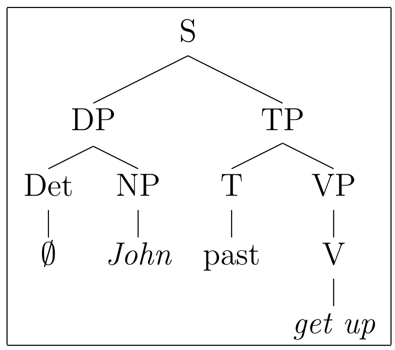

b.

$u \quad r_0 \quad ev_1$

$u = John \quad r_0 < n \quad rp := r_0$

$(Event(ev_1)\ \&\ ev_1 \subseteq rp) \vee (State(ev_1)\ \&\ rp \subseteq ev_1)$

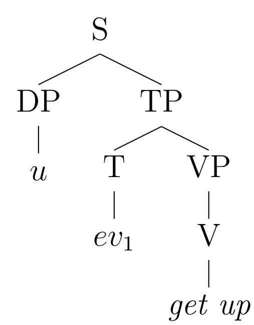

c.

$$\begin{array}{|c|}\hline u \quad r_0 \quad ev_1 \quad r_1 \\ \hline u = John \quad rp < r_1 < n \quad rp := r_0 \\ \\ Event(ev_1) \quad ev_1 \subseteq rp \\ \hline \end{array}$$

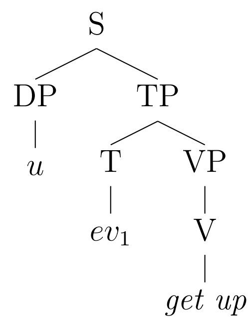

d.

$$\begin{array}{|c|}\hline u \quad r_0 \quad e_1 \quad r_1 \\ \hline u = John \quad r_0 < r_1 < n \quad e_1 \subseteq r_0 \\ \\ e_1 : get - up(u) \\ \hline \end{array}$$

The processing of the next sentence, *He went to the window*, is almost exactly like that of the first sentence. The only difference is that this time the reference point identifying mechanism doesn't have to accommodate a time but finds the suitable candidate r_1. The same goes for the third sentence. (111.a) shows the representation just after the tense of the second sentence has been processed, (111.b) after the second sentence has been processed fully and (111.c) the representation after full processing of the third sentence, *He raised the blind*.

(111) a.

$$\begin{array}{|c|}\hline u \quad r_0 \quad e_1 \quad r_1 \quad u_1 \quad ev_2 \\ \hline u = John \quad r_0 < r_1 < n \quad e_1 \subseteq r_0 \\ \\ e_1\text{: } get - up(u) \\ \\ u_1 = u \quad rp := r_1 \\ \\ (Event(ev_2)\ \&\ ev_2 \subseteq rp) \vee (State(ev_2)\ \&\ rp \subseteq ev_2) \\ \end{array}$$

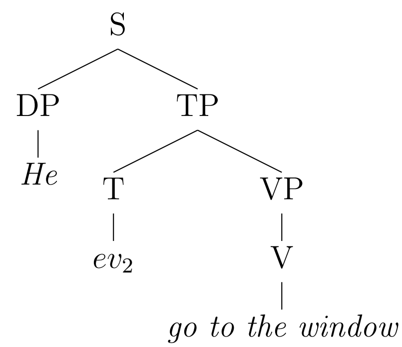

b.

$$\begin{array}{|c|}\hline u \quad r_0 \quad e_1 \quad r_1 \quad u_1 \quad e_2 \quad r_2 \\ \hline r_0 < r_1 < r_2 < n \\ \\ u = John \quad e_1 \subseteq r_0 \\ e_1\text{: } get - up(u) \\ \\ u_1 = u \quad e_2 \subseteq r_1 \\ e_2\text{: } go - to - the - window(u_1) \\ \hline \end{array}$$

c.

$$
\begin{array}{|c|}
\hline
u \quad r_0 \quad e_1 \quad r_1 \quad u_1 \quad e_2 \quad r_2 \quad u_2 \quad e_3 \quad r_3 \\
\hline
r_0 < r_1 < r_2 < r_3 < n \\
\\
u = John \quad e_1 \subseteq r_0 \\
e_1\text{: } get - up(u) \\
\\
u_1 = u \quad e_2 \subseteq r_1 \\
e_2\text{: } go - to - the - window(u_1) \\
\\
u_2 = u \quad e_3 \subseteq r_2 \\
e_3\text{: } raise - the - blind(u_2) \\
\hline
\end{array}
$$

(111.c) is the context DRS for the fourth sentence, *It was light out*. We have already seen, in a way of constructing a semantic representation for this sentence *in vitro*. When we combine the non-temporal aspects of that construction with the temporal processing that has just been used to obtain (111.c), then the result we obtain is that in (112.c). The crucial difference with the processing of the past tense event sentences that we have just gone through arises at the point when the process gains access to the lexical properties of the copular verb *be* and recognizes it as a state describer. This means that the eventuality ev_4 that was introduced in processing the past tense of *be* is a state. This disambiguates the temporal location relation to inclusion of the reference point within the state. (112.a) and (112.b) show the representation just before and after the lexical information is processed and (112.c) is the by now familiar simplification of (112.b).

(112) a.

$$
\begin{array}{|c|}
\hline
u \;\; r_0 \;\; e_1 \;\; r_1 \;\; u_1 \;\; e_2 \;\; r_2 \;\; u_2 \;\; e_3 \;\; r_3 \;\; ev_4 \\
\hline
r_0 < r_1 < r_2 < r_3 < n \\
\\
u = John \quad e_1 \subseteq r_0 \\
e_1\text{: } get - up(u) \\
\\
u_1 = u \quad e_2 \subseteq r_1 \\
e_2\text{: } go - to - the - window(u_1) \\
\\
u_2 = u \quad e_3 \subseteq r_2 \\
e_3\text{: } raise - the - blind(u_2) \\
\\
rp := r_3 \\
\\
(Event(ev_4) \;\&\; ev_4 \subseteq rp) \lor (State(ev_4) \;\&\; rp \subseteq ev_4) \\
\hline
\end{array}
$$

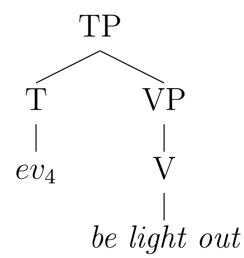

b.

$$\begin{array}{|c|}
\hline
u \quad r_0 \quad e_1 \quad r_1 \quad u_1 \quad e_2 \quad r_2 \quad u_2 \quad e_3 \quad r_3 \quad ev_4 \\
\hline
r_0 < r_1 < r_2 < r_3 < n \\
\\
u = John \quad e_1 \subseteq r_0 \\
e_1\text{: } get - up(u) \\
\\
u_1 = u \quad e_2 \subseteq r_1 \\
e_2\text{: } go - to - the - window(u_1) \\
\\
u_2 = u \quad e_3 \subseteq r_2 \\
e_3\text{: } raise - the - blind(u_2) \\
\\
rp := r_3 \\
\\
State(ev_4) \quad rp \subseteq ev_4 \\
\\
\hline
\end{array}$$

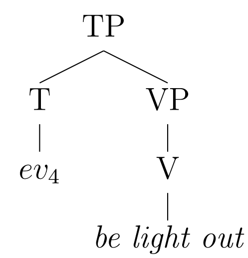

c.

$u \;\; r_0 \;\; e_1 \;\; r_1 \;\; u_1 \;\; e_2 \;\; r_2 \;\; u_2 \;\; e_3 \;\; r_3 \;\; s_4$
$r_0 < r_1 < r_2 < r_3 < n$
$u = John \quad e_1 \subseteq r_0$ e_1: $get - up(u)$
$u_1 = u \quad e_2 \subseteq r_1$ e_2: $go - to - the - window(u_1)$
$u_2 = u \quad e_3 \subseteq r_2$ e_3: $raise - the - blind(u_2)$
$r_3 \subseteq s_4$ s_4: $light - out$

Processing the remaining clauses of (107.a) lads to no new surprises and so I leave completion of the DRS for (107.a), by applying the construction principles we have so far discussed, to you.

The next example from Partee, here given as (113.a), is much like the one we have just looked at in detail (given in full in (107.a)). So I leave it as an exercise to construct DRS (113.b) for (113.a) using the principles that we have just been using in our construction of the DRS for (107.a).

(113) a. Jameson entered the room, shut the door carefully,
e_1 e_2

and switched off the light. It was pitch dark around him,
e_3 s_1

because the Venetian blinds were closed.
s_2

b.

$$
\begin{array}{|c|}
\hline
r_0 \quad e_1 \quad r_1 \quad e_2 \quad r_2 \quad e_3 \quad r_3 \quad s_1 \quad s_2 \quad r_s \\
\\
\hline
e_1 \subseteq r_0 \\
e_1 < r_1 < r_s \\
e_2 \subseteq r_1 \\
e_2 < r_2 < r_s \\
e_3 \subseteq r_2 \\
e_3 < r_3 < r_s \\
r_3 \subseteq s_1 \\
r_3 \subseteq s_2 \\
\\
e_1\text{: John enter the room} \\
e_2 : \ldots \\
\ldots \\
\hline
\end{array}
$$

There is only one respect in which (113.a) differs from (107.a). This is the conjunction *because* which connects the final two stative clauses. It presents the second of the two as an explanation of the claim made by the first. This does not affect the temporal interpretation of either clause – setting *because* aside and processing the two clauses along the lines adopted in the course of our discussions over the DRS construction for (107.a) will locate both states described by these clauses as holding at the reference point introduced by the third event clause; and that seems perfectly fitting given our intuitive understanding of what (113.a) means.

Often, however, such non-temporal aspects of interpretation – which have to do with the *rhetorical relations* that hold between successive sentences or clauses in a discourse carry implications for the temporal relations between the eventualities they are used to describe. Moreover, these rhetorical relations need not be overtly expressed by conjunctions like *because* or by any other word or feature of grammatical construction. As often as not they are implicit. But they may be clearly recongizable even so and make the impact on the temporal dimension of the interpretation as well. This is one of the central messages both of the paper by Moens and Steedman and of the paper by Lascarides and Asher and we will look at such interactions in considerable detail when we com to discuss those papers. But some of Partee's examples have to do with this as well. Consider the following pair, in which the second sentences have been modified somewhat in order to heighten the effect. In (114.a) there is a tendency to interpret the stative claus as describing the result of the event of the switch being turned. And that is consistent with locating the state as holding at the reference point introduced by the event of the preceding sentence, in accordance with the interpretation rule for simple past tense stative sentences that we have been using in the examples above. But in (114.b) this impression is counteracted by the inclination to see the second sentence as giving the reason or motivation for the action of turning the switch. And as a consequence the temporal relation between state and event is now reversed: the event terminates the state rather than bringing it about. Encoding the state as holding at a time after the event would get this wrong.

(114) a. John turned the light switch. It was pitch dark.

b. John turned the light switch. It was too dark to read.

In order to be able to deal with cases like (114.b) we have to take account of a host of non-temporal facts and relations and not just of temporal ones. Theories which do this have to be much more complex than purely temporal theories such as the one laid out in Partees paper. (As of now, the SDRT of Asher is the only serious attempt to deal with the impact of rhetorical relations on temporal interaction.)

The remaining examples from Partee's paper that we will review all have to do with the role of *when*-clauses. The central point is that a central function of*when*-clauses is to temporally locate the eventualities of their main clauses. In particular, a *when*-clause can override the narrative that a part tense event clause determines via its associated reference time the temporal location of the eventuality introduced by the claus following it. For instance, (115.b) differs from (115.a) in what it implies about the time when the janitors came in. (115.a) conveys that the janitors came in just after the the people in the room started to leave, and thus presumably well before the room was empty. In (115.b) the *when*-clause *when the room was empty* overrides this aspect of the interpretation of (115.a), locating the entry of the janitors at the time when the room had become empty. That this time is conceived as later than the reference time introduced by the first sentence, Partee notes, is shown by the fact marginal acceptability of (115.c). The second sentence seems strange here precisely because the reference time associated with the event of the people *beginning to leave* the room is understood to be a time when the filing out is still in full swing, so that the room is *not* empty at that time.

(115) a. People began to leave. The janitors came in.

b. People began to leave. When the room was empty, the janitors came in.

c. ? People began to leave. The room was empty. The janitors came in.

when-clauses constitute just one among a large variety of different types of expressions that serve the temporal location of the eventualities described by the clauses to which they are attached or of which they are part. Other types are: (i) subordinate clauses beginning with the conjunctions *while*, *after*, *before*; (ii) prepositional phrases beginning with the prepositions *after*, *before*, *during*; (ii) adverbs like *an hour ago*, *two days later*, *once upon a time*, *recently*, *soon*; (iii) dating expressions like: *(on) the 15-th of March, nineteen forty six*, *(in) August 2014*, *today*, *yesterday*, *tomorrow*, *the day after tomorrow*, *(on) Thursday*, *last September*, *next strong*, and many more.

All of these will be called *temporal locating adverbials*. The list above is far from complete. But even as it is, it is a non-trivial task to provide an exhaustive description of the semantic contributions made by the types represented by it. We won't even begin to deal with this task here. The only thing we note is that all temporal locating adverbials have the function of locating the eventuality described by the sentence of clause with which they are associated – in that way in which adverbials can be associated with (or constituents of) larger phrases, viz. by *adjunction*. (Adjunction is a syntactic rule that attaches one phrase – in the case at hand: the temporal locating adverbial – to another phrase without changing the grammatical category of the other phrase. We have already encountered one example of adjunction in the syntactic tree for the combination of an *if*-clause and a main clause in (103.a), repeated here as (116.a). (116.b) repeats the syntactic tree adopted earlier, which was then displayed in (103.a). In this tree the *if*-clause, formed by combining an expression of category S with the conjunction *if*, is adjoined to the S-node of the main clause, with the result of a 'higher copy'of this S-node, which now incorporates the *if*-clause as one of its constituents.

(116) a. If Pedro owns a donkey, he beats it.

b.

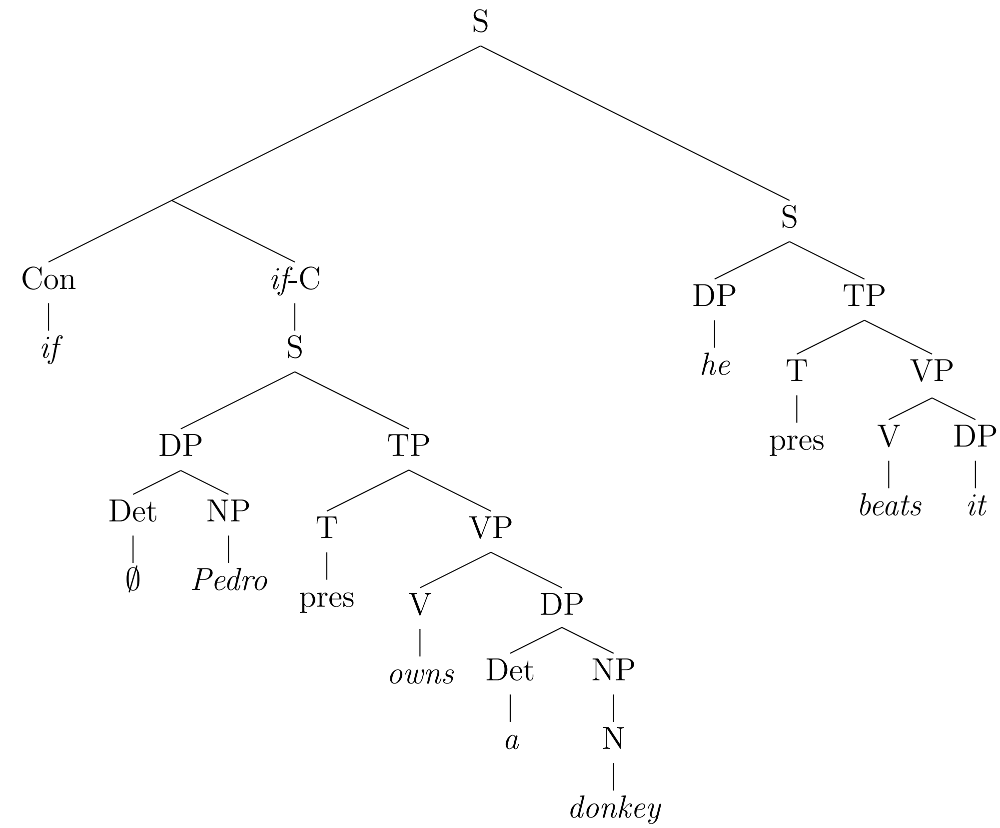

In the remainder of these comments on Partee (1984) we will only look at clause-initial occurrences of temporal locating adverbials and we will assume that these are always adjuncts to S. Adverbial clauses beginning with *when*, *after* etc. are among them, and for these the syntactic structure looks almost identical to that in (116.b).

Before saying more about the *when*-clause of (115.b), let us first have a look at the *when*-clause of (115.b) of the next example, in (117). Partee gives for this discourse the DRS in (16.b).

(117) a. Mary turned the corner. When John saw her, she crossed the street.

(16b)

(0) r_0 (0) now (2) e_1 (2) r_1 (5) u (8) e_2 (8) r_2 (11) v (13) e_3 (13) r_3

---

(0) $[r_p := r_0]$

(1) [[Mary turned the corner]]

(2) $e_1 < now$ (2) $e_1 \subseteq r_0$ (2) $e_1 < r_1$ (2) $r_1 < now$

(3) $[r_p := r_1]$

(4) e_1: $\boxed{\begin{array}{c} \text{(6)}\ w \\ \hline \text{(4) [[Mary turn the corner]]} \\ \text{(5) [[}u\text{ turn the corner]]} \\ \text{(6) } w = \text{the corner} \quad \text{(6) } u \text{ turn } w \end{array}}$

(5) u = Mary

(7) [[When John saw her, she crossed the street]]

(8) $e_2 < now$ (8) $r_2 < now$ (8) $r_1 < r_2$

(9) $[r_p := r_2]$

(10) $e_2 \subseteq r_2$

(10) e_2: $\boxed{\begin{array}{c} \text{(10) [[John see her]]} \\ \text{(11) [[}v\text{ see her]]} \\ \text{(12) } v \text{ see } u \end{array}}$

(11) v = John

(13) $e_3 < now$ (13) $r_3 < now$ (13) $e_3 \subseteq r_2$ (13) $e_3 < r_3$

(14) $[r_p := r_3]$

(15) e_3: $\boxed{\begin{array}{c} \text{(17)}\ z \\ \hline \text{(15) [[she cross the street]]} \\ \text{(16) [[}u\text{ cross the street]]} \\ \text{(17) } z = \text{the street} \quad \text{(17) } u \text{ cross } z \end{array}}$

However, we want, once again, to look more closely at how such a DRS can be constructed.

Both *when*-clause and main clause of the second sentence of (117) are event sentences and the interpretation of the combination is that the event of the main clause follows that of the *when*-clause; or, in the terms we have been using, that the event of the *when*-clause comes with an immediately following reference time, within which the event of the main clause is included. (118.a) gives the initial stage of the semantic representation construction for the *when*-clause of (117), consisting of a DRS for the first sentence and the syntactic structure of the *when*-clause. (118.b) gives the result of of the first two processing steps that I assume apply to this syntactic structure: (i) splitting the *when*-clause off from the main clause and making each into a separate constituent of the condition set of the DRS; and (ii) replacing *when* by a time discourse referent, which we mark with the subscript '$_{wcrp}$' (short for '*when*-clause reference point') to indicate the role it is going to play in the interpretation of the *when*-clause-main clause combination.

(118) a.

$$u \quad r_0 \quad e_1 \quad r_1$$

$$u = Mary \quad r_0 < r_1 < n \quad e_1 \subseteq r_0$$

$$e_1\text{: } turn - the - corner(u)$$

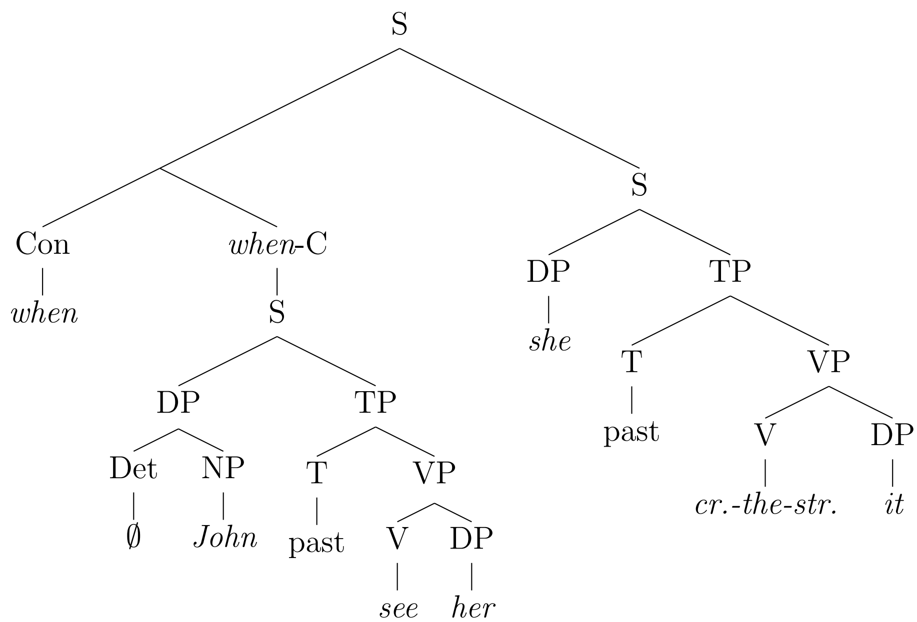

b. 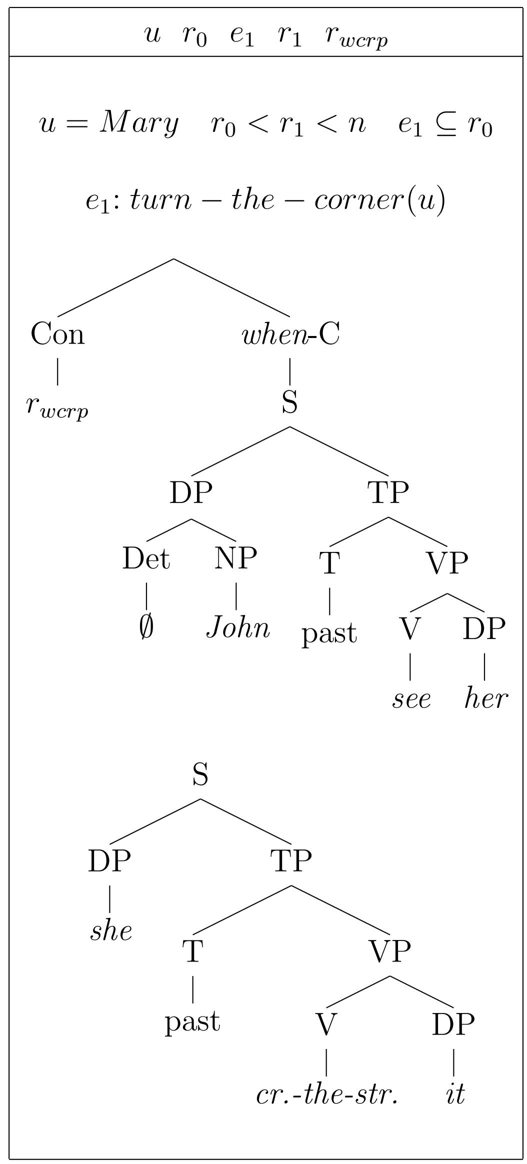

Interpretation of the structure below the S-node of the *when*-clause now proceeds in a manner that is just like what we have seen for other simple past tense clauses, but with one difference: Since we are dealing with a *when*-clause, which serves to locate the eventuality of its main clause, the *rp* for the main clause is not the reference time introduced by the previous sentence. Moreover, that reference point shouldn't be used as location time for the *when*-clause event either, Partee argues, for that event may well be much later than the event e which introduced that reference point. She seems to imply that the *when* clause event should come at *some* time *after* e, though even that seems questionable as a general principle, as shown by the following example (119), a variant of one that we will discuss at some length in the comments on the next paper, 'Temporal Ontology and Temporal Reference' by Moens and Steedman:

(119) The council did build the 39-th Street Bridge. When the architect drew up the plans, there was a good deal of resistance. But the proposal passed.

In this example, the drawing up of the plans surely didn't come *after* the time when the bridge was built.

What does seem correct, however, is that the *when*-clause introduces its own reference time, r_{wcrp}, which locates the main clause event as temporally included within it. Furthermore, r_{wcrp} comes 'in the wake of' the *when*-clause event just as main clause events come with a reference point in their wake – on that is 'immediately' after it, though Partee gives no way of making this notion precise. (She does use a special symbol for 'immediately after', something that hasn't been emulated here.)

The result of dealing with the *when*-clause is shown in (120.a). (120.b) shows the result of also processing the main clause. Again, this is almost identical to what we have seen before, but with one difference: It is now the *when*-clause reference point that serves as location time for the main clause event; so it is r_{wcrp} that will now occur on the right hand side of the reference point condition '$rp := r_{wcrp}$'. (120.c) wraps things up by giving the completed DRS for the first two sentences of (117). (120.c) is virtually identical with (120.b). The only difference is that rp has been removed.

(120) a.

$$
\begin{array}{|c|}
\hline
u \quad r_0 \quad e_1 \quad r_1 \quad r_{wcrp} \quad v \quad e_2 \quad w \\
\hline
u = Mary \quad r_0 < r_1 < n \quad e_1 \subseteq r_0 \\
e_1\text{: } turn - the - corner(u) \\
\\
v = John \quad w = u \quad e_2 < r_{wcrp} < n \\
e_2\text{: } see(v, w) \\
\hline
\end{array}
$$

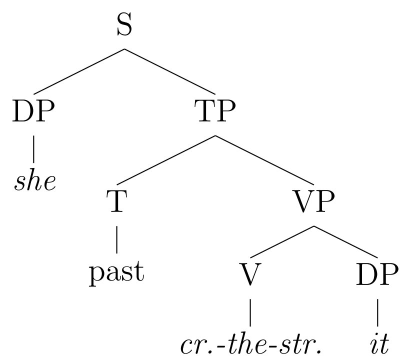

b.

$$
\begin{array}{|c|}
\hline
u \quad r_0 \quad e_1 \quad r_1 \quad r_{wcrp} \quad v \quad e_2 \quad w \quad z \quad e_3 \quad r_2 \\
\hline
u = Mary \quad r_0 < r_1 < n \quad e_1 \subseteq r_0 \\
e_1\text{: } turn - the - corner(u) \\
\\
v = John \quad w = u \quad e_2 < r_{wcrp} < n \\
e_2\text{: } see(v, w) \\
\\
z = u \quad rp := r_{wcrp} \quad e_3 \subseteq rp \\
e_3\text{: } cross - the - street(z) \\
\hline
\end{array}
$$

c.

$$
\begin{array}{|c|}
\hline
u \;\; r_0 \;\; e_1 \;\; r_1 \;\; r_{wcrp} \;\; v \;\; e_2 \;\; w \;\; z \;\; e_3 \;\; r_2 \\
\hline
u = Mary \quad r_0 < r_1 < n \quad e_1 \subseteq r_0 \\
e_1\text{: } turn - the - corner(u) \\
\\
v = John \quad w = u \quad e_2 < r_{wcrp} < n \\
e_2\text{: } see(v, w) \\
\\
z = u \quad e_3 \subseteq r_{wcrp} \\
e_3\text{: } cross - the - street(z) \\
\hline
\end{array}
$$

The *when*-clause of (115), *When the room was empty*, differs from that of (117) in that it apparently describes a state rather than an event. Intuitively it is clear in what way the *when*-clause of (115) locates its main clause event: the event is temporally included within the time when the room was empty, i.e. within the duration of the state of the room being empty. There are two ways, however, in which this result can be obtained. We can either (i) adopt the principle that when a *when*-clause is stative, then its r_{wcrp} is included within the state it describes or (ii) assume that the *when* in the *when*-clause of (115) requires an event sentence as complement and that therefore a coërcion from state description to event description is necessary before *when* can be semantically combined with it; and furthermore that coërcion in this case is from state description to description of the onset of the state. After this coërcion processing of the *when*-clause of (115) can then proceed just as for the *when*-clause of (117): the already introduced r_{wcrp} is located as immediately after this onset, and thus within the state whose onset it is.

At this point the second of these two options may seem unnecessarily prolix and round-about. In the light of further evidence it seems to be likely to be the correct one. But we lack the resources at this point to argue the point.

(121) contains a few more examples showing the same pattern as (117). (121.a) is an extension of (117) with a third sentence, which describes an event that can be seen as located within the reference point introduced by the main clause of the second sentence. (As an exercise, extend the DRS in (118.c) to a DRS for the entire discourse in (121.a) by processing this third sentence within the context provided by (118.c).) (121.b) is just like (121.a), except that in the second sentence *when*-clause and main clause have

been interchanged. Noteworthy about the effect of this reversal is that the interpretation we get isn't quite what we would get by constructing the representation for the first two sentences of (121.b) in the same way as we did for (117). A possible interpretation of the second sentence of (121.b) is that John noticed Mary as soon as she started crossing the street, or while she was crossing the street. In fact, for me these possibilities are more prominent than the one according to which John saw Mary only when her crossing of the street was complete. This is precisely the difference with (121.c), which makes the temporal succession of these two events – first the crossing, then the seeing – fully explicit.

(121) a. Mary turned the corner. When John saw her, she crossed the street. She hurried into a store.

b. Mary turned the corner. When she crossed the street, John saw her. She hurried into a store.

c. Mary turned the corner. After she crossed the street, John saw her. She hurried into a store.

d. Mary turned the corner. Before John saw her, she crossed the street. She hurried into a store.

It is tempting to think of *before* and *after* as converse relations – for any times t and t' $tbeforet'$ iff $t'aftert$, and likewise for events. But (121.d) shows that the matter is more complicated. it is true that the second sentence of (121.d) does allow for an interpretation on which it says the same thing as (121.c), viz. that there was an event of Mary crossing the street and an event of John seeing Mary and that the first event temporally preceded the second. But that is not the most salient reading for the second sentence of (121.d). The more salient reading is that according to which the second event never took place: Mary's crossing the street prevented John from seeing her.

This more prominent, so-called *non-veridical* reading of the second sentence of (121.d) – 'non-veridical' because the entire *when*-cause main clause combination can be true without the *when*-cause being true; 'veridical' is etymologically related to Lat. 'veritas' (Eng. truth) – has been perceived as a challenge for the theory of tense and aspect since the early seventies. The problem hasn't been completely solved to this day, although considerable progress has been made. I will do no more here than does Partee, by just noting that the problem exists.

The fact that (121.b) can be true even when the crossing event doesn't fully precede the seeing event is consistent with observations that were made by a number of researchers around the time of Partee's second paper or somewhat later, including three of the papers subsequently discussed in this class (those by Moens and Steedman, by Webber and by Lascarides and Asher). All these authors point out that non-temporal factors must be taken into account if the temporal relations are to be identified with greater accuracy; but do this in a theoretically convincing way, let alone in a computationally tractable one, has proved extremely difficult.

Sentences with *after*-clauses, on the other hand, are immune to the variability that is illustrated by the comparison of (121.b) and (121.a). This is shown clearly by (121.c), which gets the same kind of unequivocal interpretation that our algorithm yielded (problematically as we have now seen) for (121.a). This difference is evidently a difference between the lexical meanings of the conjunctions *when* and *after* and should be reflected in different processing of those conjunctions. On way to guarantee complete temporal precedence of *after*-clause event to main clause event, which will be unaffected by the modifications in the processing of *when*-sentences for which we have just acknowledged the need, is to introduce two time discourse referents as part of dealing with *after*, one, r_{acet} (for 'after clause event time'), to locate the after-cause eventuality, and a second one, r_{acrp} (for 'after clause reference point'), to locate the main clause eventuality, with the two standing in the relation $r_{acet} < r_{acrp}$. With two event clauses this will have the effect that the *after*-clause event is included within r_{acet} and the main clause event is included within r_{acrp} and thus that th first completely precedes the second. (Exercise: Construct the DRS for (121.c) using this principle for the processing of *after*.)

4.5.1 Quantificational *when*-sentences

As Partee already observed in her 1973 paper, tenses can also play a part in sentence interpretations that involve quantificationally bound variables ranging over times. This is true in particular for sentences containing *when*-clauses, as shown by the following sentences.

(122) a. Always/Usually/Often/Sometimes, when Mary telephoned, Fred was asleep.

b. When Mary telephoned, Fred was always/usually/often/sometimes/rarely asleep.

c. Whenever Mary telephoned, Fred was asleep.

d. When Mary telephones, Fred is asleep.

e. When Mary telephoned, Fred was asleep.

As I put things just now, tenses seem to be capable of playing a part in interpretations that involve quantification over times, but both the interpretation algorithms discussed above and the forms of some of the sentences in (122) indicate that they are only one among several factors that are responsible for this kind of quantificational structure. In our treatment of the *when*-clause in (117) the tense is directly instrumental only in selecting the time r_{wcrp} as reference time and temporally relating the described eventuality to it. In that example r_{wcrp} represents a single time. but as shown by the examples in (122), it is also possible for it to get bound by some quantifier. In (122.a) the quantifying element plainly seems to be the adverb (*always*, *usually*, etc.). The usual assumption is that quantifying adverb induces a quantification structure – in DRT: a duplex condition – with the *when*- clause forming the *restrictor* (the left hand DRS of the duplex condition) and the main clause the *nuclear scope* (the right hand DRS). In particular, for the case where the adverb is *always*, we get for (122.a) the representation in (123).

(123)

$$\begin{array}{|c|}\hline u \quad v \\ \hline u = Mary \quad v = Fred \\ \\ \begin{array}{|c|}\hline r_{wcrp} \quad e_1 \\ \hline e_1 < r_{wcrp} < n \\ e_1\text{: } telephone(u) \\ \hline \end{array} \;\; \left\langle \begin{array}{c} \forall \\ r_{wcrp} \end{array} \right\rangle \;\; \begin{array}{|c|}\hline s_2 \\ \hline r_{wcrp} \subseteq s_2 \\ s_2\text{: } be-asleep(v) \\ \hline \end{array} \\ \hline \end{array}$$

Representations for the other sentences enumerated in (122.a) are like (123), except for having different quantifier specifications within the diamond of the duplex condition. The same goes for the sentences in (122.b), though here a further story is needed to explain why the quantifying adverb can do the work it evidently does do, while apparently occurring within the scope of the *when*-clause (rather than the other way round). The conjunction *whenever* in (122.c) has the conjunction role of *when* in (122.a,b) and the universal quantifier role of *always* packed together. So the operations dealing with this conjunction have to cover both the contribution that are made by *when* and the contribution that are made by *always* to the construction of the representation in (123). (These are just some of the many details that need to be accounted for by a theory that deals with the full range of forms that are made available in English for these kinds of temporal quantifications. The little that has been said here is just to give a taste of what awaits the one who wants to provide full coverage of these forms.)

A quantified sort of interpretation is also prominent, and perhaps even the only possible one, for the present tense sentence (122.d). This has to do with the fact that, as we have seen, the standard use of the simple present tense in English is incompatible with event descriptions, in that it cannot be used that the described event occurs at the utterance time. Occurrences of the simple present tense with event sentence are therefore typically reinterpreted as habitual or generic, and the presence of a *when*-clause actually facilitates such an interpretation insofar as it provides a type of situation such that the sentence can b understood as claiming that the main clause is true in the different instances of that type.

Such habitual, generic or plainly quantified interpretations are easier to obtain for present tense *when*-clause main clause combinations than when such sentences occur in the past tense. In fact, so far we have been proceeding on the assumption that past tense sentences of this kind had an episodic interpretation – the *when*-clause describes some particular eventuality in the past and the main clause eventuality is temporally located through being temporally related to the *when*-clause eventuality in some way – and that that was it. But in fact, past tense *when*-clause main clause combinations also allow for quantified interpretations, though this dos seem to require some help from the context. On such context, for the sentence in (122.e), is that provided in (124).

(124) In the end Mary just gave up on Fred. When she telephoned, he was asleep. When she rang his doorbell, he wasn't at home. It just didn't seem possible to make contact with him.

The examples with *when*-clauses that we have discussed is clearly just the tip of an iceberg – that of all the interactions between tenses and temporal adverbials that are found in a language like English. This is a topic, however, to which Partee's papers are not in the first instance devoted to and so there is no justification for pursuing it further in these comments. On the other hand there is a range of further points discussed in Partee (1984) that we have not yet touched upon. But I leave matters at this, and move on to the next paper.

4.6 Moens and Steedman: Temporal Ontology and Temporal Reference

4.6.1 Introduction

This is another of the seminal papers in the theory of tense and aspect. Its most important contributions are these:

1. The various types of aspectual *coërcion* to which verbs from the different Vendler categories can be subjected as part of tailoring them to current communicational needs.

2. The processes that lead the interpreter of a text or piece of discourse to his understanding of the temporal relations between the events and states that are mentioned and described usually involve other aspects of discourse understanding as well. Often the identification of temporal relations is a secondary effect of those non-temporal aspects of interpretation.
Aspect-modifying operations come in two forms. Some of them are *overt* in that is there is a morpho-syntactic operation that transforms the verb into a audibly different form that heralds the change. Others are *covert*: they are semantic options without ostensible morphological or syntactic manifestation. But even though they are not 'visible at the surface' of what is said or written, they may be just as important to understanding as the overt operations. In fact, languages often differ in that what is an overt operation in one is a covert operation in the other. But nevertheless, the semantic implications of the operation are as much part of the semantic competence of the speakers of the language in which it is covert as it is for the speakers of the language in which it is overt. An example close to home is the English progressive. In English this form is not only very common, it is also obligatorily used in contexts where verbs are used with the aspectual profile which is expressed by the progressive but not by the non-progressive form. Other languages, even languages spoken nearby, and genealogically close, such as German or Dutch, do not have a progressive – they only have cumbersome circumlocutions, which sound clumsy, are seldom used and are not obligatory in most of the situations in which the English progressive is used obligatorily. (Against the background of how much languages can differ in the ways they express aspectual properties and temporal relations these differences

may appear as minor, but for a proper understanding of how text and aspect in, say, English and German, respectively the difference is important enough.)

Central to the discussion of this issue in 'Temporal Ontology and Temporal Reference' is the diagram which we repeat here as (125).

(125)

Central to the interpretation of this diagram is the notion of the *Nucleus*, which Moens and Steedman introduced as the core of aspectual structure. The Nucleus consists of a *preparatory phase* (a process), a *culmination* and a *consequent state*. It corresponds closely to what Vendler identifies as the aspectual structure of accomplishment verbs such as transitive *write* or *build*, except that the resulting states of the events described by such verbs – that of the letter being completely written or the house that was under construction being ready for habitation – are now included with the events that produce them.

As Vendler characterizes his other three categories, the eventualities described by those can also be seen as parts of the Nucleus. Achievement verbs and phrases describe events that consist just of the culmination point of the Nucleus. Activity verbs and phrases can be seen as describing what corresponds just to the preparatory phase of the Nucleus, and state verbs and phrases can be thought of as describing what corresponds to its consequent state.

The last correspondence – between stative verbs and phrases and the consequent state part of the Nucleus – is perhaps the least immediately persuasive one, and what makes it less than persuasive has to do with a point that is implicit in what Moens and Steedman say about culmination, but that hasn't been made explicit so far. First consider the achievement verb *die*. We can distinguish between two kinds of events that this verb can be used to describe. On the one hand there are the cases where someone is in the process of dying for hours – or days or weeks or even longer than that – until death at last occurs. And on the other there are the cases where someone dies instantaneously, e.g. when he gets his head blown off by a grenade. In the second case there is no period leading up to the moment of death of which we can say that the victim 'was dying'. But what the two cases have in common is that up to the time t of which it is true to say: 'x died at t', x was not dead. In other words, t separates the state of x not being dead from that of y being dead. And thus it marks not only the starting point of a new state – that of x being dead – but also the termination of some other state – that of x not being dead.

This is a general feature of culmination points, not only those that are described by achievement verbs and phrases, but also those that are part of the event complexes described by accomplishment phrases. For instance, consider one of our stock examples of accomplishments, the phrase *write a letter*. The culmination of each complete letter writing event marks the beginning of a letter (the one that was being written) existing as a finished letter. But at the same time, and by the same token, this culmination marks the end of the opposite state, that of there not being a finished letter.

The general moral of this is that Moens and Steedman's Nucleus can be thought of as involving two states, the *post-state* (or consequent stat) which starts at culmination time, and the *pre-state*, which ends at culmination time and which is the opposite of the post-state, in that th conditions that characterize it are incompatible with those characterizing the post-state. And as for the question how states fit into the general schema of event structures provided by the Nucleus, states can be thought of as corresponding to the pre-state part of the Nucleus just as plausibly as they can be thought of as corresponding to its post- or consequent state. And each of these two possible correspondences suggests its own coëcion: from a state to its starting point (stats as consequent states) and from the state to its endpoint (states as pre-states). The first of these coërcions is well-attested in many languages. In the literature it is often referred to as *inchoative*. Well-known examples are those involving the verb *know*. On of them, 'When Pete entered the room, I knew there was trouble.', we already encountered. Here the verb *know* is used to refer to the event of the speaker *coming to know* that there was trouble – the transition from not knowing this to knowing it. In this case the reinterpretation is suggested by the structure of the sentence combined with the content of *when*-clause and main clause. (The suggestion can be overruled by adding *already*, as in 'When Pete entered the room, I already knew there was trouble.' and with heavy stress on *I* the sentence without *already* seems neutral between the 'coming to know' and the 'already knew' interpretation.) But it is also possible to make the inchoative interpretation explicit by embedding the main clause under the verb *come to*, or by replacing *know* by *realized* as in 'When Pete entered the room, I realized/came to know that there was trouble.'

I do not know of coërcions to the ends of states (that is: of cases in which a state describing verb or verb phrase V is coërced to a verb or verb phrase describing the ends of the states described by V). But there are many cases where a state describing verb or VPs is coërced to a verb or VP which treats the described state as an events, with a well-defined end, which functions as a kind of pseudo-culmination. Moreover, the event descriptions that result from such coërcions treat the events they describe as 'punctual': from the perspective of the discourse in which the coërced description occurs the event described behaves like a single, undivided temporal instant.

What we have just been saying about coërcion from state to event descriptions is at odds with the Moens and Steedman diagram, which doesn't have any arcs going from states to events. (I have no explanation for why such arcs are absent from the diagram.) On the other hand the diagram indicates that coërcions to descriptions of punctual events do not always start from descriptions of states: quite a number of arcs land on the node labeled 'point'. All these arcs are justified insofar as there are transformations of eventuality descriptions that instantiates those arcs.

But what is a 'point'? Or better: What is a punctual event? Evidently, what is meant cannot be 'punctual in the sense of physical time', the points of Moens and Steedman's diagram cannot be events that last for just one and only one instant of that time. Such events are too short for us to notice; if there are any, they are hypothetical transitions from one state to another, such as crossing the exact midpoint on the path from A to B. But even for such events the idea that they are instantaneous in the sense of physical time seems problematic; for where exactly does the path from A to B begin, and where does it end? And if we cannot answer those questions we cannot determine the exact midpoint either. In any case, the notion of 'treating' an event as 'a point' by describing it in a certain way (or by interpreting its description in a certain way) seems to refute the suggestion that 'point' could mean 'instant of physical time in this context. If the very same event can be either 'treated' as point-like or as temporal extended, clearly these properties cannot be properties of the event as such, properties that it has independently from how we choose to look at it or describe it.

We can get some kind of grip on what could be meant by 'point-like' and 'temporal extended' in the present context of discussion by returning to the Russell-Wiener construction of temporal instant structures from underlying event structures. Recall the following features of that construction: temporal instants are constructed as maximal sets of pairwise overlapping events and an event e from the base structure 'is going on' at a thus constructed instant i iff i (as a set of events) contains e as a member. So an event will be instantaneous in the sense of the time structure obtained via the Russell-Wiener construction if it belongs to just one of the constructed instants and non-instantaneous if it belongs to more than one.

Now consider a bit of narrative discourse, such as the toy examples discussed in Partee (1984). Let us focus, for instance, on (107.a), repeated here as (126.a), with the DRS that Partee proposes for it, also repeated here under (126.b).

(126) a. John got up, went to the window, and raised the blind.
e_1 e_2 e_3

It was light out. He pulled the blind down and went back to bed.
s_1 e_4 e_5

He wasn't ready to face the day. He was too depressed.
s_2 s_3

b.

e_1 e_2 e_3 e_4 e_5 s_1 s_2 s_3 r_s
$e_1 < e_2 < e_3 < e_4 < e_5$
e_1: John get up
e_2: John go up to the window
e_3: John raise the blind
e_4: John pull the blind down
e_5: John go back to bed
s_1 O e_3
s_2 O e_5
s_3 O e_5
s_1: It be light out
s_2: John not be ready to face the day
s_3: John be too depressed

(126.b) mentions, through the discourse referents in its DRS universe, a small number of events, temporally ordered in the way the DRS indicates and thus form a structure to which the Russell-Wiener construction can be applied. But since there are only a handful of events in the base structure, there will only be a small number of instants in the resulting instant structure. And because there are so few instants, it is comparatively easy for an event to emerge as instantaneous in the sense of this sparse instant structure. In fact, it is not hard to verify that when the instant structure is built just from the *events* mentioned in (126.b), then each of those events will give rise to an instant all its own: if e is such an event, then the corresponding instant will be $\{e\}$ and so e will be instantaneous in the sense of the constructed time structure, with $\{e\}$ the one and only instant of the structure, in the sense of this instant structure, at which e is going on. (Moreover, since all the events are instantaneous in this case, it is also possible to say something about at which instants of the constructed structure the states mentioned in (126.b) are going on: when (126.b) says of a state s that it overlaps some event e, then it follows that s must hold at the one time $\{e\}$ at which e is going on. That is of course only partial information about the duration of those states. As Partee points out, all the discourse in question implies is that each of the states it introduces overlaps with the event of the last event sentence; with which other events the state might overlap the discourse leaves open, and that is also true of the DRS that she constructs for it.)

So much for the 'discourse-internal' time that can b constructed from the DRS in (126.b). But all this *is* all 'discourse-internal': all it says about any real events and states in the world that a discourse like (126.a) could be used to describe is through the perspective that the discourse imposes on them. And that isn't everything. To see this, let us assume that (126.a) is a truthful description of some real world episode, consisting of the actual events e_1,.., e_5 and states s_1,.., s_3. What can we say about these real events? Are *they* instantaneous? The answer to that question depends of course once more on what the instant structure is with respect to which the question is intended. But if that instant structure is, as it most naturally would be, that of the of the physical time of the real world to which these events belong, then the answer can be, and almost certainly will be, that these events are *not* instantaneous, that they occupy intervals of time, which contain many instants; in fact, for all we know these intervals may contain non-denumerably many instants.

If that is the right answer, then we arrive at the following picture: (i) the events represented by the discourse referents e_i are instantaneous in the sense of 'discourse time'; (ii) the real events e_i represented by these discourse referents are non-instantaneous in real time (the time of physics that orders events in the external world); (iii) nevertheless the description that the discourse provides of the episode to which the events e_i belong is a truthful one. If the combination of (i), (ii) and (iii) has a ring of the paradoxical about it, paradox can be easily dispelled: the discourse describes the episode correctly iff the correlation f between the discourse referents e_i and the events e_i *verifies* the properties and relations attributed to e_i by the DRS-conditions of the semantic representation (126.b) of the discourse: the events $e_i = f(e_i)$ have those properties and stand in those relations. (More accurately: it must be possible to extend f to the remaining discourse referents in the DRS universe of (126.b) so that this extension f' of f satisfies all the conditions in the condition set of (126.b); this is the standard definition of *truthful*, or *correct* description in DRT.) That even when this correctness condition is fulfilled the events represented by the e_i are non-instantaneous in one sense (that of physical time) and instantaneous in another sense (that of the discourse time) need not surprise us. For the discourse time, which is constructed from a only a small selection of the events that can be found in the real world, is therefore much *coarser* than the time we can obtain by applying the Russell-Wiener construction to the totality of real world events: The instants of this latter time structure will be much richer; it is the product of, and therefore reflects, many more instances of temporal precedence and overlap than does the much simpler instant structure generated by the small selection of events that are represented in (126.b).

This is a general fact about the Russell-Wiener construction: When you apply it to two event structures $< E, \prec, O >$ and $< E', \prec', O' >$ and $< E', \prec', O' >$ is a substructure of $< E, \prec, O >$, then the instant structure $< I', <' >$ obtained from $< E', \prec', O' >$ will be a 'homomorphic contraction' of the instant structure $< I, < >$ obtained from $< E, \prec, O >$: map each instant i from $< I, < >$ to the set of those events in it that belong to E': this will in general map several instants from $< I, < >$ onto one and the same instant i' from $< I', <' >$, but the order is preserved:

if i_1 is mapped onto i'_1, i_2 is mapped onto i'_2 and $i'_1 <' i'_2$, then $i_1 <' i_2$.

When, as in the kind of case we are discussing, $< E', \prec', O' >$ is only a tiny part of $< E, \prec, O >$, then there will be a huge amount of contraction: many, many instants from $< I, < >$ will be mapped onto the same instant i' from $< I', <' >$. So it is possible for an e from $< E, \prec, O >$ to b going on at only this one instant i' of the (very coarse) instant structure $< I', <' >$, while going on at many instants from the much more finely grained instant structure $< I, < >$, viz. all those that are mapped onto i' by the contraction map. And this will be so in particular for the events e that belong to the substructure $< E', \prec', O' >$.

In these last two paragraphs we have compared the instant structure $< I', <' >$ that we can think of as the discourse tim for a given discourse with the instant structure $< I, < >$ generated by the totality E of all the events in the world of which the discourse speaks. But is this second structure the same as physical time? Not necessarily. As we saw when discussing the relations between times and events in the first part of the course, different views can be and have been held about this relationship, with Leibniz as early advocate of the 'relational' view, according to which physical time is nothing but the reflection of the temporal relations between the events that make up our world and Newton as an advocate of the 'non-relational' position according to which time is given in advance of and thus independently from what happens in our world as it happens to have evolved in the course of (that) time. It is reasonable to assume that for the Leinizian physical time just is the structure $< I, <>$. For the Newtonian there is no reason to make this identification. But if it is true that whenever two real world events e and e' stand in a relation of temporal precedence or overlap, then this manifests itself in their respective relations to physical time – when $e \prec e'$, then every instant t of physical time at which e is going on will be earlier than every instant t' at which e' is going on, and when $e \; O \; e'$, then there is an instant t of physical time such that both e and e' are going on at t – then physical time will be at last as fine-grained as $< I, <>$ and stand also in the same kind of homomorphic contraction relation to $< I', <'>$ as $< I, <>$ does. In short, what w have ben saying about the relation between $< I', <'>$ and $< I, <>$ also applies to that between $< I', <'>$ and physical time.

The moral of all this for Moens and Steedman is that when they speak of 'points' (as they do in the diagram we are discussing), then it is in the sense of discourse time that they should be understood. And the same applies to pretty much all uses of 'point' and 'punctual' in the tense and aspect literature. If punctuality in the sense of physical time plays any role in matters of language and language-driven conceptualization it is at best a marginal one. But it is nevertheless important for us not only to have a grasp of what the notion of punctuality is that *is* important for language and the conceptualizations that come with it, but also how that notion is related to the properties of real time.

So much for the notion of 'point' as it occurs in the Moens and Steedman diagram. There are a couple of further remarks about what the claims are that are made by this diagram, before we move on to the other major aspect of the paper, the intimate intertwining of temporal and non-temporal relations. The first question concerns the nature of the 'transitions' that are represented by the arcs of the diagram. What form do such transitions take? Answer: That varies - between transitions and between languages. First, let us focus on English and consider the diagram as only making claims about it. What does the diagram, understood in this narrower sense, tell us about English?

There are two main forms the transitions attested by the arcs of the diagram can take: either a verb or verb phrase used to describe events or event complexes of the type of the arc's starting node are actually changed (such as when we put the verb into its progressive form or turn it into a perfect); or alternatively, the verb or verb phrase isn't changed, but the linguistic or extra-linguistic context in which it appears forces or invites a different interpretation of it from the one it gets in those contexts that speakers conceive of as the more basic ones. We have already found reasons why these two options must be sharply distinguished: some transitions, such as for instance that from accomplishment descriptions to descriptions of the preparatory processes of accomplishments, cannot be left to the interpreter; they *must* be marked overtly, by turning the verb into a progressive. (That is why uses of accomplishment phrases in the simple present so often sound bad or even ungrammatical.)

But there are also many transitions that can be left to the interpreter and where no overt transformation is required. An example is the interpretation of an event verb phrase as expressing habitual occurring of the type of events that are described by the phrase in its basic use, or of a disposition to perform such events. An instance was the use of *beats* in the two sentence discourse 'Pedro owns a donkey. He beats it.', which we have discussed at length. There an 'episodic' interpretation, according to which a beating by Pedro of his donkey is going on at the time of speech is impossible – that would require the present progressive *is beating*. So an alternative interpretation of the verb has to be found, and this option is indeed available, in the form of recasting the verb as describing a habit or disposition towards events of its basic sort.

In the context of the present discussion this example is especially instructive because it shows the two types of aspectual transition side by side: A progressive reinterpretation of *beats* would be compatible with the special demands of the standard use of the present tense but in English it isn't available for the form *beats*; that interpretation would require the overt transformation into a progressive. On the other hand, interpreting the verb as description of a habit or disposition is possible without overt transformation; and since such an interpretation is also compatible with the standard use of the present tense, that is the one that we get.

Arguably the story about *beats* in our example is even more involved than this. It has been suggested (and I believe the suggestion is right) that progressive forms of events should regarded as descriptions of states and that this is true both for the progressives of telic event verb phrases (accomplishments and achievements) and for those of the non-telic event verb phrases (the activities, or process descriptions in the terminology of Moens and Steedman). On this view a progressive *be running*, formed from the activity verb *run*, also involves a transition, from an activity description to the description of states which consist in such activities going on. If this is right, then forming the progressive of a telic verb phrase such as *write a letter* involves two transitions: (i) from the type of event complex that is described by the basic use of the phrase (the full nucleus of an accomplishment verb phrase) to the type of the preparative processes of such complexes, and (ii) form the description of such processes to the description of the corresponding state. Progressives of achievement phrases such as the verb *die* can be analyzed along similar lines: first a transition to the description of an event complex that includes a preparatory phase (when the subject is in the process of dying); then a narrowing down to this preparatory phase by dropping the culmination point and consequent state and then the transition from process to state. (If this is the analysis we adopt, then the progressive of an achievement phrase involves three steps.)

This multi-stage analysis of the progressive of telic phrases can be made more explicit by seeing them as a succession of tacit and overt transformations. Consider the somewhat simpler case of accomplishments. With the formation of progressives of accomplishment it has in common the transition from process to state. But it differs from progressives like that of *run* in that this transition is preceded by a the transition from accomplishment description to activity/process description. One way to account for both the similarities and the differences between these two cases of progressive formation is to see the progressive as an operation that is directly applicable only to process descriptions. When it is given an activity verb like *run* as input, then it can apply directly. But when its input is a verb phrase like *write a letter*, then a preliminary adjustment is needed, in which the input phrase is first transformed into a process description. (As it stands, the Moens and Steedman diagram doesn't perhaps make fully explicit what this transition amounts to; but I believe it is fair to assume that the arc from 'culmination' to 'process is intended to capture this transition.) (A similar story can evidently be told about what happens when the progressive is applied to an achievement phrase; I won't elaborate the point.)

The preliminary transition from accomplishment description to process description that, according to what was said in the last paragraph, is involved in forming the progressive of an accomplishment is a case of coërcion: A given operator is presented with an input that dies not conform to its requirements, but an adjustment of that input is possible and will therefore be carried out as part of the application of the operator. The progressive is one of a number of aspectual operators of English that may involve coërcion in that they allow for certain kinds of adjustments of inputs that do not fit their input requirements. The present tense (in its standard use) is another. In the light of the discussion above we can now make more precise what its power to coërce consists in. The standard use of the present tense also selects for state descriptions. And that means that when it is combined with any kind of event description coërcion to a state description is needed. (Note that this is just as true for non-telic (process/activity) as for telic (achievement or accomplishment) event descriptions. 'He runs' is also (virtually?) impossible to interpret as describing a current bout of running, and is more naturally understood as saying that 'he' is a runner; in this regard activity phrases are like accomplishment and achievement phrases.) In other words, the present tense, in its standard use, goes without further ado only with stative verb phrases. These can be either lexical statives such as *know*, *love* or *be sick* or stative phrases derived from underlying event descriptions. Of the latter we have already identified progressive forms as one variety. But if our story is to hold water, then habitual and dispositional reinterpretations must also be regarded as descriptions of states. But that is not all that implausible on intuitive grounds. Dispositions are generally treated – in philosophy, and also if perhaps less emphatically so, in linguistics – as states. And habituals are most naturally thought of as holding over extended periods of time in virtue of repeated, regular and often predictable occurrences of events of the relevant type.

Part of our story about the aspectual restrictions placed on the standard use of the present tense is thus that both progressive forms and habitual or dispositional (re-)interpretations of non-progressive forms of event descriptions are formally treated by the grammar of English as descriptions of states. For a speaker who wants to make a standard use of the present tense and wants to use an event verb or verb phrase in conjunction with it, there are thus (at least) two options: talk about a current process or use the verb or verb phrase in a habitual or dispositional sense. The first of these requires an explicit transformation (from non-progressive to progressive); the second option forbids it. It is part of the grammar of both the progressive and of the standard use of the present tense that they 'select for' process descriptions and state descriptions, respectively, but also that they allow for a particular range of adjustment operations for input descriptions that do not fit their selection requirements. In this sense, and to this extent, aspect coëcion is a grammatical phenomenon.

Up to this point we have been speaking about the Moens-Steedman diagram in relation to English. The diagram could also be taken as saying something about human languages more generally. It is hard to see how it could be interpreted as making a statement about *all* human languages, as there ar quite a few that don't have verb tenses of the kind we find in English and other Indo-european languages. And it wouldn't be easy to circumscribe the totality of those languages spoken around the globe to which the diagram could be regarded as applicable. But on the other hand there are certain (comparatively small) families or groups of languages for which such a diagram clearly makes sense. For instance, we might consider the diagram as applying to th Germanic languages (in the broad sense of 'Germanic' in which it includes English and the Scandinavian languages as well as German, Dutch and Yiddish) and the Romance languages. And if it were claimed that the diagram is applicable to the members of such a family, that would mean that exactly the same aspectual transitions are possible in each of them.

Even for the family just suggested there is no hope of finding the predictions made by the Moens&Steedman diagram confirmed. For we have already seen that the diagram doesn't seem fully adequate for English, as it ignores all transitions to states. But perhaps another diagram could be drawn up that is adequate for English and that turns out to be equally adequate for the other languages in our chosen group. But note well that if such a diagram can be found at all, it will only capture what transitions are possible in *some* form. That we cannot hope to do better than that is shown by the fact that what in English can and must be expressed by using the progressive is not overtly expressible – not at least in a comparably simple and straightforward way – in other languages of the group. Neither French nor German have a progressive of the sort English does, but they allow progressive reinterpretation of event verbs without any accompanying change in form. So the English dialogue in

(127.a) is translated into French and German as in (127.b) and (127.c), with simple presents where English has present progressives. (As noted earlier, this is why native speakers of French, German and many othr languages so often make the mistake of using a simple present in English where the correct or natural form is a progressive.)

(127) a. A: What are you doing?
B: I am eating an apple.

b. A: Qu'est-ce que tu fais là?
B: Je mange une pomme.

c. A: Was machst Du da?
B: Ich esse einen Apfel.

In short, even closely related languages may differ in how they implement the realization of certain aspect transitions. It should be added that languages which lack simple ways of expressing a certain transition and allow for reinterpretation without morphological change instead, will as a rule have certain ways of making the transition explicit nevertheless. But these transitions will tend to be verbose and clumsy and speakers will avoid them unless they want to make it absolutely clear that it is the transformed aspectual meaning they want. For instance, in the German example B could have replied with 'Ich bin dabei, einen Apfel zu essen', which means something like 'I am engaged in eating an apple'. But in the given situation such a reply would be rather marked, as if B wanted to say that he was engaged in some sort of project.

This discussion of the Moens and Steedman's diagram has been quite critical. So it is well to conclude this discussion with the following: even if some details of the diagram may be hard to justify, its over-all message is a very important one; and it was that especially at the time when the paper was written. The over-all message is that the aspectual properties of lexical verbs are just one of the inputs to the complex process that determines the aspectual properties of complete clauses and sentences. Even after the verb has been chosen, a considerable variety of further factors may be involved in determining what the ultimate aspectual properties will be and many of these take the form of operations that modify the aspectual profile of the expressions on which they operate into a different profile. Aspectual determination and modification are part of the computation of meaning all the way from the germ provided by the lexical verb to the highest levels of the syntactic structure of clause or sentence.

This is all that these comments have to say about aspectual transitions. The remainder of these comments on Moens and Steedman will be concerned with the interactions between the temporal and non-temporal relations between eventualities.

4.6.2 Temporal and Non-Temporal Relations

The second main point of 'Temporal Ontology and Temporal Reference' is one that is also the central issue the next two papers in our sequence, those by Webber and Lascarides and Asher. This is that the temporal relations between the eventualities mentioned in a discourse interact with non-temporal temporal relations, and therefore cannot be determined on their own. In part this is, Moens and Steedman argue, because some of the words that others had treated as making purely temporal contributions actually do more than that. A prominent example in their paper is the word *when*. In Partee's papers *when*-clauses are analyzed as temporal locators of the eventualities described in their main clauses, with an emphasis on the similarities between complex sentences of the form '*When* S_1, S_2' on the one hand and sentence sequences of the form 'S_1. S_1' on the other. On her account the difference between '*When* S_1, S_2' and 'S_1. S_2' coms out only when they are embedded in a larger context. For instance, she claims, a discourse of the form 'S_0'. '*When* S_1, S_2' is subject to other principles of interpretation than one of the form 'S_0. S_1. S_2'. (When S_0 describes an event e_0, then the *when*-clause eventuality, though normally taken to be later than this event, need not be understood as 'next in the sequence of events' in the way that, according to her, one should interpret the eventuality introduced by the next clause a series of simple main clauses.) But the temporal relation between the eventualities described in S_1 and S_2 is supposed to be the same whether they are combined as '*When* S_1, S_2' or as S_1. S_2'. And in either case the relations are assumed to be determined by purely formal properties of the sentence or discourse.

To show that the temporal relations between main clause eventuality and *when*-clause eventuality depend in more than the temporal meaning of the sentence conjunction *when* and the tenses of the two clauses Moens and

Steedman focus on examples like those in (128).

(128) When they built the 39th Street bridge ..

- a. a local architect drew up the plans.
- b. they used the best materials. (*)
- c. they solved most of their traffic problems.

Each of the events described by the main clauses in (127), Moens and Steedman observe, stands in a different temporal relation to the event of the *when*-clause. There is a sense in which that seems right: What is described in (127.b) must have been part of, and thus simultaneous with, the actual building of the bridge; (127.a) must have taken place (one would hope) before the actual building began and (127.c) can be understood as the result of building the bridge and thus – perhaps – as a consequence of the building and therefore simultaneous or overlapping with the consequent state of the *when*-clause event. These judgments are obviously based on our understanding of what the *when*-clause and the different main clauses mean and on our knowledge of how bridges are normally built and what they are for. They show that the processing of *when*-clause-main clause combinations proposed by Partee cannot be right in general. And the same objection applies to sequences of simple past main clauses, as we will see at length when discussing the paper by Webber which is the next one on our list. What examples like those in (127) indicate is that the best that can be concluded in general about the temporal relation between the event of the main clause event of a *when*-clause-main clause combination (and likewise the event introduced by the next sentence in a discourse consisting of simple sentences in the simple past tense) must be in some temporal proximity to the *when*-clause event (or to the event of the preceding event sentence). What cannot be concluded in general is that the former follows upon the second.

One way in which we can do justice to this is not to assume that the reference times that are introduced by *when*-clause or main clause events follow those events, but that they are intervals that not only include the events that they are used to locate, but also the events by which they are introduced themselves. Of course this will not impose any real constraints on temporal location unless we say something about what these intervals are.

But here the Moens-Steedman notion of the nucleus may come to our rescue: We might assume that the reference time is just the time spanned by the nucleus described by the *when*-clause or preceding main clause, and thus say that the new event, by being temporally included within this duration, is temporally included within some part or parts of the nucleus of its predecessor. *Where* within this nucleus the new event is situated may then depend on considerations that have to do with causal relations (or 'contingency' relations, as Moens and Steedman put it, but without saying much about what 'contingency' really is) and on all sorts of information about how our world normally works, which typically is part of any competent speaker's intellectual baggage but notoriously hard to formalize. Thus, in the example before us the nucleus introduced by the *when*-clause would consist of the building of the bridge and its culmination (the completion of the bridge), together with the consequent state – that if the bridge having been built and, perhaps, the decongested traffic situation that resulted. This much would be given by an algorithm that makes use of the same kind of information about linguistic structure that Partee relies on, together with information about the structure of the nucleus, which, it could said, is just a more elaborate version of the aspectual properties of verbs and verb phrases on which we have been relying all along (for instance, in deciding which verbs and verb phrases describe events and which describe states). 'Contingency' relations might then determine the location of the new event within the old nucleus more closely, but this would be a further refinement, on top of the rougher temporal structure that can be identified on the basis of the more limited information that is delivered by grammar and lexicon. And perhaps it will then also be possible to spell out those additional (non-temporal conditions) under which such a two-step procedure for temporal relation determination reduces to the kind of narrative progression to which Partee confines herself.

But how much of an improvement is this as compared to simply abandoning the first, form-based step of the procedure and relying exclusively on the contingency-based considerations that Moens and Steedman argue we cannot do without? That depends on how much of a real constraint it is to insist that the new event e' is included within the duration of the nucleus described by some previous clause. There are two problems with the suggestion that leap at us when we try to apply it to the examples in (128) which prompted this discussion. The first is: For how long does a nucleus go on? It consists, we said following Moens and Steedman, of preparatory phase, culmination and consequent state. But when does the consequent state come to an end? That is a question for which there is no definitive answer to this day. In some cases it seems plausible to say that the consequent state has a well-defined end. Take the event that consists of you moving from A to B. In this case the culmination is your arrival at B and the consequent state is your being at B. When you leave B again, that is the end of the consequent state. And when you leave soon, then that will make the nucleus quite short. But what if you don't?. What, if you never leave B again? Does in that case the nucleus go on forever? If so, then the rule that the next event e' mentioned in a discourse after the statement that you went from A to B is temporally included within the nucleus described by the preceding sentence isn't going to tell us all that much; it only tells us that e' didn't start before the beginning of the nucleus. In general we would like some curtailment on the duration of the reference times provided by nuclei that prevents them from being unlimited in a case like this.

And when is the beginning of a nucleus. In a way we have been told: that is the beginning of the preparatory process. For some cases, like that of your moving from A to B, that may seem clear enough. The preparatory process is that of you moving towards B and that process starts the moment you set off from A. But in other cases the matter would seem to be less straightforward. Take the one that got us into this discussion, of the council building the bridge. One assessment that has been made on relation to (128.a) is that its main cause event – that of the architect drawing up the plan – *preceded* the event of building the bridge described in the *when*-clause. If that were the final word then the current proposal, according to which the main clause event must be included within the nucleus described by the *when*-clause, would be wrong for this case. But arguably this isn't the right verdict. It is based on the assumption that the building only started with the physical work that culminated in the finished bridge. But is it really all that obvious that the building of the bridge started only then? Wasn't the drawing up of the plans part of it as well? Or, for that matter, the debates of the Council that led to the commission for those plans from the architect that was chosen? If we allow for the possibility that the nucleus began at such an earlier time, then the event of the architect drawing up the plans does lie within it and there is no conflict with the rule we proposed. But now we have a new problem on our hands, which is how far back a nucleus can be assumed to extend. If we cannot come up with clear constraints on where nuclei can start, then our proposal is in danger of telling us nothing at all: not only could the nucleus on forever, it could also have begun arbitrarily far before the time of its culmination. For an event e' to be included within the duration of such a nucleus is to say very little indeed.

To sum up, the proposal of identifying reference times with the nuclei of *when*-clauses or preceding main clauses does remove the conflicts between Partee's processing principles and examples like those in (128.a) and (128.b). But when we try to spell it out in detail, we find that it is in danger of becoming wrong in its turn or else of evaporating into vacuity, in which case all the real work of temporal location will have to be done during the second stage of the procedure and the burden will fall squarely on the contingency-based principles by which this second stage of processing is governed.

This issue is the central one that dominates and motivates not only much of 'Temporal Ontology and Temporal Reference', but also the next two papers, by Webber and by Lascarides and Asher.

. The remaining comments on Moens and Steedman focus on what they have to say about (i) progressives, (ii) perfects and (iii) temporal adverbs.

4.6.3 Progressives

Much of what might be added to what Moens and Steedman have to say on this topic has already been said. In particular, the examples in (129) cover familiar ground. (129.a,b,c) exemplify the possibility of applying the progressive to activity phrases, accomplishment phrases and achievement phrases. In the first part of these comments on Moens and Steedman we noted that the details of the interpretation process may be different for each of these three cases; and to the extent that I am able to judge that suggestion is in keeping with Moens and Steedman's views on the possibility of iterated aspect modification, corresponding to chains of two or more arcs in their diagram. Furthermore, (129.d) is another testimony to the possibility of applying the progressive to what is basically a stative phrase, but only if and when this phrase is first coerced into an activity description. (129.e) confirms that when such a state-to-activity coercion is not possible, then applying the progressive goes awry.

(129) a. The president is speaking.

b. The president is giving a speech.

c. The president is reaching the end of his speech.

d. The president is being obscure.

e. The president is knowing what he wants to say. (*)

But the examples in (130) add a further twist. The salient interpretation of (130.a), when it is offered without context as it is here, is that at some unspecified time Roger was engaged in the action of running a mile. And when followed by the *but*-clause of (130.b), this interpretation becomes inescapable because the elliptical phrase *gave up* requires the interpreter to recover some plan or project that Roger was engaged in at the time and that he could have given up on, and the first clause of (130.b) can supply this only when it is interpreted in the way just described. But nevertheless, (130.d) shows that another interpretation of *was running a mile* is possible as well. The natural interpretation of *was running a mile* here is that last week Roger was being able to run a mile (but, it is implied, not more), but that between last week and this week his ability greatly improved so that he is now able to run three miles in one go.

(130) a. Roger was running a mile.

b. Roger was running a mile, but he gave up after two laps.

c. Roger was running a mile yesterday.

d. Roger was running a mile last week. This week he is up to three.

e. Quite a few of my friends run one mile. But only Roger runs three.

This interpretation of the past progressive in (130.d) is closely related to the one that is strongly suggested by the sentence in (130.e) for the non-progressive forms *run* and *runs*. Evidently what is meant here is that many of the speaker's friends can run and/or regularly do run on mile, but that Roger can go (and/or regularly does go) up to three. These interpretations of the phrase *run a mile* are on a par with what we said about *beats* in the two-sentence discourse 'Pedro owns a donkey. He beats it.', a point that was taken up in our discussion of the Moens-Steedman diagram. There we noted that this transition – from episodic to habitual/dispositional reading – is something that (at least in English) doesn't involve change of form.

(130.d) indicates that the episodic-to-habitual/dispositional interpretation is also possible when the progressive is involved, and this complicates the story about the progressive and the Simple Present that we told as part of our discussion of the diagram. We seem committed at this point to the view that re-interpretation may involve a number of aspectual transitions and that these have to happen in a certain order.

The most plausible account that I can think of goes as follows. The first transition is from the basic episodic reading of the accomplishment phrase *run a mile* to its habitual/dispositional reading. I assume that it is to the habitual/dispositional reading that the progressive is then applied. But for this to be credible something more needs to be said. In our earlier discussion it was suggested that when an event verb or verb phrase is given a habitual or dispositional reading it functions grammatically as a state description and that this explains why such verbs or verb phrases can appear in the Simple Present (in cases where the Simple Present has its standard use) as long as they are given such a reading. If this is right, and it is also true that in (130.d) the progressive is applied after re-interpretation to a habitual/dispositional reading has already occurred, then an intermediate step will be needed that transforms the state description that results from the re-interpretation the verb or verb phrase as a description of a habit or disposition into a kind of process or activity description, which fits the input requirements of the progressive.

But what is this transformation from a habit/disposition description into a process description? What process is being described by the transformed phrase? One effect of state-to-activity coercion that is often mentioned is agent control: 'When we say of Bill that he 'is being obnoxious' we mean not that he is an obnoxious individual, but that right now he is behaving in an obnoxious fashion, something that he could stop of he wanted to and could be asked to put a stop to. Volition doesn't seem the relevant factor in the transformation we are considering – from the state of being in the habit, or of having the ability, to run a certain number of miles to an activity derived from it. But Moens and Steedman's own example (129.d) of state-to-activity coercion triggered by the progressive can be seen as a kind of bridge between the two cases. 'The president is being obscure.' can be said on an occasion where the president consciously expresses himself in an obscure manner in order to obfuscate some facts that are embarrassing to his or his Government. But it can also be used to express that he is making statements that are obscure whether he wants them to be or not: obscurity is ascribed to his current performance, rather than to the president himself, as along-standing property. In fact, for all I can tell, this may be the more salient interpretation.

I take it that the force of state-to-activity coercion of the habitual/dispositional interpretation of *run a mile* is similar to this second interpretation of (129.d) in that agent control plays no part. Rather, the transformation is from what would normally be taken to be a long term property of the subject to a more limited, changeable and short term property – one that a subject can have at some time without anything following from that about it having the property at times at some considerable remove.

This assessment of the semantics of state-to-activity coercion in the case at hand seems to fit our intuitions about what (130.d) means, and I will leave the matter at that. But note well that if this story is correct, then the interpretation of a progressive form like that in (130.d) involves no less than three successive aspectual transformation operations: (i) reinterpretation of *run a mile* as description of a habit or disposition; (ii) reinterpretation of this habit/ disposition description into a corresponding activity description; (iii) transforming this activity description into a state description through direct application of the progressive. This may seem a rather baroque account for something that feels like such a simple and natural matter to English speakers. That isn't much of an argument for a simpler account, for linguistic processing often proves to be a complicated and intricate affair when looked at closely and in the context of a wider range of phenomena all of which a speaker of the language must be able to handle and for all of which he presumably makes use of tools from a single kit. But in the present instance a slight modification of our story makes it look more straightforward and this modification may therefore be the better story. The modification simply consists in assuming that the result of reinterpretation of an event description as description of a habit or disposition is neutral on the question whether the resulting description is of the aspect type 'state' or of the type 'process/activity'. When placed in a given linguistic context it will then take on whichever of these two options the context requires; and when the context doesn't impose such a requirement, then 'the state-or-process' question simply presumably isn't resolved, but with no loss to semantic content. Thus when the a habit or disposition is used in the Simple Present, then its capacity to function as a state description is activated and when it is used as input to the progressive, then what is activated is the process/activity option. For the progressives in (130.d) this means that just two aspectual transformations are involved, one covert and one overt, whereas the Simple Presents in (130.e) require just a single covert operation.

Nothing was said so far about (130.c). This sentence is interesting in that it seems to allow for both an episodic interpretation (at some particular time yesterday, Roger was in the process of running a mile') and for a habitual/dispositional reading ('yesterday one mile was as far as he did/could run without pausing'). The second interpretation may not have occurred to us until we noted the obvious similarity with (130.d), but once we have seen that the sentence can be interpreted in this way it seems to express this reading perfectly naturally.

4.6.4 Perfects

The central observation that Moens and Steedman make about the Perfect is that the Perfect is an aspectual operator, which transforms an eventuality description into a description of the consequent states of the eventualities described by the untransformed description. The result of this transformation can then be put into the Present, Past or Future tense just as verbal descriptions that have not been subjected to the Perfect transformation. (In this regard the situation is analogous to the possibility of putting the different tenses onto progressivized descriptions just as they can be put on non-progressivized descriptions. In this they were (in my opinion) clearly right, and at the time when they made this claim, it was by no means received wisdom that this was the right way to look at perfects. Rather, the more prevalent opinion was that the perfect tense forms of English are genuine tense forms, fully on a par with non-perfect forms such as the Simple Present and the Simple Past, much as is implied by Reichenbach's analysis of 'perfect' and 'non-perfect tense forms'.)

But while they were right, and importantly right, with this basic claim, Moens and Steedman seem to have been over-optimistic in how much mileage can be got out of this claim all on its own. They point in particular to two notorious puzzles about the English Present Perfect illustrated in (131) and (133).

(131) shows that the Present Perfect is incompatible with an adverb like *yesterday* which denotes a period that is wholly in the past of the utterance time (see (131.b)). In this respect, the Present Perfect differs from the Simple Past, which is compatible with adverbs like *yesterday* ((131.a)); and *yesterday* is different from *today* in that the latter (whose denotation does include the utterance time) is compatible with the Present Perfect ((131.c)), as well as with the Simple Past ((131.d)).

(131) a. They married yesterday.

b. They have married yesterday. (*)

c. They have married today.

d. They married today.

In the years that have passed since 'Temporal Ontology and Temporal Reference' appeared it has become clear that a conclusive account of the prohibition illustrated in (131.b) is not easy to come by. Such an account must bring into play quite a few assumptions about the way in which tenses and aspectual operators interact with temporal adverbs such as *yesterday* and *today*, it must explain what it is about the English Perfect as opposed to the Perfects found in many other languages and it also has to explain why in English the prohibition is specific to the *Present* Perfect. (It doesn't seem to hold for the Past Perfect, for instance, as shown by (132).)

(132) In late afternoon they landed on Gran Canaria for their honeymoon. They had married the day before.

I believe a much more satisfactory account of the prohibition shown in (131.b) can now be given. That the Perfect is an aspect operator that transforms eventuality descriptions into descriptions of their consequent state plays its part in that account, but only as one of several ingredients. But the story is complex and that this is not the place to tell it.

The puzzle presented by the contrast between (133.a) and (133.b) is also more complex than Moens and Steedman seem to make it out to be, even if here too their view of Perfects as consequent state descriptions must presumably part of any viable explanation.

(133) a. Einstein has visited Princeton. (*)

b. Princeton has been visited by Einstein.

c. Einstein is under six feet tall. (*)

d. Princeton is in New Jersey.

e. Einstein visited Princeton.

f. Princeton was visited by Einstein.

The contrast between the infelicitous (133.a) and the felicitous (133.b) is surprising given that for so many purposes turning an active sentence (a sentence in the active mood) in the corresponding passive leaves its semantics unaffected. That there is the difference which (133.a) and (133.b) so undeniably display clearly has to do with the fact that at the present time, at which we take the sentences to be uttered or displayed, Princeton still exists, but Einstein does not, in the sense that he is no longer alive. Furthermore, the contrast shows that the verb phrase of a sentence must be construed as a predicate of the subject at the time indicated by tense. When sentences are in the Simple Present, as in (133.c) and (133.d), then this time is the time of utterance (or the time of presentation) and indeed we get the same kind of contrast as between (133.a) and (133.b). With past tenses, on the other hand, the contrast disappears, as shown by (133.e,f), both of which are felicitous.

Comparison of (133.a,b) with (133.c,d) and with (133.e,f) on the other can be seen as a further confirmation that (133.a,b) are present tense sentences in a sense, and thus that the Present Perfect is a present tense of sorts. (Though a Reichenbachian might respond by pointing out that the relevant question is whether the Reference time does or does not coincide with the Speech time.) But even if we accept this (as I have already said we should in any case), there are still a number of matters that need clearing up. First, we need an account why it is that the sentences of English, or at last some of them, involve predicating the verb phrase of the grammatical subject at the time indicated by tense. Second, like the previous puzzle the Einstein-Princeton puzzle is specific to the present tense, which needs an explanation as well. And finally, we need an account for why it is that not all Simple Present and Present Perfect Sentences with *Einstein* as grammatical subject aren't infelicitous. The sentences in (134), for instance are perfectly acceptable.

Are we to conclude that Einstein is still with us as scientific icon, if no longer as a human being of flesh and blood?

(134) a. Einstein has had as profound an influence on the development of modern physics as any physicist since the end of the nineteenth century.

b. Einstein is the most famous of all modern scientists. (*)

Once again, therefore, it seems right that the Moens-Steedman treatment of the Perfect can be an important piece in the puzzle that (133) presents, and for all we know an essential one; but it is just one of several pieces. I am not sure that at this point in time (2013) all the pieces have been identified and fitted properly together. But from what little Moens and Steedman say about this puzzle we may probably conclude that they didn't have all the pieces when they wrote 'Temporal Ontology and Temporal Reference'.

4.6.5 Adverbials

The last part of these comments on Moens and Steedman is concerned with what they have to say about temporal adverbials. The interaction between tenses aspect operators and the different types of adverbs and complex adverbials that affect either or both of tense and aspect is a topic with many sides and many complexities. None of the papers that we have looked at so far – not for that matter those that are still to come – do more than reveal glimpses of these complexities. 'Temporal Ontology and Temporal Reference' is no exception to this, and it would be unreasonable to try and go much beyond what thy have to offer on this topic.

But there is one point that deserves to be taken up. It concerns the adverbials of the types *for an hour* and *in an hour*. We have seen that Vendler identified these adverbial types to test verbs and verb phrases for membership in his aspectual categories. We already noted in our comments on Vendler that what is missing in his account is that when a verb phrase is compatible with either of these adverbial types, i.e. when they can be felicitously combined, then the combination is itself a phrase with aspectual properties (and thus one for which the question of Vendler-like classification comes up as well. Vendler was not alert to this, or simply not interested, because his primary concern was the classification of verbs. But for Moens and Steedman this is a natural and important question, as for them aspectual properties are an important feature of semantics from the lexical constituents of a sentence all the way up to the sentence itself. We will see the implications of this different perspective especially in connection with the last example of these comments to Moens and Steedman.

But the first point to be commented on here is one much closer to Vendler's program. Vendler notes that many verbs cannot be assigned to just one of his four categories: such verbs are 'cross-categorial' in that they sometimes behave as if they belonged to one category and in other situations as if they belonged to another. Now that we have seen more clearly how verbal descriptions can be transformed into aspectually different ones (in accordance, roughly, with Moens and Steedman's transition diagram), it is possible to say something more about at least some of the cases of 'aspectual multi-functionality' to which Vendler draws attention. In particular, aspectual coercion an occur under the influence of the 'test adverbials' *for x amount of time* and *in x amount of time* (which, by the way, entails that we have to be much more careful with the application of these tests than is directly apparent from the applications Vendler shows, although he cautions against oversimplifying the issues connected with application of his tests).

(135), (136) and (137) give examples in which a verb phrase can be combined felicitously with a *for an hour* type adverbial. In (135) the verb phrase is *work in the garden*. Both of these are activity verb phrases, so according to Vendler they should be combinable with a *for*-adverbial like *for five hours*, and as (135) indicates, they are indeed. But the focus now is not on what *for*-adverbials can combine with, but on the properties of what results when they can. One of these properties shows up in the contrast between (135.b) and (135.d). The oddity that Moens and Steedman perceive in (135.b) has to do with the fact that the perfect wants a description of culminating events as input – it wants such inputs because the events they describe supply the consequent states that are described by its output descriptions, those that result when the perfect is applied. *work in the garden* squarely fits the category of activity phrases, those which describe non-culminating eventualities. This accounts for why (135.b) is jarring, unless the context allows us to coerce the phrase into the description of culminating events, such as that of doing one's daily spell of garden work. (The Simple Past tense sentence (135.a) is not subject to such a restriction. The Simple Past simply serves to locate some event of John working in the garden somewhere in the past of *n*.) Combining *work in the garden* with *for five hours* turns this phrase into the description of activities with a well-defined end point, reached five hours after onset. This description is an unproblematic input to the perfect, with the effect that (135.d) is straightforwardly felicitous, no less than it Simple Past alternate (135.c).

(135) a. John worked in the garden.

b. John has worked in the garden. #

c. John worked in the garden for five hours.

d. John has worked in the garden for five hours.

(136) and (137) give examples to show the powers of coercion and the semantic flexibility that this adds to the language. *play the sonata* is intuitively an accomplishment phrase: it describes events that reach their culmination when the sonata reaches its end. According to a simple application of Vendlerian doctrine combination with a *for*-adverbial should therefore be impossible. But that it is possible doesn't prove Vendler wrong; it only shows that, in accordance with his general words of caution, the accmplishment phrase *play the sonata* can be (re-)interpreted and thereby made to fit another of his categories for which combination with *for*-adverbials is possible according to his principles.

Moreover, (136.a) and (136.b) show that more than one reinterpretation is possible in principle and that the one the interpreter opts for depends on additional considerations. In (136.a), in which *for a few minutes* denotes an amount of time considerably shorter than a typical sonata from the classical repertoire, coercion is to the activity of being on the process of playing the sonata from beginning to end. In (136.b), where the amount of time – eight hours – exceeds typical sonata length by a considerable margin, the coercion is more in the spirit of iterated complete performance events. (Though perhaps the more correct image here is that of the 'universal grinder' which we discussed as part of our struggles with Bach's 'Algebra of Events': Just as we can say 'There is apple in the salad.' to describe a situation in which the salad contains bits of some part of a single apple, or one in which the salad contains the bits from grinding one complete apple, ir one in which, say, two and a half apples went into the salad, the amount of complete sonata playing that can be described by the accomplishment-to-non-telic-eventuality coercion involved in (136.a) and (136.b) could be a fraction of one complete performance, one or several compete performances, but also some fraction greater than one, consisting perhaps of two complete performances plus a playing of part of the first movement. And that isn't all. Just as the apple-containing salad may contain bits of many apples without there being an apple of which it contains all the bits, the coerced interpretation of *play the sonata* can be used to describe a long sequence if snippets, with long sequences of repeated playing of the same hard passages that can be so trying for those within earshot.)

(136.c) is another example of the same sort. Here the verb phrase in question, *arrive late for work* is presumably an achievement. Achievements too are unsuitable for combination with a *for*-adverbial. So here too reinterpretation is needed, and possible. The reinterpreted *arrive late for work* functions as a description of sequences of successive achievements described by the basic interpretation of *arrive late for work*. This is again a non-telic description – it doesn't put an upper limit on the number of iterations, so that any sequence satisfying it could in principle be extended to a longer sequence that satisfies it as well. So *for several days* should be combinable with it, and it is.

(136) a. Sue played the sonata for a few minutes.

b. Sue played the sonata for about eight hours.

c. John arrived late at work for several days.

(137) gives some further illustrations of the same phenomenon. The felicity of (137.a) is predicted by the theoretical commitments already made: *be winning the race* is of the type 'state description' and thus combinable with *for the first few minutes*. That the same content cannot be expressed by the non-progressive form is also in keeping with earlier commitments. But it is nevertheless worth noting the apparent contrast between this case and the felicitous uses of the simple form in the superficially similar sentences in (136). It appears that morphologically unmarked non-telic reinterpretations of achievements (such as *win the race* or *arrive late for work*) is only possible in an iterative sense: the reinterpreted phrase serves as description of iterations of the kinds of event that are described by the phrase qua achievement phrase. Such an interpretation is blocked when the event described by the achievement phrase is unrepeatable, like the winning of some particular race: Once you have won that race, that's it, for that particular race cannot be run

again. Note that (137.c) confirms this assessment rather than refuting it, for the noun phrase *the Monaco Grand Prix* is the name of a generic event – a race that is run each year again and that you an therefore win repeatedly. That's why (137.c) is felicitous while (137.b) is not.

That (137.a) is felicitous as well needs a further comment. Progressives of accomplishment and achievement phrases serve to describe bits of activity that are performed with the intention to turn it into a completed event, of the kind described by the non-progressive form, or that raise the expectation that they will develop into such a completion. This future-oriented, prospective dimension to the meaning of such progressives is missing from the morphologically unmarked 'universal grinder' type of interpretation that is presumably involved in (136.a,b) and arguably also in (136.c) and (137.c). (The future-oriented dimension of progressives like that in (137.a) (and most of our earlier examples of the progressive) has ben prominent in the semantic literature on the progressive. I do not go into that literature here.)

Nevertheless, there are many instances where the two kinds of non-telic reinterpretation – the morphologically unmarked habitual/dispositional/'universal grinder' type and the future-oriented type that is overtly marked by the progressive – are hard to tell apart. This appears to be the case in particular when the reinterpreted phrase is originally an accomplishment. An illustration is furnished by the first two sentences of (136). When we replace in these sentences *played the sonata* by *was playing the sonata* the results are again felicitous, and there seems to be no appreciable difference in meaning. The correct diagnosis of why this should be so escapes me, but I am reasonably confident that one could be found.

(137.d) exemplifies a case of adjustment to felicitous combination with a *for*-adverbial that is of a kind we haven't so far encountered. Intuitively it is clear that what *for a few minutes* applies to semantically is the consequent state of the achievement phrase *leave the room*. But in this case preparing the ground for this does not take the form of reinterpreting the achievement phrase into one that describes the consequent states of events described by it in its basic interpretation. Rather, the achievement phrase seems to make one part of the nucleus of the events it describes in its basic interpretation – viz. the consequent state – available for modification by the *for*-adverb, but nevertheless retains its status as event description, but now of events whose consequent states have the additional property (that of lasting a few minutes).

I have no more to say about this kind of combination pattern, which does not consist in (i) reinterpretation to a type of phrase that satisfies the constraints imposed by a certain operation and then (ii) applying the operation, but which, rather, makes a certain part of the eventuality complexes it is used to describe available as input to the given operation but then absorbs the result of applying the operation of that part without yielding its basic aspectual profile. But it is probably good to be alert that this kind of 'coercion pattern' is found as well.

(137) a. Red Rum was winning the race for the first few minutes.

b. Red Rum won the race for the first few minutes. (??)

c. Niki Lauda won the Monaco Grand Prix for several years.

d. John left the room for a few minutes.

The sentences in (138) are examples of aspect coercion to fit the requirements of *in*-adverbials. In (138.a) *in two hours* is evidently used to make a claim about the duration of the ascent. So we get a statement that comes to much the same thing as (138.b). But whereas in (138.b) the adverb *in two hours* can be applied directly to the event described by the accomplishment phrase *climb the mountain*, application to the achievement phrase *reach the top* in (138.a) requires that we first expand the profile of the described event complex with some suitable preparatory phase; the actual process of climbing the mountain whose top is said to have been reached is the obvious preparatory phase to co-opt for this purpose.

(138.c) presents a more complicated case. First, what the sentence describes is that the *beginning* of the state of John being ready occurred no more than five minutes. So let us suppose that it is part of the interpretation of the sentence that the verb phrase *be ready* is reinterpreted in this 'inchoative' manner. That turns the phrase into an achievement phrase and thus puts the interpretation of (138.c) on the same footing as (138.a). But as we saw in connection with (138.a), something more needs to happen before the *in*-adverbial can be applied in the way it wants to: we need to supplement the mere culmination described by the achievement phrase – here: the beginning of the state of John's being ready – with a preparatory phase to which the *in*-adverbial can then actually be applied. But in the case before us there is, unlike in the case of (138.a), nothing in the content of the achievement phrase that suggests what this 'preparatory phase' could be. Rather, it is the context in which (138.c) is used that should tell us what to take as 'preparatory phase': The context should suggest a certain time at which John's being ready is wantd or expected, and the question is then how long it took from that point until this state was reached.

Exactly how such 'times from which to start counting' are inferred from the context is an interesting question in its own right, but it is on that doesn't quite belong here. So all that I want to retain from this discussion of (138.c) is that once again we see a (two element) chain of readjustments of the input phrase that are needed before the operation that triggers those adjustments – in this case: combination with an *in*-adverbial – can be carried out.

(138.d) is in essence an instance of the same pattern as (138.c). The only difference is that now the state description that has to be coerced into some sort of accomplishment is the consequent state of an event of Harry spilling his coffee, which is described by the perfect phrase *have spilled one's coffee*. If combining this perfect with *in less than three minutes* follows the same principles that we just proposed for the interpretation of (138.c), then this state description is first subjected to inchoative reinterpretation and then the context has to provide a time which can serve as starting point for the 'preparatory phase'.

It is interesting to compare (138.d) with (138.e), which may at first sight look so much like it that one might think the differences may be ignored. In a way this is of course true, for (138.d) and (138.e) seem to mean pretty much the same thing. But in fact, when we look more closely we see that the principles involved in the interpretation of (138.e) are quite different from the ones that were just proposed for (138.d), and that they are actually a good deal simpler. First, it seems that the event described by the verb phrase of (138.e), *spill one's coffee*, should be seen as what Moens and Steedman refer to as a 'point' - an event that can be thought of as happening at some particular time, such as the time three minutes from a given time t (or some time less than three minutes from t). The only further step that is then needed for a full interpretation of (138.e) is then to identify t in the given context.

(Moens and Steedman annotate (138.f) with two question marks, signalling some considerable degree of infelicity. It is true that the sentence does sound a little weird, But given that it has the exact same structure as (138.e), the only reason there could be for the difference is that fifteen minutes is an oddly long time for the kind of thing the sentence is trying to say. The sentence seems to improve when *in* is replaced by *within* or the Simple Past is replaced by a Past Perfect, as in (138.c), or both. As it is I have no good story why these changes do improve the sentence.)

(138) a. Laura reached the top in two hours.

b. Laura climbed the mountain in two hours.

c. John was ready in five minutes.

d. In less than three minutes Harry had accidentally spilled his coffee.

e. In less than three minutes Harry accidentally spilled his coffee.

f. In fifteen minutes Harry accidentally spilled his coffee. (??)

In (139) we see a verb phrase that is modified by both an *in*- and a *for*-adverbial. With a sentence such as this we have left the original concerns of Vendler well behind us. *play the Minute Waltz* is an accomplishment phrase to which the *in*-adverbial *in less than sixty seconds* can be applied without any preliminary adjustments. But in order to then subject the resulting phrase *play the Minute Waltz in less than sixty seconds* to combination with a *for*-adverbial adjustment is necessary. This adjustment is once again of the dispositional/iterative kind. Thus adjusting and then combining with *for more than an hour* gives us once again a kind of accomplishment phrase, and this phrase is now in its turn the input to the aspect operator *it took me two days*. This operation is again unproblematic because an accomplishment phrase is just the kind of input that *it took me two days* wants.

(139) It took me two days to play the 'Minute Waltz' in less than sixty seconds for more than an hour.

Before we move to the next paper, Webber's 'Tense as Discourse Anaphor', first a brief interlude about temporal relations in discourse and the notion of 'Reference Time'

4.6.6 Tenses, Temporal Relations and Reference Times

The following discourse was the subject of one of the Problems on the 2nd Homework assignment:

(140) We reached the top of the mountain at a quarter to one.

We had got up at 4.30. We had had an unhurried breakfast. Then we had got together all that we would need and had taken the first bus to the bottom of the trail. The climb had taken us just over six hours.

After we had had a short rest and enjoyed the spectacular view, we set off on the way down.

This last example shows that the original notion of Reference time as it was introduced by Reichenbach serves a different function from that of the Reference Points of Hinrichs and Partee. In (140). The temporal structure of the part consisting of the first five Past Perfect clauses is of a kind that closely resembles the Simple Past Tense narrative progressions that much of their discussions focus on. If we take this analogy seriously, then we should assume that each of the first four Past Perfect clauses introduces a reference point situated 'immediately' after the event described by the clause an which serves to locate the event described by the clause following it. But at the same time each of the Past Perfect clauses requires, according to Reichenbach, a Reference time that is situated in the past of the Speech time and that locates the eventuality described by the clause as situated within its own past. Now let us focus on the Past Perfect clauses two to five. According to the assumption we just made in order to capture the analogy with the proposals of Hinrichs and Partee for each of these the event e_i it describes is located by the reference point r_{i-1} introduced by the immediately preceding Past Perfect clause, and the temporal relation between the two of them is that of inclusion: $e_i \subseteq r_{i-1}$. Evidently, the Reference time in the sense of Reichenbach cannot be the same as r_{i-1}, for e_i must be situated *before* the Reichenbachian Reference time and that is *not* the relation between e_i and r_{i-1}. In any case it is intuitively clear what the Reichenbachian Reference time should be in this case: for each of the first five Past Perfect clauses it is the (past) tim of the set of people indicated by *we* reaching the top.

What this example teaches us is that when we combine the insights of Reichenbach and those of Hinrichs-Partee, then at least for the Past Perfect clauses 2 to 5 in (140) their interpretation requires two distinct times, Reichenbach's Reference time *and* a reference point in the sense of Hinrichs-Partee, and not just one time.

(In DRT-based work on tense and aspect the two roles are often kept distinct by referring to Reichenbach's Reference time as *Temporal Perspective Point* or *TP-pt* and to the reference times of Hinrichs-Partee as *Reference Points* or *R-pts*. I will switch to this terminology at some later point in this course.) Once we have seen the need for distinguishing between Temporal Perspective Point and Reference Point in connection with 'extended flashbacks' the question naturally arises whether this need – for both types of temporal points – shouldn't arise in connection with all tenses. The answer to that question clauses can only be 'yes'. But exactly what form this answer takes in its details is another matter, to which we will have to return later.

4.7 Webber: 'Tense as Discourse Anaphor'

Webber's paper can be seen as a link, or bridge, between the last paper discussed, that of Moens and Steedman, and the two papers of Partee discussed before it. With Moens and Steedman Webber shares the leading insight that temporal relations between events in discourse are often inferable only as part of a package that also includes non-temporal relations, and where it is often the latter that carry the primary load: once those non-temporal relations have been grasped, the temporal relations follow, as a kind of secondary effect. The insight shared with Partee is that tenses play the part of discourse anaphors, which seek their antecedents in other sentences belonging to the same discourse. And like Partee Webber points in this connection to striking parallels between amphora in the temporal and anaphora in the nominal domain.

But on this last point there is nevertheless an important difference. Partee emphasizes the similarities between tenses and pronouns. For Webber the parallel is that between the anaphoric behavior of tenses and the behavior of anaphoric *noun phrases* – not just anaphoric pronouns but also anaphoric definite descriptions. Noun phrase anaphora in this more general sense is a more diversified phenomenon than singular pronoun anaphora and for that reason it provides, Webber claims, a better counterfoil for the different forms that tense anaphora can take.

Since the analogy between nominal and temporal anaphora plays such a central part in what Webber wants to say about the mechanisms of temporal reference, the first part of her paper is devoted to those features whose analogues will be prominent in what follows about the anaphoric potential of the tenses. The central starting observation of this first part of the paper is that while the interpretation of singular anaphoric pronouns typically involves referential identity of the pronoun with its antecedent, this is often true only in a modified sense for plural pronouns and it does not need to be true at all for anaphoric definite descriptions. These points are illustrated by the examples in (141).

(141) a. Wendy gave Eliot a T-shirt for Christmas. Unfortunately *it* had the logo 'You ate it, Ralph'.

b. Wendy gave each boy a T-shirt. *They* each had a different logo on the front.

c. The vice president must be over 35 years old *He or she* must be able to count.

d. The dachshund down the block bit me yesterday. *They* are really vicious beasts.

e. A bus came around the corner. I signaled *the driver* to stop.

f. The driver stopped the bus when a passenger began to sing 'Aida'.

g. The driver stopped the bus when the passengers began to sing 'Aida'.

The *it* of (141.a) is a classical example of a pronoun that is anaphoric to a referential antecedent (*a T-shirt*) and 'coreferential' with that antecedent: the pronoun 'refers' to whatever its antecedent 'refers' to. (141.b-d) are variations of this 'coreference' relation: The *they* of (141.b) serves to refer to the respective T-shirts of the boys who got one from Wendy. An exact account of cases like this is complicated because of the way in which the quantification expressed by the first sentence – that each boy got a T-shirt – is extended to the second sentence, with *they* referring to distribution over the same set of T-shirts that are involved in the semantics of the first sentence (viz. the T-shirts given to the respective boys). But it is nevertheless true of this example as well that for each of the boys the second sentence talks about the same T-shirt as the first sentence. (141.c) is like (141.b) in that the first sentence is talking about the different vice-presidents that are possible in different possible worlds. Of each of those possible vice-presidents (male or female) the second sentence says what properties this vice-president must have if the rule expressed in (141.c) is to be obeyed. And (141.d), which like (141.b) has the plural pronoun *they*, refers to dachshunds in general. Here the most plausible account is one that analyzes *they* as referring to the kind denoted by the noun *dachshund*.

As we have indicated, the analysis of each of these four cases, they all involve referential identity between the anaphoric noun phrase (the pronoun) and its antecedent. For (141.d-f) this is different. In connection with (141.e) there is a sense in which *the driver* in the second sentence is 'anaphoric' to *the bus* in the first: it is only in relation to *the bus* that *the driver* can be given its intended interpretation. But the two noun phrases are clearly *not* coreferential: no driver is ever identical with the bus she or he is driving. Anaphoric relations like that between *the driver* and *the bus* in this example are now most often referred to as *bridging* (after Clark (1997)). A bridging relation between an anaphoric noun phrase and its antecedent is a relation of non-identity which is indicated by the descriptive content of the anaphoric phrase. In many cases (though by no means all) the head noun of a bridging anaphor is a relational noun and in that case the bridging relation is given by the noun: the referent of the anaphoric noun phrase stands in the relation expressed by its head noun to the referent of the antecedent. (That is, in the example before us the referent of *the driver* stands in the 'driver-of' relation to the referent of *the bus*.)

Much the same is true for the 'bridging anaphors' *a passenger* in (141.f) and *the passengers* in (141.g). The point of (141.f) is that because bridging doesn't involve referential identity, the bridging anaphor need not be a definite noun phrase. It is really the head noun *passenger*, you might say, that is the 'anaphoric' element in *a passenger* in (141.f) – and likewise in *the passengers* in (141.g), and the same thing can also be said about the noun *driver* in (141.e). *passenger* in (141.f) and (141.g) and *driver* in (141.e) select the sets of entities that stand in the relation they express to the referent of the antecedent phrase. In the case of *driver* in (141.e) this set is a singleton (i.e. has a single member) and for that reason the only suitable bridging phrase that can be built from it is the singular definite description *the driver*. With *passenger* in (141.f) and (141.g) this is different. The presumption implied in these two examples is that the set of of passengers on the bus referred to in the first sentence is not a singleton. That allows the formation of a different and wider range of bridging anaphors. The singular indefinite *a passenger* is one of them and the definite plural *the passengers* a second one. But other bridging phrases would have been possible too, such as *some passengers*, *most passengers*, *more than one passenger*, *five passengers*, *at most seven passengers*, *at least four passengers* and so on.

I have discussed these cases in somewhat more detail than Webber does herself, but only to bring out what I take to be her central point: that anaphoric relations in the nominal domain can take a substantial number of different forms, and that only some involve 'referential identity' while many others do not, involving various kinds of bridging relations instead. Arguably Webber's most important point here is that it is this *variety* of possible anaphoric relations in the nominal domain that is essential to a correct understanding of the parallels between nominal and temporal anaphora.

With an eye on these parallels Webber formalizes nominal anaphora in a way that brings out the different reference relations they involve. This formalization presupposes a notion of Discourse Model – the listener's model of the discourse that he has processed at any given point – in which the referents of noun phrases belonging to the part of the discourse that has already been processed are represented in a manner that fits the general form of such Discourse Models. Given what we have been assuming in these notes about discourse representations – i.e,. that they take the form of DRSs in the sense of Partee's 'nominal and Temporal Anaphora' and of the comments on that paper in these Notes – we may as well persist with that assumption here and assume, for the purpose of the present discussion, that Webber's discourse Models can be identified with DRSs (even if Webber herself did not think of her Discourse Models in these terms). And if we do this, then the representations of the 'referents' of earlier noun phrases that are part of the Discourse Models which enter into Webber's formalization will be the discourse referents that are introduced for the antecedent noun phrases in the construction of the DRS-shaped Discourse Model.

Given these assumptions the formal representation Webber adopts for the relation between anaphoric NPs and their antecedents takes the following form: of a function '$\alpha(\mathrm{NP}_b,\mathrm{E}_a)$', where NP_b is the anaphoric noun phrase that needs to be interpreted, E_a is the dref for the anaphoric antecedent, and α itself is a function that reflects the relation between the entity referred to or described by NP_b and the entity represented by E_a. Often α reflects the identity relation, something that Webber formally expresses by writing '$\alpha(\mathrm{NP}_b,\mathrm{E}_a) = \mathrm{E}_a$'. When NP_b is a third person pronoun this 'identity function' is the rule (albeit one with some exceptions). But when NP_b is a definite description, then α is often some relation other than identity, and that is especially common when the head noun N of NP_b is relational and denotes a relation R. In these cases the application of the function α denotes one or a set of more than one entities that stand in the relation R to the referent of the antecedent noun phrase and we can make that explicitly by

denoting the function as α_{R}. It follows from what we have said that in all such cases $\alpha_{\mathrm{R}}(\mathrm{NP}_b,\mathrm{E}_a) \neq \mathrm{E}_a$.
It is worth noting, especially in view of the parallels with temporal anaphora which are the point of what Webber has to say about the nominal case, that different instances of noun phrase anaphora will between them involve a wide and open-ended range of different functions α. Assuming that what has been said here about bridging descriptions with relational nouns as head nouns is correct, then as we have seen there is for each such noun N denoting a relation R a corresponding function α_{R} that maps NP_b and E_a onto the unique entity that in the given context stands in the relation R to E_a. The total range of possible α's will include all these functions α_{R}, and much more.

The last part of Webber's paper that is devoted to nominal anaphora addresses the question how anaphoric antecedents are identified – in her formal terms: how does the interpreter identify the E_a that she needs as argument to the function *alpha* that according to Webber resolves the anaphora of NP_b? This is a topic in its own right, and has been treated as such from times preceding Webber's paper (see in particular Sidner (1979), Grosz & Sidner (1986) to the present day. (One particular direction this work took is that of *Centering Theory* (Grosz et al. (1983)), (Beaver (2004) et al (2004) among many others). The focus of Centering Theory is on what noun phrase a speaker or author should use at a given point in a discourse or text for the purpose of referring back to a given antecedent – for instance, should the choice be a pronoun or a suitable definite description? This is somewhat different from a focus on antecedent detection, but there is an obvious close connection: for example, if Centering Theory is right in saying that in order to refer back to a certain antecedent from a certain position in the text a definite description should be used rather than a pronoun and the text contains a pronoun in that position, then the pronoun must have some other antecedent than this one.)

Whatever the details of the constraints that are imposed on the range of possible antecedents for a given anaphoric noun phrase in a discourse or text, there is no question that these constraints are quite strong in that they restrict the set of possible antecedents to what is usually only a small 'local' subset of the totality of antecedents that are made available by the discourse or text as a whole. And that, as Webber is right to observe, is no less true of temporal than it is of nominal anaphora. But as she makes quite clear, the antecedent constraints on temporal anaphora are for the most part quite different from those on nominal anaphora and require their own story. So I won't go into this part of her discussion of nominal anaphora any further.

Temporal Anaphora

The theoretical foundation of Webber's approach to temporal anaphora shares with that of Partee's 'Nominal and Temporal Anaphora' the central part it attributes to Reichenbach's Reference time. But the role that reference time plays in her account is not the same. In particular, she does not assume, as do Hinrichs and Partee, that each new event sentence in a narrative introduces its own reference time, located shortly after the event the sentence describes and that, in the Hinrichs-Partee algorithm, can then serve as 'temporal antecedent' for the next sentence. Rather, for Webber it is the event itself that provides the temporal anchor for the next sentence. But it does that by anchoring the next event via *its* (i.e. that second event's) Reference time and that involves relating the Reference time to some part of the nucleus of the antecedent event. In this way a greater flexibility can be secured in the temporal relations between the events of two successive sentences, that flexibility that we have seen is needed in our discussion of Moens and Steedman.

Webber proposes to account of the interpretation of past tense sentences in discourse formally along the lines of her use of the different functions α that do a central part of the work in her proposal for a treatment of noun phrase anaphora. The corresponding functions for temporal anaphora are represented by means of the letter β. Webber assumes that each case of temporal anaphora involves three arguments: (i) the new tensed clause C_b that has to be interpreted in relation to the given Discourse Structure, (ii) the relevant discourse entity E_a and (iii) the Reference time RT_b of the new clause C_b. C_b, Webber says, enters into the application $\beta(C_b,E_a,RT_b)$ of the relevant function β in two different ways: (a) because C_b determines (via the principles of compositional semantics) the event E_b that must be temporally related in the right way to the 'antecedent event' E_a; (b) because of the Reichenbachian (ST,RT,ET)-structure that is determined by C_b's tense form. The temporal relation between E_b and E_a which should be the result of the interpretation of the tense of C_b is, on this account, the product of two factors: (i) the choice of the function β and (ii) the relation between E_b and

RT_b as determined by the tense of the new sentence. More specifically, the choice of β determines where RT_b is to situated in relation to E_a. Combining this with the information about the relation between E_b and RT_b then gives the temporal relation between E_b and E_a.

At an intuitive level there is much to be said for this kind of analysis, and by and large the examples we will look at below will confirm this. But it is difficult to make sense of the particular way in which Webber proposes to capture this analysis. In analogy with what she says about arguments and values of the α- functions in terms of which she accounts for the different types of noun phrase anaphora, she wants the values of her β- functions to be the events E_b determined by C_b, and that for different β- functions. I do not quite see how this can be made sense of. The value of $\beta(C_b,E_a,RT_b)$ should be that which tells us how E_b is temporally related to E_a. In fact, once E_a and the function β have been determined, the temporal relation between E_a and RT_b is fixed as well, and with that, via the (ST,RT,ET)-structure determined by the tense of C_b, also the temporal relation between E_a and E_b. So it seems that if we want to make the analogy with the α-based analysis of nominal anaphora as close as possible, then we should take the functions β to return relations between RT_b and E_a as values and let β operate on just these two arguments. The really hard work, and more so than in most cases of noun phrase anaphora, is to determine the choice of β and it is here that information about what kinds of events that E_a and E_b are plays a crucial part. In our discussion of Moens and Steedman we have seen that the temporal relation in which the events described by successive sentences in a discourse are understood to stand to each other depends crucially on world knowledge of how two events of the given descriptions could be related; and Webber is as much aware of this as Moens and Steedman and equally concerned to integrate such world knowledge-based factors into her account of tense anaphora.

As a matter of fact, there is only a quite small number of different β-functions that Webber distinguishes; in this respect the options for temporal anaphora are much more surveyable than those which arise in the case of nominal anaphora. In her more detailed observations about temporal anaphora Webber makes use of just three such functions, which she denotes as 'β_0', 'β_{prep}' and 'β_{conseq}'. (For some of the later examples in the paper it isn't entirely clear whether they can be handled with one of these three functions. Webber doesn't address this question – at the point where these examples are adduced, they are used to illustrate other phenomena. Without further input I see no way of addressing the question either.)

To see what the three functions β_0, β_{prep} and β_{conseq} do, it will help to look at some examples, for which the time has come in any case.

(142) a. John partied until 3 am. He came home and went to bed.

b. The elderly gentleman wrote out a check, tore it from the book and handed it to Costain.

c. John played the piano. Mary played the kazoo.

d. For an encore John played the Moonlight sonata. The opening movement he took rather tentatively, but then ...

e. John went into the florist shop. He picked out three red roses, two white ones and one pale pink.

f. John bought Mary some flowers. He picked out three red roses, two white ones and one pale pink.

All sentences and clauses in these examples are in the past tense and all of them describe events. The first two examples, (142.a) and (142.b) are of the kind that would get the right treatment in the account of Partee(1984): each next sentence or clause describes an event that is understood to have occurred shortly after the event described by the previous clause or sentence. This is also true for not true for (142.e). But it is not true for (142.c), (142.d) and (142.f). Webber wants to account for these differences in terms of different β-functions: The β-function relevant to be used in (142.a), (142.b) and (142.e) is the function β_{conseq}, whereas the function involved in (142.c), (142.d) and (142.f) is the one she calls β_{prep}. Both the names of these functions and their output are explained by the setting in which the choice between them is made. This setting is provided by an eventuality selected by the interpretation process as E_a, which in all these examples is the eventuality introduced by the immediate clause or sentence. Webber makes about these events the same fundamental assumption that is also made by Moens and Steedman, viz. that it has the profile of a *nucleus*, with preparatory phase, culmination and consequent state. The choice of β is in essence a choice between that part of the nucleus of E_a to which RT_b is to be related (either by the relation of temporal inclusion or that of temporal identity). The cases in which the event of the next sentence or clause is understood as following that of the preceding clause or sentence are those where what is selected is the consequent state of E_a, and in this case the relation determined by the selected β-fiunction β_{conseq} is that of temporal inclusion (of RT_b) within the consequent state of E_a. When the sentence or clause C_b also describes an event, E_b and the tense of the clause is one that determines the relation between RT and ET to be that of coincidence, then E_b will be interpreted as temporally included within its reference time RT_b, so that we end up with an interpretation in which E_b is temporally included within the consequent state of E_a. The choice of β_{conseq} thus produces the same net effect as the Hinrichs-Partee algorithm of Partee(1984).

The cases of (142.c), (142.d) and (142.f) require different β-functions. In the case of (142.d) and (142.f) this is the function β_{prep}, which locates RT_b, and with it, for the examples considered at this point, E_b, within the preparatory phase of E_a. In the case of (142.c) the relevant β-function is the function β_0, which identifies RT_b with a time including the duration of E_a. Here the ultimate result is that E_b and E_a occur within the same 'reference time'; a special case of which is that where they cover the exact same time, something that it is reasonable to assume for the two events of (142.c).

It is plain from this discussion that the hard, and as yet unaccounted for, part of the interpretation of the tense anaphora in these examples is the choice of β-function. This choice requires semantic representations of the two event descriptions that are provided by the two successive sentences or clauses and will as a rule depend on various kinds of linguistic and would knowledge. The formalization of how these choices are made is not among the goals of Webber's paper – not unreasonably so, since it is notoriously complex and its complexity was understood as clearly at the time when her paper was written as it is today. (With the next paper, by Lascarides and Asher, we will encounter an impressive attempt to get a formal grip on this problem, even if the application of the proposals Lascarides and Asher make also remain – inevitably – fragmentary.)

Schematically the process of tense interpretation, as Webber conceives of it and as it is illustrated by what we have just seen must happen in the case of the examples in (142), involves the following steps:

(i) identification of E_a (with its semantic representation);

(ii) identification of the (ST,RT,ET)-structure determined by the tense form of the new sentence;

(iii) determination, using the event descriptions C_a and C_b and whatever relevant information is needed, of the relevant β function;

(iv) determination of the temporal location of E_b in relation to E_a by combining the location of RT_b relative to E_a – given as output of the β function application – with the relation between RT_b and ET_b – given by the (ST,RT,ET) structure.

In principle we could pack all these different operations into a single function β, with the arguments that Webber suggests, but that is possible only if we take the choice of β-function to be part of the output. In fact, it is this choice which is the major task that needs to be accomplished by the interpretation process, once the antecedent event E_a has been determined. We can also, staying somewhat closer to the formal proposal that Webber makes, take the output of the central function to be a pair, consisting of something that specifies the choice of β-function and something that identifies the relevant relation between E_b and RT_b. This is what I will assume as part of the present attempt to reconstruct Webber's proposal. More specifically, I will assume that the outputs of the β-function are pairs consisting of (i) one of the labels '$_0$', '$_{conseq}$', '$_{prep}$' that Webber uses to distinguish between her different β-functions and (ii) something which identifies the relation between RT_b and ET_b.

What we need as inputs to a function that is capable of returning these outputs are (i) E_a (with its semantic representation) and (ii) the full finite clause or sentence S_b. The latter could arguably be broken up into two parts, (ii.1) its tense, which determines the relation between RT_b and ET_b, and (ii.2) the tense-free part, which constitutes the event description properly, without any information about where the described event occurred (or is occurring or will occur) in time. (If we assume that semantic representations are DRS computed from the types of syntactic structures that were adopted in our comments to Partee's second paper, then the event description proper provided by the sentence/clause could be identified with that part of the DRS which remains when the conditions introduced by the tense features – those associated with the T-node – are removed.)

As far as I can tell, thinking of S_b as decomposed in this way into tense (as determining (ST,RT,ET)-structure) and C_b as the tense-free description of the event E_b, is consistent with what Webber has in mind (although I do not think that this can be inferred conclusively from any of her actual formulations). But let us assume that I am right in this and thus that C_b is the tense-free description that S_b provides of the event E_b. Then we can think of the β-function as a 3-place function with as arguments (i) C_b, (ii) E_a (with its description C_a) and (iii) something that determines the relevant part of the semantics of the tense of S_b. Webber assumes that the tense is represented by the (ST,RT,ET) structure which gives its complete Reichenbachian interpretation, consisting of a temporal relation between ST and RT and one between RT and ET. But as we have already seen, all that really matters for locating E_b relative to E_a is the second of these two relations. It might be argued that for the entire temporal interpretation process that is needed for S_b the first relation is important as well. For it is only when this relation reveals the tense form as a 'past tense' – i.e. one with the relation $>$ between ST and RT – that interpretation of the tense of S_b via an anaphoric relation to the event of some earlier past tense clause is a viable option at all. It seems, however, that this information is relevant only to an earlier stage of the interpretation process, during which E_a is selected as temporal antecedent for E_b, and which must thus precede the application of the β-function; and for this reason it seems more natural to restrict the input to the third argument slot of our revised β-function to just the relation between RT_b and ET_b. We will assume that this input is specified as one of the symbols '$=$', '$>$' and '$<$', to indicate which of the three possible relations between RT_b and ET_b obtains by virtue of the tense form of S_b.

If this is how the inputs and outputs of the revised β-function are determined, we are facing a certain awkwardness: on the one hand the relation between RT_b and ET_b is one of the inputs to the function (i.e. as the specification of its third argument) and on the other it is part of the value the function that the function returns, viz. as second component of that value. So as far as this pierce of information is concerned the function does no more than to pass it on unchanged from what it receives as inputs to what it returns as output. But let us not be deterred or irritated by this oddity. Our present task, that of a viable reconstruction of the account Webber presents, is not an easy one, and oddities of this kind should be the least of our worries. More important is this: It is plain from these considerations that what really matters about the output of the new β-function is the specification of which of the particular β-functions in Webber's sense is involved in the given interpretation. This output component is, we already decided, given in the form of one of the subscripts that distinguish Webber's β-functions from each other. Since this part of the output is the result of a non-trivial computation, and will vary between the three possible values $_0$, $_{conseq}$, $_{prep}$ for different combinations of C_b and E_a, this information evidently cannot be a fixed part of the function itself. That is, our new β-function has to neutral between what distinguishes β_0, β_{prep}, β_{conseq}, etc from each other. Let us make this explicit by representing our new function as '$\beta_?$'. The net effect of applying this function to a set of arguments is then that of answering the question implicit in the subscript '?'.

To give an impression of what this revised formalization looks like in actual cases, let us apply it to some of the examples in (142), viz. (142.a), (142.c) and (142.e). First consider the interpretation of the second clause, *He came home*, of (142.a). Let E_a be the semantic representation of the event introduced by the first sentence of (142.a) (*John partied until 3am*) and C_b the semantic representation of the event description provided by the second clause. The relation between RT_b and ET_b is coincidence in this case. So the interpretation task presented by the second clause can be formally presented as the left hand side of the equation in (143.a) and the outcome of this task as the value that the function returns for the given arguments; the value lis given as the right hand side of the equation.

(143) a. $\beta_?(C_b, E_a, =) \quad = < conseq, = >$

b. $\beta_?(C'_b, E'_a, =) \quad = < conseq, = >$

c. $\beta_?(C''_b, E''_a, =) \quad = < 0, = >$

d. $\beta_?(C'''_b, E'''_a, =) \quad = < prep, = >$

Exactly the same type of equation characterizes the task and outcome of interpreting the third clause of (142.a) (where E'_a is now the semantic representation of the second clause and C'_b the semantic representation of the event description provided by the third clause (see (143.b)). The corresponding representations of the task and solution of the interpretation of the tenses of the second sentences in (142.c) and (142.f) are given in (143.c) and (143.d). The differences between these last two equations in (143) and the first two only concern the first components of the returned values. For the last two equations these are '=' and 'prep', respectively, whereas for the both of the first two the component is 'conseq'.

Temporal Focus

The last part of Webber's paper is devoted to what she refers to as 'Temporal Focus' (TF). Temporal Focus is the problem of how the antecedents are selected that are involved in particular instance of tense anaphora. In the comments above we separated this issue from the one on which we have been concentrating so far: that of locating the new event in relation to the antecedent once that antecedent has been selected. We also noted, at the end of the comments on what Webber has to say about nominal anaphora, that the selection of the antecedents for nominal anaphors shows only superficial similarity with the selection of TFs. Rightly therefore, Webber devotes a separate section to the latter topic.

A central question about tense anaphora is when and how it produces the effect of temporal progression'. Partee and Hinrichs account for the temporal progression effect by assuming that with each new event comes a Reference Time that is situated shortly after that event and that serves as location time for the next event (which in its turn will introduce a new Reference time shortly following it, and so on). But as Webber emphasizes – we have seen this more than once in these comments – not all instances of temporal anaphora involve narrative progression, not even those where both the new eventuality E_b and and the 'anaphoric antecedent' E_a are events. One of the things that Webber wants her theory to account for is what distinguishes the cases of temporal anaphora that do involve temporal progression from those that do not.

As we have seen, there are two factors in her account that enter into this distinction, the semantics of the tense form of the new sentence or clause (i.e. whether the relation it determines between RT and ET is = or >) and the particular β-function involved. The only cases that yield temporal progression are those in which the relation between RT and ET is = and the β-function is β_{conseq}. In this case the new event E_b is situated within the consequent state of E_a and the Temporal Fopcus can shift from E_a to E_b. (When the β function involved is β_0 or β_{prep} there is no forward movement of the TF: with β_0 the TF remains exactly what it was (and so also where it was), and with β_{prep} there are two possibilities: either the previous event E_a is retained as TF or the new event E_b replaces it, but neither possibility is one of forward movement.)

So far so good. But as Webber observes, discourses are often more complex than the simple examples considered so far, with embedded discourse segments which create their own temporal structures, with or without narrative progression. The next examples provide some comparatively straightforward illustrations of texts with embedded discourse segments. (144.a) is the earlier two-sentence discourse (143.e). In (144.b) the second sentence with its Past Perfect tense, has been inserted between the two sentences of (144.a). There is a sense in which this sentence interrupts the forward movement of the narrative. This second sentence does not never the time forwards, the event of the first sentence remains the TF and it is only the last sentence of (144.b) that causes it to move.

But in addition the second sentence introduces a new event (the promising event) which is situated at some time before the current TF – i.e. the event introduced by the first sentence – and in this way creates an additional potential TF. In (144.c) this option is exploited by the third sentence whose event (Mary's saying that she would never forgive John if he forgot) is naturally understood as following upon the promising event, and not on that of John entering the florist shop introduced by the first sentence. The use of the promising event as TF (and thus as temporal antecedent) in the interpretation of the third sentence of (144.c) is what makes the combination of the second and third sentence into what Webber calls an *embedded discourse segment*.

(144) a. John went into the florist shop. He picked out three red roses, two white ones and one pale pink.

b. John went into the florist shop. He had promised Mary some flowers. He picked out three red roses, two white ones and one pale pink.

c. John went into the florist shop. He had promised Mary some flowers. She said she wouldn't forgive him if he forgot. He picked out three red roses, two white ones and one pale pink.

d. John went into the florist shop. He had promised Mary some flowers. She had said she wouldn't forgive him if he forgot. He picked out three red roses, two white ones and one pale pink.

Before we turn to the difference between (144.b) and (144.c) first a remark about the second sentence that they share. This is a sentence in the Past Perfect, for which Webber adopts the Reichenbachian analysis according to which it denotes the pair of relations > between ST and RT and > between RT and ET. This evidently gives the right result, so long as we can infer that in this case $\beta_?(\mathrm{C}_b,\mathrm{E}_a,=) \quad = <0,=>$. This value of $\beta_?$ establishes RT_b as coincident with the duration of E_a and because of $\mathrm{RT}_b > \mathrm{ET}_b$ (where ET_b is the interval occupied by E_b) it then follows that E_b is located entirely before E_a. But what justifies the assumption that $\beta_?(\mathrm{C}_b,\mathrm{E}_a,=) = <0,=>$?

There is a way of looking at perfect tense forms that differs from Reichenbach's and that helps to clarify this question. We already mentioned this alternative in our comments on Moens and Steedman. According to it all sentences involving perfect tense forms are descriptions of result states. For instance, the sentence 'John has promised to buy Mary flowers' describes a result state of the event of John making this promise, and this sentence, in which the tense form is that of the *Present* Perfect, asserts that this result state holds at the utterance time. Analogously, the Past Perfect sentence 'John had promised to buy Mary flowers' asserts that the result state of John promising to buy Mary flowers held at some past Reference time.

By analyzing sentences with perfect tenses as result state descriptions one classifies these as a special category of state descriptions. And with state descriptions, about which Webber does not say anything specific, the default assumption is that the states they describe always hold over a period that includes what Webber calls the Temporal Focus, irrespective of whether the state described is a result state or some other kind of state (such as, say, that of Mary being sick). True, we have seen that there are exceptions to this principle, as in 'John turned off the light. It was pitch dark.', but in the light of research since Webber's paper it has become increasingly plausible that these are to be regarded as special constructions in which the state description is first coerced into an event description, in the spirit of Moens and Steedman. In the light of this research it seems that temporal anaphora involving state descriptions is governed by different principles than temporal anaphora involving event descriptions – that it is governed by just one basic principle, according to which the described state includes TF, but with the possibility that this principle can be overridden in certain cases where its application leads to absurdity. Cases where the described state is a result state are special only in that they never allow for this kind of overriding.

On this view of Past Perfect sentences, the contribution to (144.b) by its second sentence is like that of any other past tense state description: the described state is said to hold at TF. In Webber's terms this should mean that the β-function that is selected in such cases is β_0. Or better perhaps: I assume that this is so, even though Webber doesn't explicitly relate the choice of β-function to the distinction between state and event descriptions. If we make this assumption then with state describing past tense sentences there is no forward movement (unless they are coerced into event descriptions). Result state descriptions are of course special in that there can be be no result state without there being an event of which it is the result state. Thus one of the implications of (144) is that there was an event of John promising Mary that he would buy her flowers before the time when he was in the shop to make his promise true. That is, the Discourse Model built by an interpreter of the first two sentences of (144.b) will contain a representation of the event of John making his promise as well as the event of his going into the florist shop.

This means that when it comes to processing the next sentence of the discourse, both these events are in principle available as potential TFs. When the next sentence is as in (144.b) ('He picked out three red roses, two white ones and one pale pink.'), then the natural choice of TF is the event of John entering the florist shop. But when the third sentence is as in (144.c) ('She said she would never forgive him if he forgot.'), it is plainly more natural to choose the event of John promising Mary to buy her flowers as anaphoric antecedent for its tense and therewith make it the new TF. What determines these choices in the two cases is, once again, a matter of plausibility reasoning on the basis of world knowledge, about which the paper doesn't have anything specific to say. But Webber observes that there is one important asymmetry between the two choices: When the event of John entering of the florist shop is chosen as TF then that's it and the story about discourse interpretation is just as it has been told up to this point in these comments. But when, as in (144.c), the choice falls on the promising event, then that has the effect of taking the combination of the sentence that introduced this event and the one that now exploits the event as TF as an embedded discourse segment. In such cases the current TF is not fully discarded, but retained by the interpreter as a potential TF for subsequent processing. (In the terminology Webber uses the TF is 'cashed'.) In (144.c) the option that is retained in this way is then exploited right away by the last sentence, which uses the cashed event as TF, with the effect that the picking of the roses is interpreted as following the entering of the florist shop, rather than following Mary's reaction to John's promise.

In the light of what we have been saying about possible analyses of sentences with perfect tenses it is of some interest to compare (144.c) with (144.d), which intuitively has the same meaning, but which differs in that the third sentence is also in the past perfect (like the second sentence). In (144.d) the second and third sentence form what is known as an *extended flashback*. Given what we have just suggested about the semantics of perfects, the third sentence of (144.d) should, like the second one, be construed as asserting that a result state – that resulting from an event of Mary's saying she wouldn't forgive John if he forgot – holds at the TF introduced by the first sentence, viz. that of John going into the florist shop. This locates the event of Mary's saying she wouldn't forgive John if he would forget to buy her flowers in the past of the event of his entering the florist shop, as intuively it should, but it does not establish any relation temporal relation between between Mary's saying what she said and john's promise. However, in this case too the temporal relation between the event of John's promising and that of Mary's saying she would not forgive him if he forgot is, again intuitively, a crucial part of the way the discourse is interpreted.

There are in principle two ways to deal with this aspect of the interpretation of the third sentence of (144.d). One is to adopt the principle that in extended flashbacks relations between a new eventuality and a TF arises at two different levels - a primary level which involves the result states and a secondary level which concerns the events that give rise to those result states. (For a version of this see Kamp & Reyle (1993), Ch. 5.) A second possibility would be to treat the temporal relations between the events as a matter of pragmatically driven inference that falls outside the realm of semantic processing.

Our remarks about the possible analyses of sentences with perfects makes this a natural place to have a look also at another pair of contrasting examples that Webber discusses, given in (145).

(145) a. John went to the hospital. He had broken his ankle on a patch of ice.

b. John went to the hospital. He broke his ankle on a patch of ice.

The meanings of (145.a) and (145.b) are clearly different. The natural interpretation of (145.a) is that the breaking of the ankle happened before John set off towards the hospital and presumably was the reason for his going there. (145.b) cannot be interpreted in this way; its only interpretation, it would seem, is that John broke his ankle on the way to the hospital. Webber's theory can account for this difference in that the 'earliest' possibility for the event described by its second, Simple Past Tense, sentence is that it is located within the preparatory phase of the event of John going to the hospital. And the most natural assumption about this preparatory phase is that it consists of John's moving from wherever he happened to be (home, presumably) and just that. That assumption then entails that the breaking of the ankle happened in the course of this move.

Any account which predicts that the RT of the second sentence coincides with the complete nucleus of the event described in the first sentence (that of John going to the hospital) will get us the intuitively correct interpretation for (145.a): the event of John breaking his ankle happened before this entire nucleus. Webber obtains this result on the assumption that interpretation of the second sentence of involves the β-function β_0; the account which treats the second sentence as the description of a result state and assumes a general default interpretation mechanism for state descriptions makes this prediction because the result state will temporally wrap around the TF. Here we see once more the merit that Webber's account (and likewise of the one by Moens and Steedman) has when compared with attempts to explain temporal relations without reference to non-temporal notions: her account gives us a handle on why (145.b) *cannot* have the interpretation of (145.a).

But as Webber notes, matters are more complicated yet. Consider the following two pairs (146.a), (146.b) from her paper and also the variant (146.c) of the examples in (145), which I have added as a way of elucidating my cureent intuitions about what may be going in these cases.

(146) a. John went to the hospital. He took a taxi because his car was in the shop.

b. John went to the hospital. He had taken a taxi because his car was in the shop.

c. John went to the hospital. He had twisted his ankle when getting out of his car, so now he had to go, as well as he could, to two different parts of the hospital complex.

Contrary to what we saw in connection with (145), the second sentences of (146.a) and (146.b) can be used to describe the very same part of the same episode – that of John going to the hospital by taxi. That the Simple Past sentence in (146.a) can be used for this purpose is in agreement with what we observed in relation to (145): The taking of the taxi seems to be located within the preparatory phase of the event of John going to the hospital and that, we have seen, is one of the things for which Simple Pass sentences in Past tense sentence sequences can be used for. But in the light of what we have been saying about (145) it may seem surprising that the Past Perfect sentence in (146.b) can be used for this same purpose as well. As far as I can see there are two possible reasons for why the use of the Past Perfect is compatible with this scenario. The first is that 'take a taxi' can be understood as 'decide to take a taxi', which can be thought of as describing an event that preceded the nucleus of the event E_a described as 'John go to the hospital; and the second is that, for some reason, the second sentence of (146.b) can be understood as locating its result state at a TF within the consequent state of E_a. Perhaps these two explanations should be seen as working in tandem. How delicate the facts surrounding these examples are is further illustrated by (146.c), which seems quite similar to (145.b), but which can be understood as describing an event of John twisting his ankle that occurred on his way (in)to the hospital. Here too it seems that the contents of the two sentences allow the interpreter to situate the reference time of the new sentence within the consequent state of the TF – the event of John going to the hospital – so that it is possible for the new event to be situated at any time before this RT, including times within the preparatory phase of TF.

But *why* is there this apparent difference between (146.c) and (145.a)? Comparison of the two examples suggests that the interpretation of (145.a) according to which John's breaking his ankle precede his going to the hospital (and which presumably was the reason for his going) rather than something that happened on his way to the hospital is a default interpretation, which can be overruled when more specific context information is available, as it is in (146.c). BUt *when* can the default be overwritten by additional information, and what must that information be like to have the power to overwrite? These are questions that arise whenever a certain interpretation has default status, with the possibility of being overwritten in certain situations. But saying that this is a general problem can't be a substitute for explicit solutions in particular instances. I have no idea how to solve the problem in this instance.

Other ways of providing Temporal Foci and introducing Embedded Discourse Segments

The last two aspects of her general topic that Webber broaches in this paper are related. One is the range of different ways in which Temporal Foci can be made available for the interpretation of tensed clauses and sentences; and the other is the variety of different kinds of segment embedding that may be encountered when discourses become more complex.

(147) is an example of the first of these aspects.

(147) I was at Mary's house yesterday. We talked about her sister Jane. She spent five weeks in Alaska with two friends. Together they climbed Mt. McKinley. Mary asks whether I would like to go to Alaska some time.

In this discourse the third and fourth sentence are naturally understood as forming an embedded segment. We are cued to such an interpretation by the fact that the second sentence refers to an event of the speaker and Mary *talking* about Mary's sister Jane. There seems a clear preference for interpreting the subject *she* of the third sentence as anaphoric to *Jane* and in the light of what we have just been told – that Mary and the speaker were talking about Jane – it is natural to take this sentence and the sentence or sentences following it as elaborating on what the speaker and Mary were discussing about Jane. Starting a new part of the Discourse Model to represent the content of this embedded segment has the effect that the established TF at this point, viz. the event of the speaker and Mary talking about Jane, isn't used to locate the event of the new sentence. it is of course natural and perhaps inescapable to infer that the event of Jane spending five weeks in Alaska preceded the time ('yesterday') when the speaker and Mary were talking about her, but presumably this inference is based on a different interpretation mechanism, in which tense anaphora of the kind discussed in most of Webber's paper (and in most of these comments) doesn't play its usual part.

Once the interpretation of the third sentence of (147) is in place, the fourth sentence can be interpreted by the kind of mechanism that has been in focus in most of these comments, in which the event of the third sentence plays the part of TF. (I note in passing that once again it is not clear to me which particular β-function will give the right result in this case.) That this interpretation of the fourth sentence involves a genuine choice of TF, which must be made on the basis of non-temporal relations between the contents of the third and the fourth sentence, is shown by the interpretation of the fifth and last sentence, which requires a similar choice but where the TF is once again the event introduced by the second sentence.

The relation in which the embedded segment of (147) stands to the 'main narrative' that surrounds it isn't just a purely temporal one. We noted that the events described in the segment must have happened before the meeting between the speaker and Mary that is spoken about in sentences one, two and five. But that these events actually did happen rests on the inference – plausible enough, perhaps, in this case, but not a logical entailment – that what was being said about Jane's trip to Alaska actually did take place, in other words, that what (presumably) Mary told the speaker about Jane and Alaska was the truth. In a discourse like (148), with a structure that is very similar to that of (147), such an inference is clearly not justified.

(148) I was at Mary's house yesterday. She told me all manner of things about her sister Jane. Among them: She spent five weeks in Alaska with two friends. Together they climbed Mt. McKinley. They nearly froze to death, but where saved in the nick of time by a helicopter. Later it turned out that all this was fabricated. In fact Mary doesn't even have a sister.

(147) and (148) are just two examples of embedded segments, and what Webber has to say about this phenomenon is meant to do no more than alert the reader to the fact that embedded discourse segments exist and that they complicate the theory of discourse interpretation in general and of the rules that govern temporal anaphora in particular. Such observations point towards all the work that remains for a comprehensive theory of the semantics and pragmatics of discourse interpretation and they are sobering as comments on the comparison between what it is still to be accomplished and what has been accomplished so far. The hope for automatic text processing, in which such a comprehensive account has been turned into a algorithm that you can run on a computer and that returns a semantic representation (e.g. a DRS or a more complex structure built from DRSs) when you feed it a text, is still what it was then, a distant and fanciful dream.

Another point Webber makes, and the last one to be mentioned in these comments, is that the TFs needed for the interpretation of new tensed clauses, aren't always introduced by earlier clauses but that noun phrases referring to events can do this just as well. Her example to illustrate this is the following variant of (147) .

(149) I was talking to Mary yesterday. She told me about her trip to Alaska. She spent five weeks above the Arctic Circle with two friends. The three of them climbed Mt. McKinley.

In this discourse the TF for the interpretation of the third sentence is the event denoted by the noun phrase *her trip to Alaska*. (Here too it is not entirely clear from what Webber says how the principles of temporal anaphora resolution she discusses earlier on are to be applied; but I presume that the β-function involved is β_{prep}.)

The more general moral here is that reference to events can be expressed by means of noun phrases no less than it can be expressed by means of full sentences and clauses and that we often go back and forth between these two modes of expression. At least as common as the direction exemplified in (149) – from the nominal description *trip* to the tensed sentences that follow – is the opposite direction, in which events are introduced by full clauses,to be followed by nominal references to those events or to events that stand in some perspicuous relation to them. Examples are (150.a) and (150.b). (The second is in Webber's paper, the first has been made up as part of these comments.)

(150) a. Mary climbed Mt. McKinley. The climb/The event took several days.

b. Mary climbed Mt. McKinley. The preparations took longer than the ascent.

(150.a) is not only an example of the general pattern just mentioned – an event introduced by a full clause is resumed through the use of a noun phrase – but also shows one of the reasons for this pattern: if you want to convey the information carried by (150), and you want to do that in two steps, by first saying that the event occurred and then giving information about how long it took, then referring to the event by means of a noun phrase, as in the second sentence of (150.a), offers an advantage of conciseness combined with stylistic variation: the subject NP of the second sentence can be combined with the phrase *several days* that gives the duration of the trip via the verb *took*. Suppose you wanted to phrase the second sentence while avoiding nominal reference to the climb spoken of the first sentence: How could you have done that? The only option, it seems, would have been something like the following sentences.

(151) a. Mary took several days to climb Mt. McKinley/the mountain/this mountain.

b. Climbing Mt. McKinley took Mary several days.

c. When she climbed Mt. McKinley Mary took/needed several days.

All these variants sound awkward. They all make one feel that the same event is introduced twice over, once by the first sentence of (150.a,b) and then a second time by the various sentences in (151). And that is odd. One should't knowingly introduce the same entity twice. Once an entity has been introduced and you then want to say something more about it then the construction you use should make it clear that what you are doing is *resume* reference to it. Resumption is one of the primary functions of definite noun phrases such as pronouns and definite descriptions, and the definite descriptions in the second sentences of (150.a) and (150.b) are naturally interpreted in tat way. In particular, the natural interpretation of *the climb* (or *the* *event* in (150.a) are readily interpretable as resuming reference to the event introduced by the first sentence. But resumption of previously introduced eventualities is not among the primary functions of full finite and infinite clauses. Their primary function is to introduce new eventualities into the discourse. This is what causes the sense that as follow-ups to the first sentence of the discourses in of (150) the sentences in (151) are somehow misused.

(150.b) shows that noun phrases that are anaphoric to events introduced by full clauses need not be 'coreferential': the events they denote need not be those events themselves but can also be events that stand to those in some relation other than identity. In this regard the noun phrases *the preparations* and *the ascent* in (150.b) differ from *the climb* and *the event* in (150.a) in the same way that *the driver* and *the passengers* differ from *the bus* and the third person pronouns considered in the discussion of nominal anaphora in the first part of Webber's paper: *the preparations* and *the ascent* are bridging anaphors in the realm of events, just as *the driver* and *the passengers* are bridging anaphors in the realm of individuals.

This covers the main points of Webber's paper. As we have seen, the paper contains a number of important insights. Its formal proposals can do with some improvements and refinements, and many of its important observations are little more than pointers towards clusters of problems that one was beginning to discern at the time when Webber wrote and that remain equally hard challenges today. The next paper we will consider, by Lascarides and Asher, is the most ambitious and serious attempt to meet many of these challenges. (It still qualifies that today, after the more than 20 years that have passed since its appearance in 1993). But precisely because that paper is much more specific about the non-temporal aspects of discourse interpretation than any of the papers we have considered so far, it shines a glaring light on how very difficult the general enterprise of formulating a comprehensive account of discourse interpretation really is.

4.8 Lascarides and Asher: 'Temporal Interpretation, Discourse Relations, and Commonsense Entailment'

General aims and approach of the paper

Both the paper of Moens and Steedman and that of Webber emphasize the importance of non-temporal factors in the determination of temporal relations between eventualities. As demonstrated by example after example they bring and discuss, world knowledge is often of the essence. And one ends up feeling (and I think, is meant to feel) that it always plays a role, even in those cases which appear to confirm purely temporal interpretation algorithms like that of Partee's *Nominal and Temporal Anaphora*. But while they make a persuasive case for the importance of non-temporal information in the interpretation of tensed discourse, neither M & S nor Webber have much to say about precisely how non-temporal information plays its part in the determination of temporal relations. (All that we get from M & S is that world knowledge enables the interpreter to choose between the different operation that can be performed on the nucleus introduced by sentence S1 to provide the intended temporal anchor for the event (or nucleus) contributed by S2. Webber formalizes this same idea in the form of her different β-functions, one of which selects the consequent state of the nucleus introduced by S1, while a second selects the preparatory process, a third the nucleus was a whole; perhaps more such relations are needed than just these three. But how does world knowledge determine those choices? Even for the examples that those authors discuss this matter is left to the creative imagination of the reader.

The paper by Lascarides and Asher is an attempt – a heroic attempt, one might be tempted to say – to fill this gap by developing a formal account of the interpretation processes involved, in which specific bits of 'world knowledge' play their specific, and explicitly described, parts in the inferential processes that lead to interpretations like those that M & S nor Webber consider. It should be obvious from the start that a detailed formal account of the sort L & A develop cannot get by without making certain assumptions about the form of the bits of world knowledge that they suggest interpreters make use of in order to arrive at their interpretations of tensed discourse. But that is no objection. You cannot expect from a single pair of researchers, whose aim it is to lay bare the general architecture of the reasoning that goes into the determination of temporal relations between events a systematic account of the world knowledge that is available to competent speakers; and in any case there would have been no room for the presentation of such a 'world knowledge module' within a paper of journal length, even assuming one had actually succeeded in spelling out what information goes into such a component and how the information that is relevant for the interpretation of a particular piece of discourse can be retrieved from it effectively.

There are other problems with the paper as well, which can be seen as problems within the context of L & A's own brief. We will come to those eventually. But there are some core cases which had been a problem for everyone until L & A developed their account and much can be learned by looking at their solution to those cases carefully and in detail.

One such example, the first L & A analyze in full detail, is the pair of two-sentence discourses in (152).

(152) a. Max stood up. John greeted him.

b. Max fell. John pushed him.

L & A observe in relation to these examples that when one is offered (152.a) in a context which does not specify any unexpected, out of the way information, we are inclined to understand the second event, that of John greeting Max, as happening after the first event, that of Max getting up (or perhaps as simultaneous with it, but certainly not before it). In contrast, it is natural to understand the second event of (152.b) (that of John pushing Max) as preceding the first event (that of Max falling). And the reason for this difference is intuitively obvious: for (152.b) we get the reverse temporal order from the one we get for (152.a) for the obvious reason that it is natural to understand the pushing as the cause of the falling.

Connected with this inference that an interpreter of (152.b) is likely to draw about the causal relation between the two events it mentions is another aspect of the interpretation of this discourse, which, however, must nevertheless be sharply distinguished from it. This is the *rhetorical* relation between the second and the first of its two sentences. We understand the second sentence as giving some kind of *explanation* for the truth of the first sentence (i.e. for why the event described by that sentence occurred). It is this rhetorical relation of causal explanation that, L & A point out, is directly responsible for the temporal order of the events in (152.b): Awareness that bering pushed by someone can be the cause of falling makes it possible for the interpreter of (152.b) to see the second sentence as describing the cause of the event described by the first sentence and thus to seethe second sentence as standing in the rhetorical relation providing an explanation for the claim made in the first sentence. Since causes do not follow their effects, and more often than not precede them. one infers for the two events ev_1 and ev_2 the reverse temporal order from what appears to be the most plausible interpretation of (152.a).

No Discourse Interpretation without Rhetorical Relations

L & A's point – that it is the rhetorical relation of explanation which is ultimately responsible for the interpretation of (152.b) according to which e_2 cannot be later than e_1 – is part of a general thesis about discourse interpretation, which they take over from *Rhetorical Structure Theory* (RST; see Mann & Thompson (1988)):

(153) In order to perceive a discourse D as coherent the interpreter must be able to assign a rhetorical relation to every new sentence or clause of D, which relates it to some other sentence, clause, or group of sentences or clauses.

For all but the initial sentence or clause of D this clause, sentence or group of sentences or clauses will precede the given sentence or clause: only the initial sentence/clause will be rhetorically connected only to something that follows it.

(153) is one of those sweeping claims that are difficult to prove. But once you have become alert to it, it becomes compelling. And it does that even in the absence of a very clear picture of the repertoire of possible rhetorical relations between which the interpreter can and must choose. This, in fact, is one of the main challenges of RST: what is this repertoire? And the second question is: Which part or parts of the discourse are available as rhetorical relata to a new clause or sentence, so that relating it rhetorically to one or more of those parts can count as a justification of the new clause or sentence

as a coherent contribution to the discourse?

Proponents of RST haven't always been very explicit on these two points. But it is here that the theory L & A present in the paper we are discussing is an important improvement on earlier discussions of rhetorical relations. They explicitly specify a (comparatively small) list of possible rhetorical relations, and also make formally explicit to which part or parts of a given discourse structure a new clause or sentence can be 'rhetorically attached'. This enables them to use their version of RST as part of an account of the semantics and pragmatics of discourse that has real bite (i.e. that can make predictions about when a discourse is coherent and about when it is not, and what its meaning is when it is coherent).

Principle (153) entails that a rhetorical relation between the second and the first sentence is also involved in the interpretation of (152.a). This of course must be a different relation from the one involved in (152.b), for it is one with different implications for the temporal relation between the two events. L & A call this relation *Narration*. It is the relation that holds between a new event sentence and an earlier one – often if not always, its immediate predecessor in the discourse – when the events described by the two sentences stand in what Moens and Steedman call a contingency relation. For L & A Narration carries the entailment that e_2 follows e_1.[6]

The interpreter of (152.b) will infer that its two sentences stand in the rhetorical relation of Explanation on the basis of the plausible inference that the falling was caused by the pushing. It is part of our knowledge of the world that falls can be caused by pushes and that pushes can cause falls and, they further assume, this knowledge can be seen as taking the form of a defeasible 'Push Causal law'. Note well, however, that this 'law' isn't just a piece of knowledge about the world as such. Rather, what it says is that *if* the sentences in question stand in *some* rhetorical relation to each other and *if* the events described by the first and the second sentence are the kinds of events such that an event of the second kind can be a cause of an event of the first kind, *then* the events tare to be seen as standing in this causal relation and

[6] L & A postulate a quite restricted set of possible rhetorical relations, viz. (i) *Narration*, (ii) *Explanation*, (iii) *Elaboration*, (iv) *Background* and (v) *Result*. In these comments only the first two will be directly relevant.

the second sentence as standing to the first in the relation of (causal) Explanation.

In somewhat different words, the interpretation process for (152.b) goes as follows:

(i) Assume that there is some rhetorical relationship between the second sentence S_2 of (152.b) and its first sentence S_1. (There has to be a rhetorical relation between S_2 and S_1 because S_1 is the only sentence or clause that is available for S_2 to rhetorically related to.)

(ii) If there is such a rhetorical relation between S_2 and S_1, then the evenualities ev_2 and ev_1 that S_2 and S_1 describe must be *e-connected.* e-connectedness covers a number of relations between two eventualities, either (a) one is a consequence of the other (I take it that this relation can only obtain when ev_1 is an event and ev_2 is a *consequent state* (or *result state* of ev_1; or (b) one of the two is part of the other (I take this to mean that the first is part of the preparatory phase of the second); or (c) the two stand in some causal relation; or (d) the two temporally overlap.

(iii) Given that ev_1 and ev_2 are e-connected, it can be inferred, via the Push Causal Law, that ev_2 is the cause of ev_1.

(iv) But if that is how ev_1 and ev_2 are related, then the rhetorical relation between S2 and S1 must be the relation of Explanation.

(v) The temporal relation between ev_1 and ev_2 is now doubly supported as it were, by the causal relation between them and by the rhetorical relation between the sentences that describe them: ev_2 cannot be after ev_1, and, for all we know, preceded it. (There remains a certain kind of indeterminacy on this point. Some causes wholly precede their effects; others temporally overlap with them. This seems true also of the special case where the cause is a pushing of *a* by *b* and the effect a's falling down. The pushing may accompany some initial part of the falling or it may initiate it, by making *a* lose balance, so that the falling happened only after the pushing ended.)

So much for L & A's account of (152.b). To summarize: we arrive at its interpretation by (i) assuming that there is some rhetorical relation between

S2 and S1 and thus an e-connection between their events; (ii) by using our world knowledge to infer that in this case e-connection takes the form of e_2 being the cause of e_1; and (iii) inferring from this that the rhetorical relation must be Explanation and the temporal relation between ev_1 and ev_2 one that is commensurate both with the way they are causally related and with the Explanation relation between S2 and S1. In short, a certain amount of world knowledge relating to the kinds of events that S2 and S1 describe (i.e. that they are a pushing of a by b and a falling of a, respectively) is an essential ingredient to this interpretation.

For (152.a) this is different. To arrive at the interpretation that its two sentences are connected by Narration, and that the second event came after the first, no world knowledge is involved. The interpreter's stock of world knowledge may be assumed to contain nothing that points to some special relation between greetings and gettings up. Narration, L & A maintain, is the 'default' rhetorical relation, which an interpreter will assume in all cases where there is no information to overrule this default (as there is in the case of (152.b)). In other words, interpretation settles on Narration as the rhetorical relation between S2 and S1 *in the absence of any overriding information.*

Non-monotonic Reasoning is the Engine that drives the derivation of Rhetorical Relations

A general theory that can handle both (152.a) and (152.b) in the way that L & A want must be able to make sense of the notion of overriding information and of what it is for an inference to be valid in the absence of such information. That is, it must allow for a notion of *non-monotonic reasoning* – and thus must be built on a system of *non-monotonic logic*. A logic is *non-monotonic* when the inferences it licenses are not necessarily preserved when additional premises become available. (That is what we want for the cases at hand: in the absence of further information about special relations between the events described by S_1 and S_2 we are entitled to infer that the rhetorical relation between them is Narration: but when additional information is available, such as that contained in the Push Cause Law, then that inference is valid no longer and has to be replaced by an inference to another rhetorical relation, which is compatible with that information.) But precisely because they must allow for such possibilities non-monotonic logics are complex, and much more so than the traditional deductive logics, which are *monotonic*. (In particular, classical First Order Predicate Logic, the formal logic that has become a kind of standard in applications of logic since it received its first formulation in Frege's *Begriffsschrift*, is monotonic.)

Monotonicity is a simple property of logical systems that can be formally stated as follows:

(154) A logical system is *monotonic* iff its relation $|=$ of valid inference (of a conclusion B from a set of premises Γ) is preserved under the addition of new premises:

if $\Gamma \models$ B and $\Gamma \subseteq \Gamma'$, then $\Gamma' \models$ B.

Non-monotonic systems are logical systems which allow for exceptions to this principle: For certain Γ, B, Γ' it max be that $\Gamma \models$ B and $\Gamma \subseteq \Gamma'$, but not $\Gamma' \models$ B. The problem that confronts the designers of non-monotonic logics is that there are all sorts of reasons why the monotonicity principle (154) might fail. That is why there are all sorts of 'non-monotonic logics', which differ from each other in that one permit failure of monotonicity where another does not.Furthermore, and more importantly, many systems of non-monotonic differ in fundamental aspects of their architecture, and it is because of that that the monotonic inference patterns they preserve are not the same. In these comments little will be said about the architecture of the system of non-monotonic logic that L & A adopt as part of the theory of discourse interpretation they propose, and limit ourselves to a few observations about the validity or invalidity of certain inference patterns that are directly relevant to L & A's analyses of particular bits of discourse.

For now there are just two things about this system that should be pointed out. The first is that the non-monotonicity of the system manifests itself only when the set of premises of an argument contains formulas that have a certain kind of defeasibility built into their meaning. There is only one source of this type of 'semantic defeasibility': a conditional connective $>$, with the property that 'A $>$ B' is true iff B is true in all of the most normal cases in which A is true. For instance, if A is short for 'x is a bird' and B for 'x flies', then 'A $>$ B' can be read 'if x is a bird, then, normally, x flies', or also as 'if x is a normal bird, then x flies'.

From this last defeasible conditional and the further premise 'Tweety is a bird' we can draw the defeasible inference that Tweety flies. But this inference is defeasible insofar as it can be ascertained to be incorrect when it becomes known in addition that Tweety turns out to be not a normal bird, e.g. because he is a penguin. Not all birds fly, and penguins are among the exceptions.

The inference that Tweety flies can be seen as an instance of a general principle that L & A refer to as *defeasible Modus Ponens*. To wit:

(155) a. Defeasible Modus Ponens: A > B (Premise i), A (Premise ii) $|\approx$ B (Defeasible Conclusion).

b. Special instance: (i) Tweety is a bird > Tweety flies (Premise i), (ii) Tweety is a bird (Premise ii) $|\approx$ Tweety flies (Defeasible Conclusion).

$|\approx$ is the symbol L & A use to denote the defeasible, non-monotonic inference relation of their system. But the system comes not only with its relations of defeasible inference but also with a classical notion of non-defeasible inference for which L & A use the symbol '$|=$'. The non-defeasible inference patterns form a proper subset of the defeasible ones; in other words, if $\Gamma \models$ B, then $\Gamma \mathrel{|\approx}$ B, but in general not conversely. (As we already noted, in L & A's system patterns that are defeasibly but not non-defeaibly valid must involve occurrences of '>.) It is of course the patterns of defeasible inference that are of special interest in connection with the use L & A make of their system in the their theory of discourse interpretation.

Non-monotonic logics allow for a three-way distinction between putative inference patterns:

(i) those that are non-defeasibly valid;

(ii) those that are defeasibly valid but not non-defeasibly:

(iii) those that are not even defeasibly valid.

An example of a pattern that is defeasibly bot not non-defewasibly valued in L & A's system is the so-called *Penguin principle*, which says that if an individual a is both a P and a Q, but Q is a special kind of P, and we are given as premises two defeasible conditionals, one of which says that if a is a P then it is an R, while the other says that if a is a Q then it is a non-R, then it is the latter conclusion – that a is a non-R – that we are entitled to (defeasibly) draw:

(156) a. The Penguin Principle:
$\P(a)$, $\mathrm{Q}(a)$, $(\forall \mathrm{x})(\mathrm{Q}(\mathrm{x}) \rightarrow \mathrm{P}(\mathrm{x}))$, $\mathrm{P}(a) > \mathrm{R}(a)$, $\mathrm{Q}(a) > \neg\mathrm{R}(a) \mid\approx \neg\mathrm{R}(a)$

b. Application:
bird(Tweety), penguin(Tweety), $(\forall$x)(penguin(x) $\rightarrow$ bird(x)), bird(Tweety) $>$ flies(Tweety), penguin(Tweety) $>$ $\neg$flies(Tweety) $|\approx$ $\neg$flies(Tweety)

An example of a 'negative' inference pattern (i.e. one in which DMP os prohibited from applying) is the one known as the *Nixon Diamond*: Nixon was both a Republican and a Quaker. Quakers are normally pacifists and Republicans are normally non-pacifists. What can we infer from this about whether Nixon was a pacifist? Answer: Nothing. Formally:

(157) The Nixon Diamond:

Suppose we are given the premise set $\mathrm{P}(a)$, $\mathrm{Q}(a)$, $\mathrm{P}(a) > \mathrm{R}(a)$, $\mathrm{Q}(a) > \neg\mathrm{R}(a)$.

Then neither $\mathrm{R}(a)$ nor $\neg\mathrm{R}(a)$ can be (defeasibly) inferred.

The difference between the Penguin Principle and the Nixon Diamond is at once glaring and subtle. The extra premise $(\forall$x)(penguin(x) $\rightarrow$ bird(x)) in the Penguin Principle, which says that being a Q entails being a P, thereby giving priority to Q when it comes to inferring properties for an individual a that satisfies both, is absent from the Nixon Diamond. In the premise set of the Nixon Diamond the classifications P and Q of a compete, so to speak, on equal terms; neither is in a position to overrule the other.

Note also in this connection that if the premise set only contained the information that Nixon was a Quaker we could have defeasibly inferred that he was a pacifist, even in the presence of the (then irrelevant) information that

Republicans are normally not pacifists. Likewise, if the premises only specify that Nixon was a Republican, then it could have been defeasibly inferred that he was not a pacifist. But as soon as it is known that he was both a Quaker and a Republican, then both those inferences have to be discarded.

What has so far been said about the non-monotonic logic that L & A make use of indicates some of its properties. Theirs is a system with a connective $>$ that has the properties typical of conditional connectives. This connective is the sole source of non-defeasible validity, that is: an inference pattern with premises A_1,...,A_n and conclusion B can be defensibly valid without being non-defeasibly valid only if it contains one or more occurrences of $>$. Furthermore, inference patterns instantiating the Penguin Principle are among those that are defeasibly valid and inference patterns in which the premise set instantiates the premise configuration of the Nixon Diamond and the conclusion instantiates either $\mathrm{R}(a)$ or $\neg\mathrm{R}(a)$ (in the schematic representation used above) are among those that are not defeasibly valid.

It should not be taken for granted that it is possible to construct a logical system with these features. In fact, the construction of such a logic and proving formally that it has the mentioned properties (as well as a number of others, which do not matter here) is a far from trivial matter. That this can be done was shown by Asher in joint work with Michael Morreau not long before the L & A paper we are discussing was published. (See **?**. The system developed in this and other papers from these authors is known as *Common Sense Entailment*.)

In these comments we will assume that such a system is in place and only refer to those properties of it that are relevant to analyses of particular examples if and when we need to. Before we can turn to a detailed discussion of those analyses,we still need to say something about a couple of further assumptions that L & A make. First, the *discourse representations* that are the outputs of L & A's interpretation algorithm are structures built form building blocks that are like the DRSs we encountered in our discussion of Partee's *Nominal and Temporal Anaphora*. The discourse representations L & A envisage – so-called *Discourse Representation Pairs* or *DRP*s – consist of sets of these building blocks, moulded into a single cohesive structure by discourse relations connecting them. (We will continue to speak of 'rhetorical relations' in this connection but with the intention to capture the same notion. At least within the context of the present paper I cannot see anything wrong with this.) These relations make up the second members of DRPs. That is, a DRP is a pair $< \mathcal{K},\mathcal{R} >$, where $\mathcal{K}$ is a set of DRSs and $\mathcal{R}$ is a set of tuples consisting of the name of a rhetorical relation followed by as many elements as the relation takes arguments; these elements can be either members of $\mathcal{K}$ or small subsets of $\mathcal{K}$. The only rhetorical relations we have encountered so far, Narration and Explanation, are 2-place relations between individual members of $\mathcal{K}$, so tuples in $\mathcal{R}$ involving those two relations will have the form $<$Narration,$\alpha,\beta >$ or the form $<$Explanation,$\alpha,\beta >$, where $\alpha,\beta \in \mathcal{K}$. In fact, for the DRPs we will consider $\mathcal{R}$ will invariably be made up entirely of triples $<$Rh,$\alpha,\beta >$, consisting of the name of a rhetorical relation and two DRSs. In particular, the DRPs for (152.a) and (152.b) are structures of the form $<\{\alpha,\beta\}$, $\{<$Rh,$\alpha,\beta >\}>$, where Label is one of 'Narration,' and 'Explanation'.

The rhetorical dimension to discourse interpretation, we noted, involves for each new sentence or clause a couple of decisions: (i) to which part of the DRP as it has been built thus far should the new sentence or clause be attached by a rhetorical relation; and (ii) which rhetorical relation is to provide that link? While these two decisions cannot be kept separate in the practice of discourse interpretation, it is nevertheless possible, and conceptually expedient, to distinguish between them. L & A make this formally explicit by adopting a ternary relation, satisfaction of which they represent by simply juxtaposing terms denoting triples of satisfying arguments within angled brackets. Thus '$<\tau,\alpha,\beta>$' means that the DRP τ, the constituent α of τ and the new sentence or clause β are related in the following way: β is attached to τ through some rhetorical relation between it and the constituent α. The claim expressed by '$< \tau,\alpha,\beta >$' is thus the result of making the first of the two decisions that are needed to interpret a new clause or sentence β of a discourse D: It paves the way for the second decision that must be made as part of interpreting sentence or clause with semantic representation β, viz. that of deciding in which rhetorical relation β stands to α.
In L & A's formalization statements of the form '$<\tau,\alpha,\beta>$' play an important role, as antecedents to conditionalized claims that would have little or no plausibility – or even make no real sense – without such a qualification. This is true of, among others, one of the 'laws' that are invoked in the reconstruction of the interpretation of (152.a), the so-called law of *Narration*:

(158) **Narration**: $\langle \tau, \alpha, \beta \rangle > Narration(\alpha, \beta)$

For both examples (152.a) and (152.b) this principle has a direct and unambiguous application. Right now we state this application just for (152.a), on which we are focusing at this point. We will turn to its role in the interpretation of (152.b) below. Suppose that the first sentence of (152.a) has been processed and that α is the DRS that resulted from this. So at this point the DRP is the pair $< \{\alpha\}, \emptyset >$. Since this DRP has only one constituent, vix. α, to which the interpretation (that is: the DRS) β of the second sentence could be attached, '$< \tau,\alpha,\beta >$' *must* hold, as α is the only possible attachment point. From this and the Narration principle (158) we can defeasibly infer:

Narration(α, β).

The next inference towards the conclusion about the order of the events described in (152.a) makes use of a second principle connected with Narration, which L & A refer to as the *Axiom of Narration*.

(159) **Axiom of Narration**: $Narration(\alpha, \beta) \rightarrow me(\alpha) < me(\beta)$

One bit of notation in this formula that has not yet been explained are the expressions 'me(α)' and 'me(β)'. For any clause or simple sentence γ, 'me(γ)' stands for 'the main eventuality of γ'. This is a notion with which we are by now well familiar. For us, the main eventuality of a clause or sentence γ is the one that, in the implementations of DRS-construction went through in connection with Partee's *Nominal and Temporal Anaphora*, gets introduced by tense. In fact, we may as well build on our comments to that paper in our presentation of L & A's analyses, by assuming that the DRSs for clauses and simple sentences they appeal to are constructed along the lines shown in our discussion of Partee's paper. Applying the rules considered there, but ignoring all reference to r-points (which are essential to the Hinrichs-Partee story but are not considered explicitly by L & A), we get for the first sentence of (152.a) the DRS α shown in (160). The main event e_1 has been marked in bold face.

(160)

$m \quad \mathbf{e_1}$
m = Max $\quad \mathbf{e_1} < n$ $\mathbf{e_1}$: *stand-up*(m)

Similarly, if we process the second sentence of (152.a) in the same way (except that we rely on α for the resolution of *him*), then we obtain the DRS β.

(161)

$$\begin{array}{|c|}\hline j \quad \mathbf{e_2} \quad u \\ \hline \\ j = John \quad \mathbf{e_2} < n \quad u = m \\ \mathbf{e_2}\text{: } greet(j,u) \\ \hline \end{array}$$

The crucial step in the derivation of the conclusion '$e_1 < e_2$' we are aiming for comes now – although it won't be until we reach the same point in our reconstruction of true interpretation of (152.b) that the importance of this step can be fully appreciated. The true nature of the step is invisible as it were, insofar as it relies on the *absence* of information in the interpreter's 'Knowledge Box" about events e_1 and e_2 of the kinds described in α and β – absence of information that could have interfered with the conclusion that is to drawn. In the absence of such information – which means that all information about the present example that is available to the interpreter is given by the premises of the instantiation of the Axiom of Narration to the current α, β, me(α) ($= e_1$) and me(β) ($= e_2$) – we are licensed to draw the conclusion that $e_1 < e_2$. (L & A use 'me(γ)', where γ is the DRS for some sentence or clause S, to denote the 'main eventualtiy of γ'. This iOS the discourse referent ev in the universe of γ that stands for the eventuality described by S (and that, in the DRDSs construction algorithm used in our comments on Partee's 'Nominal and Temporal Anaphora' is introduced by the tense of the main verb of S). Frop mthe construction of γ from S it is always clear which discourse referent in the universe of γ is its main eventuality.)

There is a slight problem about where in the DRP for (152.a) the condition '$e_1 < e_2$' should be added. The DRS set of this DRP consists of the DRSs α and β and '$e_1 < e_2$' expresses a relation between drefs belonging to α and β respectively. One reasonable solution is to add the condition to the new DRS β. That has the effect that this DRS now contains an occurrence of the discourse referent e_1 that belongs to the universe of α but not to its own universe. This makes β referentially dependent on α. But as a matter of fact we already allowed for such a dependence anyway when we used the discourse referent m from α to resolve the pronoun *him* from the second sentence. In general the DRS set of a DRP weill contain DRSs that referentially depend on one or more other DRSs in the set. This is just a consequence of the general *raison dêtre* of DRT as a method capturing the effects of trans-sentential anaphora.

Let β' be the DRS we get when adding the condition '$e_1 < e_2$' to β. Then the DRP for the two-sentence discourse (152.a) becomes: $<\{\alpha, \beta'\}$,$\{$Narration$(\alpha, \beta)\}>$.

Let us summarize the succession of steps that lead to this DRP from the DRP for the first sentence. Available as premises for the inferences that are involved in this interpretation process are:

(162) (i) The statement $\langle \tau, \alpha, \beta \rangle$, where α and β are the DRSs in (164.a) and (164.b) and τ is the DRP $< \{\alpha\}, \emptyset >$ for this new α.

(ii) The principle $\langle \tau, \alpha, \beta \rangle > Narration(\alpha, \beta)$ for the α and β and τ in question.

(iii) The relevant instance of the Axiom of Narration: $Narration(\alpha, \beta) \rightarrow e_1 < e_2$, where e_1 and e_2 are the particular discourse referents occurring in α and β.

Since these are all the premises available to the interpreter, he is entitled to infer, first, by an application of Defeasible Modus Ponens, $Narration(\alpha, \beta)$ and from that in combination with the Axiom of Narration, by an application of (non-defeasible) Modus Ponens that $me(\alpha) < me(\beta)$, i.e. that $e_1 < e_2$.

The formal result of these inferences is a DRP τ for the two sentences of (152.a) the first member of which is the set of the two DRSs α and β and the second member consists of one relation specification, viz. that α and β are connected by Narration. There is a slight complication connected with the condition '$e_1 < e_2$': where should this condition be added to the new DRP? We make the somewhat arbitrary decision that it be added to the condition set of the new DRS β. The result then is the DRP given in (170).

$$
< \{\ \alpha : \quad \begin{array}{|c|}\hline m \quad \mathbf{e_1} \\ \hline \\ m = \text{Max} \quad \mathbf{e_1} < n \\ \mathbf{e_1}\text{: } \textit{stand-up}(m) \\ \hline \end{array}\ , \qquad \{\ \text{Narration}(\alpha, \beta)\ \}\ >
$$

(163)

$$
\beta : \quad \begin{array}{|c|}\hline j \quad \mathbf{e_2} \quad u \\ \hline \\ j = \text{John} \quad \mathbf{e_2} < n \quad u = m \\ \mathbf{e_2}\text{: } greet(j, u) \\ e_1 < e_2 \\ \hline \end{array}\ \}
$$

The Interpretation of (152.b): How to override the Conclusions of (152.a).

The analysis of (152.b) parallels that of (152.a) in a number of ways. α and β are now the DRSs in (164.a,b).

(164) a.

$$
\begin{array}{|c|}\hline m \quad \mathbf{e_1} \\ \hline \\ m = \text{Max} \quad \mathbf{e_1} < n \\ \mathbf{e_1} : fall(m) \\ \hline \end{array}
$$

b.

$$
\begin{array}{|c|}\hline j \quad \mathbf{e_2} \quad u \\ \hline \\ j = John \quad \mathbf{e_2} < n \quad u = m \\ \mathbf{e_2}\text{: } push(j, u) \\ \hline \end{array}
$$

Once more we have that $<\tau, \alpha, \beta>$, where the DRP τ after the interpretation of the first sentence is again the pair $<\{\alpha\}, \emptyset>$. So once again we can infer defeasibly that Narration(α, β). But at this point there is a difference. The interpreter's Knowledge Box does contain, L & A assume, special information about the possible connection between events e_1 and e_2 as described by the current α and β: pushing's are natural causes of falls. To repeat, that doesn't mean that every fall a person makes is caused by a push from someone else, or that every pushing of a person leads to their falling down. But on the assumption that there must be some kind of connection between the two events – which is implicit in the assumption that there exists a certain rhetorical relation between the sentences of which they are the main events – the potential causal relation between pushing and falling acquires additional bite. L & A capture this in the form of their 'Push Causal Law', which they first state informally as in (165.a) and then formally as in (165.b).

(165) (Push Causal Law)

a. If β is to be attached to α, and α describes an event e_1 of x falling and β describes and event e_2 of y pushing x, then, normally, e_1 causes e_2.

b. $(< \tau, \alpha, \beta > \ \& \ me(\alpha) : fall(x) \ \& \ me(\beta) : push(y, x)) > cause(me(\beta), me(\alpha))$

In addition there is a non-defeasible connection between causation and temporal order: the occurrence of the cause of an effect cannot be entirely after the occurrence of that effect. L & A give this in the following form:

(166) 'Causes Precede Effects' :

$$cause(e_1, e_2) \rightarrow \neg(e_2 < e_1)$$

Furthermore, if two sentences S1 and S2 are rhetorically related and the main event of the second sentence is the cause of the main event of the first, then the rhetorical relation between them is (not Narration but) *Explanation.* (167) states this principle and (168) makes explicit that Explanation and Narration are mutually exclusive principles.

(167) Explanation:

$$\langle \tau, \alpha, \beta \rangle \ \& \ cause(me(\beta), me(\alpha)) > Explanation(\beta, \alpha)$$

(168) Axiom of Explanation:

$$Explanation(\beta, \alpha) \ \rightarrow \neg Narration(\beta, \alpha)$$

Let us assume that the interpreter's Knowledge Box contains the 'Push Causal Law' as well as the 'Causes Precede Effects'-principle and the principles 'Explanation' and 'Axiom of Explanation' and that the interpreter can readily instantiate these to the α and β at issue. Let us also assume that the Knowledge Box yields no other relevant information about $me(\alpha)$ and $me(\beta)$. Then the interpreter now finds himself with the following set of premises:

(169)

(i) the statement $\langle\tau,\alpha,\beta\rangle$, where α and β are the DRSs in (164.a) and (164.b).
(ii) the principle $\langle\tau,\alpha,\beta\rangle > Narration(\alpha,\beta)$ for the α and β in question.
(iii) the relevant instance of the Axiom of Narration:
$Narration(\alpha,\beta) \rightarrow me(\alpha) < me(\beta)$
(iv) the Push Causal Law:
$(\langle\tau,\alpha,\beta\rangle$ & $me(\alpha)$: $fall(x)$ & $me(\beta)$: $push(y,x)) > cause(me(\beta), me(\alpha))$
(v) Causes Precede Effects: $cause(e_1,e_2) \rightarrow \neg(e_2 < e_1)$
(vi) Explanation :$\langle\tau,\alpha,\beta\rangle$ & $cause(me(\beta), me(\alpha)) > Explanation(\beta,\alpha)$
(vii) Axiom of Explanation: $Explanation(\beta,\alpha) \rightarrow \neg Narration(\beta,\alpha)$.

This is a much larger set of premises to draw the inferences from that will lead to the interpretation of the discourse than we had in the case of (152.a). Crucial is that besides the defeasible conditional in (ii) the set now also contains the defeasible conditional in (iv). Note that the antecedent of (169.iv) is more specific than the antecedent of (169.ii): anything that satisfies the conjunction '$(\langle\tau,\alpha,\beta\rangle$ & $me(\alpha) : fall(x) \& me(\beta) : push(y,x))$', also satisfies its first conjunct '$\langle\tau,\alpha,\beta\rangle$', but the converse implication does not hold in general. Thus, according to the Penguin Principle (iv) 'wins' against (ii). At this point we are entitled to infer its consequent, viz, that $cause(me(\beta), me(\alpha))$. From that, (i) and (vi) we can then infer that $Explanation(\beta,\alpha)$ and from that that $\neg Narration(\beta,\alpha)$. Also from (v) we can infer that $\neg(e_2 < e_1)$.

This concludes the interpretation of (152.b). The formal result is the DRP $<\{\alpha,\beta'\}$,{Explanation$(\alpha,\beta)\}>$, where now α and β are the DRSs of (164) and β' is obtained from β by adding the condition ' $\neg(e_2 < e_1)$'. Note that the conditional in (ii), which did the main work in the interpretation of (152.a), is side-lined by the more specific (iv). The applicability of the Push Causal Law to the present case shows it to be a 'non-normal' case of succession of two Simple Past sentences.

Reflections

A crucial question, even in relation to these two quite simple applications, is what one should take to be the premise sets to which the decisive applications of Defeasible Modus Ponens (DMP) are made. In particular, there are at least two strategies that one could contemplate: (i) At any point where a defeasible inference is drawn in the course of discourse interpretation, *all* facts about the case that are available to the interpreter at that point are part of the premise set which must defeasibly validate (in the sense of $|\approx$) the conclusion inferred. (ii) Reasoning proceeds 'locally', from premise sets that are 'directly available' at that stage of the reasoning, and which need not contain all information that is available to the interpreter in principle; in this second case there may be the danger that inferences drawn earlier on in the interpretation process prove untenable when, later, additional premises are brought into play; in particular an earlier application of DMP may reveal itself to have been untenable after all because it ignored premises that in principle could have been taken into account at that point and that would have invalidated the inference then and there if they had been. The second strategy therefore requires that a number of things have to be made explicit: (a) Which premises are available at any one point in the course of the interpretation process for applications of defeasible inference principles? (b) What kinds of revision of earlier parts of a chain of inferences are required/allowed when in the course of a complete interpretation process new premises become available?

L & A are not explicit on the choice between options (i) and (ii), but since they say nothing about the complications that arise with the second option, it seems reasonable (and fair to them) to assume that they assume the first one. Note that this option also has the advantage of enabling us to say that an interpretation of a discourse D in a well-defined context c can be qualified as a 'correct interpretation of D in c', in the sense that the defeasible conclusions that are part of that interpretation – e.g. that Narration(α,β) or that $e_1 \prec e_2$ – do follow in the sense of $|\approx$ from the totality of premises that are made available by the interpretation of D up to the very point when β is attached to α in the DRP τ via some particular rhetorical relation. Exactly what that set is must of course be made explicit as well, and as we have seen that will depend among other things about what relevant information is contained in the interpreter's Knowledge Box. In fact, it would be desirable if we could maintain that all of the interpreter's Knowledge Box is available as part of the premise sets from which the defeasible inferences are drawn, but that too is a matter that L & A do not address.

All in all it is important to see in wheat sense the interpretations obtained with the help of defeasible inference in the sense of L & A are 'correct'. On the one hand the conclusions that go into such interpretations are defeasible in the sense that if the premise sets from which they were directly or indirectly drawn had been larger (in certain ways) than they actually were, then those conclusions would have been unjustified. But on the other hand – and this, I take it, is a crucial point about discourse and discourse interpretation to which L & A would also subscribe – it is part of producing a discourse D in a given context c that all the information that is to be taken account in the interpretation of D is contained in (i) some general, presupposed set of facts and principles in some generally available knowledge base; (ii) some specific information provided by c (how contexts provide such information would have to be made explicit as well) and (iii) the properties of the DRP τ as it has been constructed at the plaint where a defeasible inference has to be drawn that is needed for the attachment of a new clause interpretation β. In what we have seen the defeasible inferences arise as part of determining the rhetorical relations via which β is attached to α. As far as I can tell that is, given the general architectural assumptions that L & A make, always the context in which the interpretation of a discourse requires defeasible inference. Once this totality of premises has been specified, for each stage in the interpretation of D, the defeasible inferences must always be justified with respect to this complete set of premises. But some of the rules of correct discourse production would, on the present view, be that *no more* than what is in this set needs to be taken account in interpretation. Only what follows defeasibly (or, of course, indefeasibly) from this premise set can be regarded as germane to the discourse as the speaker or author intends it.

On this view of discourse production and discourse interpretation the discourse interpretations (DRPs) that L & A predict are to be seen as their predictions of what are *correct* interpretations of the discourses D in question *given* the assumptions that are made about the contexts in which these discourses are produced and meant to be interpreted. Thus, in particular, the discourses (152.a) and (152.b), as they are presented in L & A's paper, are to be understood as produced and interpretable in a 'neutral' context, which contains no additional information about the circumstances in which Max stood up and John greeted him. But of course, when these sentences occur as part of a larger discourse, then the earlier parts of that discourse may have provided information that defeats the inference of a narration relation between the two sentences. E. g. the earlier discourse may have specified that John never greets anybody until he has got up. And a special context in which (152.a) is uttered without antecedent discourse may convey such defeating information too.

So it is in the light of these assumptions, I propose, that we should see what L & A have to say about (152.a) and (152.b) and the other discourses they discuss. In particular, it seems reasonable, given those assumptions, that the premise set from which Narration(α, β) is to be inferred in the case of (152.a) includes the premises in (171) and, further, that all the additional members of the total premise set (viz. all that is contained in the Knowledge Box that competent speakers of the language are supposed to share) make no difference to what can be inferred from this first, much smaller set. Note well: this last claim is non-trivial, since we may assume that the presupposed Knowledge Box contains many defeasible conditionals, one of which is crucial in the interpretation of (152.a). But we will assume (with L & A) that this is indeed the case: adding some or all of the further sentences contained in the generally presupposed Knowledge Box would make no difference to the inferences that are licensed by the set specified earlier under (162) and repeated below.

(162)

(i) The statement $\langle \tau, \alpha, \beta \rangle$, where α and β are the DRSs in (160) and (161) and τ is the DRP $\langle \{\alpha\}, \emptyset \rangle$.
(ii) The principle $\langle \tau, \alpha, \beta \rangle > Narration(\alpha, \beta)$ for the α and β and τ in question.
(iii) The relevant instance of the Axiom of Narration:
$Narration(\alpha, \beta) \rightarrow e_1 < e_2$,
where e_1 and e_2 are the particular discourse referents occurring in α and β.

At this point of our reflections on L & A's proposals it is important to become even more explicit about the form of these statements. If I understand L & A correctly, these statements should be interpreted as statements about the particular discourse D. To do justice to this let us adopt an individual constant 'D' as part of the formal representation language in which the inferences involved in the interpretation of D are stated and executed (and also additional such constants if and when referenced is needed to more than one discourse at the same time). '$\langle \tau, \alpha, \beta \rangle$' is now to be understood as a predication of the discourse D (i.e. of the denotation of 'D'), to the effect that D can be partially described as having the structure given by τ and that the DRSs α and β for its first and second sentence are related by some rhetorical relation. In other words, when fully spelled out, (162.i) takes the form of a predication involving the predicate '$\langle \tau, \alpha, \beta \rangle$' and the singular term 'D' and, sticking more or less to the abbreviators format we have been using, we can represent the relevant condition as '$\langle \tau, \alpha, \beta \rangle$(D)'. Similar modifications will be needed for the other premises, by making 'D' an argument of the predications they involve. But not all the predications really depend on D, and when there is no such dependence there is no real need to mention D as part of the predication, and these predications are left in the same form in which they appear in (162). For an example, the form of (162.ii), which is now to be represented as a defeasible conditional of which both antecedent and consequent are predications involving 'D', should be revised to (170) since the satisfaction conditions of '$Narration(\alpha, \beta)$' do not depend on D.

(170) $\langle \tau, \alpha, \beta \rangle(\mathrm{D}) > Narration(\alpha, \beta)$

Note well that in this and other applications of the mechanisms of non-monotonic inference in L & A's accounts of discourse interpretation the instances of Narration and of the Axiom of Narration must be obtained through instantiation (to the given D, α, β and τ) of *general* principles of discourse interpretation which are part of the linguistic component of the interpreter's Knowledge Box. These principles allow for countless other instantiations (for instance to the discourse (152.a) and DRSs of its two sentences). But one of the background assumptions of the proposed analysis of (152.b) (and likewise of the other examples for which L & A give explicit analyses) is that addition of either of these general principles themselves or of such other instances of them won't affect what defeasible inferences may be drawn from the given premise set. All that matters is that the one defensible conditional about the particular D in question that is part of the premise set in (162) is not defeated by any new conditional that can get added in this way. With our revised format for the premises it is easier to see why the additions won't affect the defeasible inference for which the premise set in (162) was used (the one which yields the conclusion $Narration(\alpha, \beta)$): None of the new premises web have described have antecedents that take the form of predications the argument of which is 'D'. So none of them competes with (170) in the way that could have the effect of (170) being overruled. As it stands, this argument for the 'overwrite immunity' of (170) only applies to the potential effect of additional premises of the forms just discussed. But I hope and expect that it can be extended to cover all other premises that could plausibly be added to the premise sets from which discourse interpretations are derived.

We can revise the analysis of the interpretation of (152.b) along the same lines. The new set of premises, built form predications of the discourse D' of (152.b), is given in (171).

(171) (i) The statement $\langle\tau,\alpha,\beta\rangle$(D'), where α and β are the DRSs in (164.a) and (164.b) and τ is the DRP $<\{\alpha\},\emptyset>$ for this new α.

(ii) The given instance of the Principle of Narration:
$\langle\tau,\alpha,\beta\rangle$(D') $> Narration(\alpha,\beta)$, for the α, β and τ in question.

(iii) The relevant instance of the Axiom of Narration:
$Narration(\alpha,\beta) \rightarrow e_1 < e_2$, where e_1 and e_2 are the particular discourse referents occurring in α and β.

(iv) A statement expressing the relevant facts about the falling and pushing events described in the α and β of D':
$e_1 : fall(m)$ & $e_2 : push(j,m))$

(N.B. this condition can be inferred from premise (i) (by non-defeasible inference) and thus is strictly redundant. Defeasible inference is invariant under addition to the premise set of sentences that can be non-defeasibly inferred from it.)

(v) The relevant instance of the 'Push Causal Law':
$\langle\tau,\alpha,\beta\rangle$(D') & e_1: $fall(x)$ & e_2: $push(y,x)) > cause(e_2,e_1)$

(vi) The relevant instance of the 'Causes Precede Effects' Law:
$cause(e_2,e_1) \rightarrow \neg(e_1 < e_2)$.

(vii) The given instance of the Principle of Explanation:
$\langle\tau,\alpha,\beta\rangle$(D') & $cause(e_2,e_1) > Explanation(\alpha,\beta)$.

(viii) The relevant instance of the Axiom of Explanation:
$Explanation(\alpha, \beta) \ \rightarrow \neg Narration(\alpha, \beta)$.

Inspection of the premise set (171) shows that it contains only three defeasible conditionals whose antecedents entail the statement in (i), viz. that the discourse D' has the properties that (i) attributes to it. These are the formulas in (ii), (v) and (vii). Since (iv) tells us that D' satisfies the antecedent of (vi), which is at least as strong as the antecedent of (ii), the Penguin Principle would suggest that it is (v) that can be used in an application of DMP to yield the conclusion that '$cause(e_2, e_1)$'. However, as things stand, we aren't licensed to draw this inference because we haven't assessed the relative logical strength of the antecedents of (v) and (vii). That (vii) does not block the use of (v) in an application of DMP rests on another inference principle, related to the Penguin Principle. This new principle to the effect that if the antecedent A' of a defeasible conditional A' > B' is a non-defeasible consequence of the antecedent A and and consequent B of a given defeasible conditional A > B, then the presence of A' > B' does not block the application of DMP to A > B. The intuition is that when B is a defeasible consequence of A and thus holds for all normal cases of A, and A and B jointly entail A' non-defeasibly, then the normal cases of A must all be normal cases of A'. Consequently there cannot be anything wromng with inferring B, even in the presence of the condition A' > B'. for since any minimal A=world is a normal A' world, B' holds in every normal A-world, so the defensible conditional A > (B & B') is true as well and there is no competition between A > B and A' > B' of the mutually defeating sort that is characteristic of Nixon Diamond configurations.

In virtue of this additional inference principle we are permitted to proceed as follows: We first infer, non-defeasibly, the conjunction of (i) and (iv), which gives us the antecedent of (v). Then, relying on the new principle, we infer '$cause(e_2, e_1)$' using (v). This formula can be conjoined with (i) to give the antecedent of (vii). Next, DMP is to be applied to (vii) and its just established antecedents. To justify this application we have to appeal to yet another inference principle, which is a variant of the one discussed above. A defeasible conditional competing with (vii) is (v), whose antecedent implies $\langle \tau, \alpha, \beta \rangle$ and which provides quite specific information about the events e_1 and e_2. Isn't there a possibility of (v) defeating (vii)? Intuitively that would seem an unnecessary worry, since we just used (v) to obtain the missing conjunct of the antecedent of (vii). We can make this intuition more explicit by representing (v) and (vii) schematically as being of the forms (A & C) > B and (A & D) > C. From these representations it is plain that all normal worlds of A & D are normal worlds of A & C. So B holds in all normal worlds of A & D. So this lis another case where two defensible conditionals point in the same direction, so to speak, rather than b;locking each other as they do in Nixon Diamond configurations. In cases where two defeasible conditionals are semantically related in the way of (A & C) > B and (A & D) > C the inference to B is licensed, not less than it is inn cases that are covered by the inference principle introduced in the last paragraph.

With the inference of the conclusion of (vii), the condition 'Explanation(α,β)', the interpretation of (152.b) is nearly complete. This inference gives the rhetorical relation that attaches β to α. And we can use (vi) to infer the condition '$\neg(e_1 < e_2)$' that expresses the temporal relation between the eventualities e_1 and e_2.

This second reconstruction of the interpretation of (152.b) illustrates an important aspect of L & A's propsal. As soon as our premise sets get larger, the question when applications of DMP are permissible becomes more com,pled: Once we seen,as L & A go out of their way to stress, that under certain conditions, exemplified by the premises of the Nixon Diamond, such applications are not licensed, and that in instances of the Penguin Principle one such application is licensed but the other is not, it is clear that each premise set containing defeasible conditionals with instantiated antecedents presents a potentially new problem of DMP licensing. What one would therefore want as general logical support for a theory of discourse interpretation that makes use of a regime of non-monotonic inference in the way that L & A do this, are general theorems about the underlying non-monotonic logic which classify the possible premise sets into types and specify for each of those types which defeasible inferences are legitimate and which are not. It is not obvious to me from what L & A say in their paper (nor from other papers on Common Sense Entailment that I have seen) how such theorems could be stated, let alone proved.

4.8.1 Other cases

Case 1: An incoherent piece of discourse and why it is incoherent

The bit of discourse (172), L & A observe, impresses us as incoherent.

(172) (?) Max won the race. He was home with the cup.

On the one hand we want to understand the second sentence as expressing the result of his winning the race: Max is back home with his trophy, basking, it may be assumed, in his triumph. But somehow the discourse doesn't let us get away with that interpretation. The succession of these two sentences, both in the simple past, but the first describing an event and the second a state, somehow conveys the impression that the state of Max being at home that is described by the second sentence was simultaneous with the event described in the first. That would perhaps be a possible interpretation, if it hadn't been for the PP *with the cup*. The only reasonable interpretation of *the cup* that comes to mind in the context suggested by (178) is that it refers to the trophy Max got as a reward for winning the race – that interpretation is made possible and forced upon us in the given context by our world knowledge – and it is also part of our world knowledge that such a trophy is conferred upon its new holder only after he or she has won the relevant competition. It is this the world knowledge-based inference that the state described in the second sentence *must* have come *after* the event of winning which conflicts with the defeasible discourse principle, 'States Overlap', according to which the state described by a state-describing sentence in the Simple Past must have started before the beginning of the episode that the discourse describes. That is actually an extremely strong statement and one that is unnecessarily strong for the purposes of dealing with this particular example in the way L & A propose. A more modest and more reasonable variant, which was proposed for instance in Kamp (1981*a*), is that when a state-describing sentence in the Simple Past directly follows a Simple Past event sentence, then the state described by the second sentence temporally includes the event described by the first. (See also the proposal opt Partee discussed in her 'Nominal and Temporal Anaphora', and the comments on that paper in these Notes.) For present purposes the weaker principle does all that L & A need here; it contradicts the conclusion that the described state must start after the event rather than before it.

Before we continue with our discussion of L & A's proposal for dealing with (172) let us have a quick look at the version we get when we drop the PP *with the cup*, as in (173.a).

(173) a. Max won the race. He was at home.

b. Max switched off the light. It was pitch dark.

c. The people started to leave. The room was empty.

d. The people left. The room was empty.

e. The people left. The room was empty now except for a baby that was sleeping in a pram which for some reason nobody had taken.

f. The people started to leave. The room was empty now except for a baby that was sleeping in a pram which for some reason nobody had taken.

(173.a) still strikes us weird and perhaps hardly less so than (172). But the reasons aren't quite the same. First, there no longer is the hard incompatibility between (a) the linguistic principle according to which the state described by the second of two Simple Past sentences temporally includes the event from the first sentence and (b) our worlds knowledge about winning competitions and getting trophies. Perhaps Max really was at home when he won the race; perhaps the race was some kind of computer game, in which one can take part by logging on to the right website, something you can do from your laptop or computer at home. Or perhaps Max was the owner of the horse that won the derby but couldn't attend the race because he was sick in bed.

Still, although such interpretations aren't completely ruled out for (173.a), they do seem rather far-fetched and one would expect it to be possible for them to be pre-empted by an interpretation of the state as following the event, in the sense that Max has already come home from winning the race, if such an interpretation was possible at all. And of course such interpretations are sometimes possible. L & A discuss such an example themselves, given here as (173.b). This little discourse seems fine, although the state described by its second sentence clearly follows the event described by its first sentence. So why isn't a similar interpretation, in which the state of

Max being at home comes after the event of his winning the race, possible for (173.a)?

The answer to this question involves, if I am right, two very different considerations. The first comes into focus when we compare (173.a) with (173.c) and (173.d). (173.c) and (173.d) are considered in Partee's 'Nominal and Temporal Anaphora'. When a state-describing sentence in the Simple Past follows a Simple Past event sentence, then the described state can be understood as following the described event *only* if it can be conceived as the result state of the event. What disqualifies (173.c) is that the state described in the second sentence – that of the room being empty – cannot be understood as the result state of the event of the people *starting* to leave the room. The state *can* be understood as the result state of the people leaving the room, which predicts that (173.d) should be fine. As a matter of fact, (173.d) isn't all that felicitous either, but that is for a different reason. This can be seen when we compare (173.e) with (173.f). (173.e) is variant of (173.d), in which the state description gives a little more information than that of (173.d). The state described can still be understood as the result state of the event from the first sentence, but its description now contributes something about the state that cannot be inferred just from the fact that it is the result state of the event. This eliminates the redundant, or pleonastic, character of (173.d) and that is enough to make (173.e) perfectly acceptable. In contrast, (173.f), in which the second sentence has been modified in the same way as it has in (173.e), but where the first sentence is like that of (173.c) in that it speaks of the beginning of the people leaving, remains bad.

The implications of this for (173.a) should be clear. The state of Max being at home cannot be construed as the result state of his winning the race. That suffices to disqualify (173.a) as infelicitous. (If the state could be interpreted as result state, then presumably there would still have been the redundancy problem. Possibly this is idle speculation, but in any case it is not needed for the explanation we are after.) So in this case, the only interpretation that remains is one on which the state of being at home temporally includes the event of winning. That is, as we already observed, a rather outlandish interpretation, and one that could easily be defeated by additional information about the case. But if (173.a) is offered out of the blue, then it does seem possible; and nothing else is.

Insofar as this is a possible interpretation for (173.a), the account that L & A offer of (172) cannot be applied to it, for that account entails that no coherent interpretation can be obtained. Let me summarize what that account comes to (adjusted to (172), but in a wa that doesn't distort L & A's intentions, since they do not pay attention to the special role that I have argued is played by the phrase *with the cup* which makes for the difference between (172) and (173.a)). L & A argue that the interpretation of (173.a) involves two defeasible principles, States Overlap and a world knowledge-based principle that L & A refer to as the 'Win Law', which asserts that – here I quote their paper verbatim – 'If e_1 is Max winning and e_2 is Max being at home, then normally these eventualities don't overlap'. They claim that these are both defeasible conditionals, with contradictory conclusions (viz. overlap as opposed to non-overlap of the eventualities ev_1 and ev_2 and with logically incomparable antecedents: the antecedent of States Overlap asserts that the sentence S_1 which describes ev_1 and the sentence S_2 which describes ev_2 are successive Simple Past tense sentences in a discourse and the antecedent of the Win Law mentions the types of the two eventualities ec_1 and ec_2 (as a winning and a being at home) while saying nothing about the discourse configuration of sentences describing those eventualities. This combination, they say, is an instantiation of the Nixon Diamond: Neither the consequent of the one conditional nor that of the other can be inferred; and that, they contend, accounts for the incoherence of (172): the interpreter is drawn in two different directions and thus gets himself into a kind of stalemate situation, much like Buridan's ass: rather than plunge for one of the two options and hope for the best of it, the interpreter throws in the towel and declares a discourse that puts him into such a quandary as ill-formed; a coherent discourse just doesn't do such a thing to its interpreters.

Since I want to come back to this explanation be;low, let me give the formal versions of these two defeasible conditionals here.

(174) a. (StatesOverlap)

(Past-Tense(S(α)) & Past-Tense(S(β)) & Succ(S(α),S(β),D) & State(me(β)))
> Me(α) O me(β)

(Here 'O' stands for 'overlap', 'S(α)' for the sentences whose semantic representations α and likewise for'S(β)', 'Past-Tense(S(α))' for the statement that S(α) is a simple past tense sentence and likewise for 'Past-Tense(S(β))', and 'Succ(S(α),S(β),D)' for the statement that in the discourse D S(α) is followed immediately by S(β). These predicates, by the way, arena t find in L & A's paper. I have made them up for the sake of arriving at a [plausible formalizations in a format that is compatible with what has been said in these comments jul to this point. I do not think that these improvisations interfere with L & A's intentions.)

b. (Win Law)

(me(α): win(m) & me(β): be-at-home(m)) $>$ me(α) $<$ me(β)

Our discussion of (173.a)) indicates that this conclusion cannot be right. From what L & A say and do not say about the intuitive meaning of the second sentence of (172), I argued, the particular application of the Nixon Diamond L & A propose in relation to (173.a) should apply equally to (172). But (173.a) does allow for an interpretation, we saw, in which one of the defeasible conditionals, viz. the Win Law, is overruled and a way is found to make sense of Max winning while being at home. That we do not get such an interpretation for (172) has to do, I suggested, with the apparent impossibility of interpreting *the cup* in any other way than as the trophy for the victory mentioned in the first sentence. Once that interpretation is in place, world knowledge excludes the overlap of e_1 and e_2 categorically and we are in a logical situation that the Nixon Diamond has nothing to do with.

So much for this particular example. What strikes me as more disturbing about the strategy that L & A adopt in order to deal with (172) is the idea that when the premise set to an inference involved in the determination of a rhetorical relation between two clauses of a discourse has the configuration of the Nixon Diamond then that explains why the interpretation process aborts at that point. It isn't just that this doesn't seem to be right for the particular example that L & A choose to illustrate the role that they take Nixon Diamonds to play in discourse interpretation. I suspect that the notion that the Nixon Diamond could play such a role at all is misconceived in general. To see why this might be so let us go back to the 'Urexample' of Nixon Diamond inference situations: that of Nixon himself, who is described by the premises as both a Quaker and a Republican while it is known that, as a defeasible generalization, Quakers are pacifists and that, as another defeasible generalization, Republicans are not. The moral of the Nixon Diamond is that if that is all you have to go by, you cannot, even defeasibly, draw any conclusion about Nixon's attitudes towards war and peace. For all that the premises tell you he could be a pacifist or he could be the opposite. (As those who can recall the days when the Nixon Diamond made its way into the lore of non-monotonic reasoning will remember, the answer to the question 'Is or isn't Nixon a pacifist?' was plain to everybody at the time.)

But why should the circumstance that at some point in the course of the interpretation of a piece pot discourse the available premises paperer to instantiate the Nixon Diamond pattern be a reason for concluding that the discourse is *incoherent*? Perhaps there is a prima facie plausible case that could be made for this claim: if the combination of discourse content and the information about the context in relation to which thee discourse is being interpreted do not succeed in jointly pro provideviding a premise set from which the conclusions that are needed in order to arrive tav a completed interpretation can be unequivocally inferred, then something is amiss (with the discourse in relation to the given context). But as our discussion of (172) and (173) has shown, this doesn't seem the way discourse interpretation works. What this way of looking at the role of the Nixon Diamond ignores is the remarkable ability of human interpreters to *accommodate* contextual information needed to make the interpretation of a discourse complete when it isn't on the basis of the information to which they have direct access. This is true especially in connection with bits of text that are presented 'out of context', as is typical in particular of the sample discourses that are found in papers on natural language semantics and pragmatics, including the example presentations in **?**. But accommodation is also often an effective ploy for 'saving discourse coherence' when we interpret spoken language, as witnessed in the extensive literature on accommodation of presuppositions.

The question what role Nixon Diamond configurations play in discourse interpretation si contacted with another worry about L & A's approach to discourse interpretation in general. We already drew attention to the cause for this worry in our discussion of the examples in (152): What information is supposed to be available to the interpreter at which stages of the inferencing process that is to lead eventually to the determination of rhetorical relations? My own hunch is that all available information – world knowledge, linguistic knowledge and details so far established about the particular discourse in question – are available to the interpreter at all stages of the inference process, and that it is impossible to reach negative conclusions about interpretability (i.e. that the discourse cannot be given a coherent interpretation) just because at some stage of the interpretation process the totality of available information is pruned down to a premise set which, by some principle of non-monotonic reasoning, produces an unresolvable stalemate. Admittedly that is only a hunch. But it is impossible to go beyond that so long as we haven't been given the details of exactly how the premise sets at the different stages of the process are determined. And this is a topic that, as far as I can see, L & A do not even touch.

There is a further question which relates to this last worry while at the same time targeting more specifically the intended application of the Nixon Diamond to (172) and (173.a). This is the question what exact forms should be assumed for the two defeasible conditionals that together give rise to the Nixon Diamond pattern that L & A invoke in their discussion of (172) and which should equally apply to (173.a) in the light of what they say,. It is essential to their argument that the Win Law does *not* incorporate information about the fact that, in the case in question, the eventualities e_1 and e_2 are the main eventualities of two Simple Past tense sentences that follow each other in a discourse. For if that information would be added to the antecedent of the Win Law then we wouldn't be in a Nixon Diamond situation, but rather in a Penguin principle situation, which would license that, States Overlap notwithstanding, e_1 and e_2 do not overlap, but that e_2 follows e_1. The point to be made here is a general one. The point is made clearly, if only implicitly, towards the end of Morreau's 'Fainthearted Conditionals' (see **?**). In this paper Morreau discusses the following example: tea with milk tastes good, tea with lemon tastes good, tea with milk and lemon does not taste good. Each of these three premises is formalized as a defeasible conditional, with the antecedents 'M' for 'There is milk in the tea', 'L' for 'There is lemon in the tea' and 'M & L' for 'There are both milk and lemon in the tea'. Suppose that the tea that I am offered has both milk and lemon in it and that I want to decide whether this tea will be good, in order to be able to decide whether or not to accept the offer. Morreau argues correctly that if it weren't for the last premise – the premise, which says that tea with milk and lemon is no good – then no conclusion could be drawn; more specifically, no conclusion about the offered tea being good or not good could be drawn from the premise set with which we are left when this last premise is taken away from it, and that in spite of the fact that both of the remaining conditionals have the same consequent, viz. that the tea tastes good. What blocks the inference from this reduced premise set is that in the background lurk two additional premises – two 'hidden' defeasible conditionals one of which says the (normally) when tea has milk in it then it does not have lemon in it, and (normally) when tea has lemon in it then it doesn't have milk. These hidden conditionals, which should be seen as part of the premise set whether they are explicitly mentioned or not, have the effect that the two overtly mentioned conditionals, about tea with milk being good and about tea with lemon being good, cancel each other out, so that neither can be used to infer that the tea I am offered tastes good. This is as it should be. It is only when the additional information that (normally) tea with both milk and lemon does not taste good is added that this premise can overrule both other conditionals – its antecedent represents a non-normal case of tea with milk and also a non-normal case of tea with lemon – and thereby makes it possible to infer that the offered tea doesn't taste good.

While Morreau's example isn't a case of the Nixon Diamond, his discussion is highly relevant for issues around the Nixon Diamond as well. What makes the original Nixon Diamond example compelling is that Nixon does not seem to conform to the normal expectations about Quakers nor to the normal expectations about Republicans: in the background of this example lurk the hidden conditionals that when someone is a Quaker, then (normally) he isn't a Republican and when someone is a Republican, then (normally) he isn't a Quaker.

But in the case provided by the discourse interpretation of (172) the background contains no such hidden conditionals. To the extent that the Win Law is plausible it surely isn't rendered less plausible by the circumstance that the eventualities e_1 and e_2 of which it speaks are the main eventualities of two successive Past Tense sentences in a discourse. True, in a case where that is so, there is competition with the principle that L & A refer to as 'States Overlap', but why shouldn't that be the kind of competition that can be resolved as it is in the case of (152.b)? Of course we do not want this conclusion in the case of (172) (or, for that matter, in that of (173.a)), because that would yield an interpretation that the discourse doesn't have. But what blocks the inference in the one case but not in the other? L & A's answer – that (173.a) involves an inference situation that instantiates the Nixon Diamond while (152.b) involves one involving the Penguin Principle –

feels like it is based on gerrymandering of premise sets towards the desired outcomes, and for that reason just doesn't seem convincing.

In fact, the theoretical problems we are facing in connection with (172) appear to be even more severe than this last assessment of the situation has made explicit. It is hard to see for instances why the Win Law would have to be stated in the form (174.b) and why it couldn't have been stated just as plausibly with in the form in (175) below:

(175)

(Past-Tense(S(α)) & Past-Tense(S(β)) & Succ(S(α),S(β),D) & State(me(β)) & me(α): win(m) & me(β): be-at-home(m)) > me(α) < me(β)

When the Win Law were to be stated as in (175) then its antecedent would be more specific than 'States Overlap' in the sense of the Penguin Principle and it would be possible to infer that $ev_1 \ < \ eve2$. We have seen that such an inference would be wrong for the 'incoherent' (172) and also for (173.a), which, if it has an interpretation at all, has one in which the event ev_1 and the state ev_2 do overlap.

The tentative conclusions from our scrutiny of (172) and (173) partly echo those from the last subsection: It is difficult to know exactly hot to evaluate L & A's proposal because it doesn't tell us enough about which premises are supposed to be available for non-monotonic inferences when. However – this I take to be the more specific pushot of our remarks of the present subsection about the plain incoherence of (172) and the oddity of (173.a) – it seems highly unlikely that Nixon Diamond configurations can be made responsible for the discourse incoherence in the manner L & A want. If Nixon Diamonds, anywhere along the inferential path of discourse interpretation, were certain harbingers of incoherence, then many discourses would be predicted as infelicitous that competent speakers do not experience as infelicitous or odd at all.

Case 2: The Structure and Interpretation of Discourses with more than two clauses.

So far we have dealt in these Comments on L & A with just three examples – those in (152) and (172) – each of which consists of just two sentences, and in each case our focus has been on the question how the second sentence is rhetorically connected with the first. But what L & A offer is a general theory of discourse interpretation, which is to be applicable to a much wider range of cases, the vast majority of which involve more clauses than just two. In that respect (152) and (172) are quite special. Among the cases L & A want to cover are examples like (176), in which the DRSs that form the nodes of the representation, are arranged in a kind of hierarchical structure. For instance, in (176) the sentences (c) and (d) form a segment that is subordinated to (b) and in which the ;arguer segment (b)-(e) as a whole is subordinated to (a).

(176) a. Guy experienced a lovely evening last night.

b. He had a fantastic meal.

c. He ate salmon.

d. He devoured lots of cheese.

e. He won a dancing competition.

The intuition that the structure of (176) has such a hierarchical structure can be made precise by dividing rhetorical relations into two classes, those that *coordinate* and those that *subordinate*. Of the two relations that have played a part in the analysis of (152.a) and (152.b), Narration and Explanation, Narration is a coordinating and Explanation a subordinating relation. For the analysis of two sentences discourses like those in (152) this distinction makes no difference, but when a discourse gets more complex as in (176), then the difference between coordinating and subordinating relations becomes important.

The two rhetorical relations that are relevant to the analysis of (176) are Narration and *Elaboration*. Elaboration holds between the DRSs α and β of two clauses of a discourse iff β describes some details of the event described by α. For instance the DRSs DRS(c) and DRS(d) of the clauses (c) and (d) in (176) each stand in the Elaboration relation to the DRS(b), and DRS(b) and DRS(e) stand in the Elaboration relation to DRS(a). Narration is also exemplified in (176) in that DRS(e) stands in this relation to DRS(b) and DRS(d) to DRS(c). Elaboration is, like Explanation, subordinating. This means that between them Explanation and Narration impose on the representation of (176) the hierarchical structure shown in (177).

(177)

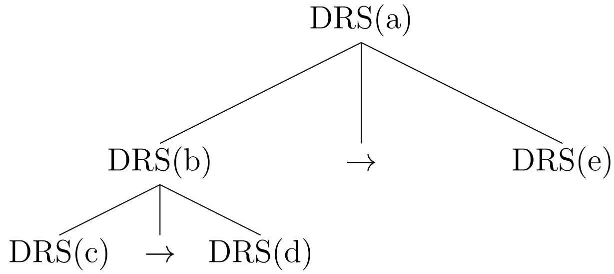

(The arcs of the tree represent subordinating rhetorical relations between the upper and lower node they connect. The arrows represent coordinating relations between the connected nodes.)

As soon as we have a discourse structure with more than one member to which the DRS of a new sentence needs to be attached a problem has to be faced that doesn't arise for two-clause discourses like those in (152): To which α should the new β be rhetorically attached? The need to choosee between different possible attachment points for a new clause non-trivially complicates the interpretation procedure as we have been describing it so far, for now we have to run the non-monotonic inference machinery that is needed to determine by *which* rhetorical relation β is attached to α simultaneously for the different possible premises $\langle\tau, \alpha, \beta\rangle$(D) corresponding to the different possible attachment points α. The problem is mitigated by the fact that in general only a subset of the nodes α of τ are possible as attachment points for β. The set of possible candidates is limited to what L & A call the *right frontier* of τ. The *right frontier of* a Discourse Representation Pair τ is defined as follows: Start at the highest level of τ's hierarchical structure. The one node at this level that belongs to the right frontier is the one that is right-most at that level. If that node has no descendants (i.e. nodes related to it by some subordinating rhetorical relation), then that is all of τ's right frontier. If the node has descendants, then the one of those descendants that is right-most at the next level will be part of the right frontier as well and the same procedure is now continued in relation to that node: if no descendants, then we are done, otherwise add the right-most descendant to the right frontier set; and so on. For the representation in (177) this gives as right frontier the node set {DRS(a), DRS(e)}, and for the representation for the first four sentences, to which the representation of DRS(e) still needs attaching, the right frontier is {DRS(a), DRS(b), DRS(d)}.

The Right Frontier Constraint – the principle that only nodes on the right frontier are possible attachment sites for the semantic representation of a new sentence or clause – has strong support from our intuitions about discourse incoherence. For illustration suppose that the speaker of (176) had continued with (176.f):

(176) f. Even more overwhelming was the tarte flambée he had for desert.

Intuitively this is an incoherent continuation of the discourse in (176) and the Right Frontier Constraint explains this. Given its content DRS(f) should be attached by Elaboration to DRS(b) and by Narration to DRS(d). But after (176.e) that is no longer possible, since neither of these nodes belongs to the Right Frontier any longer: attachment of DRS(e) to DRS(b) has made both it and the nodes subordinate to it inaccessible. So the new sentence cannot be incorporated into the discourse in the way we feel it ought to be if at all; it can't because the discourse has 'moved on' from the point where this would still have been possible, for instance, if (176.f) had followed (176.d) instead of (176.e).

As the example illustrates, the Right Frontier Constraint can reduce the set of potential attachment sites for a new clause or sentence considerably; and the size of these reductions grows disproportionally as discourses get longer: the longer the discourse the larger, on average, is the proportion of the nodes that the Right Frontier Constraint rules out. But only in some cases does the Right Frontier Constraint reduce the set of potential attachment sites back to one. Examples of where it doesn't are the cases already considered: there are three possible attachment sites for (176.e) in the discourse representation for (176.a-d) and there are still two attachment sites for (176.f) in the representation (177). These examples are still very simp;le when compared with the discourses and texts that interpreters on a daily basis as listeners and readers. But a closer look at them will nevertheless enable us to see how much more complicated things get when interpretation of a new sentence involves a choice between two or more attachment points.

Let us begin with the attachment problem presented by DRS(e). This case is discussed by L & A, who treat it as an instance of what they call 'discourse popping'. Discourse popping is involved when a new sentence of the discourse has to be attached to the already constructed discourse representation at a higher level than that of the most recently attached node of the DRP. The attachment of DRS(e) in (177) is an instance of that, in that it is an attachment to DRS(b) and thus not at the level of DRS(d), the most recently attached DRS. L & A want to account for this case (and I presume also other cases) of discourse popping by appealing to purely structural principles of discourse attachment and rhetorical relation determination. (They return to this issue more than once, but for our purposes a review of their first pass at the matter will suffice.) L & A argue as follows: There is a general, non-defeasible exclusion principle to the effect that for any α and β, if Narration(α,β), then $\neg$ Elaboration(α,β). (Narration and Elaboration exclude each other just as Narration and Explanation.) The possibility of attaching DRS(e) to DRS(b) forms the antecedent of the default principle called 'Narration': (τ,DRS(b),DRS(e)) > Narration(DRS(b),DRS(e)). Together with the non-defeasible exclusion principle just cited this instance of Narration yields the defeasible conditional (τ,DRS(b),DRS(e)) > $\neg$ Elaboration(DRS(b),DRS(e)), where τ is the substructure of (177) that consists of the nodes DRS(a), DRS(b), DRS(c) and DRS(d). Next L & A appeal to a principle involving the three nodes DRS(b), DRS(d) and DRS(e). Generally, if a discourse representation has three nodes α,β and γ, α and β stand in the relation of Elaboration – Elaboration(α,β) – and α and γ do not stand in this relation – $\neg$ Elaboration(α,γ) – then (defeasibly) β and γ cannot stand to each other in the relation of Narration – $\neg$ Narration(β,γ). Like the other principles cited in this review of what L & A have to say about this particular attachment case, this is very plausible: if we have already established that α and β stand in the relation of Elaboration, then attaching γ to β by Narration amounts to interpreting it as the continuation of the elaboration of α that is already under way through the previously established Elaboration(α,β). Applying this principle to the nodes of our structure yields the defeasible conditional: Elaboration(DRS(b),DRS(d) & $\neg$ Elaboration(DRS(b),DRS(e)) > $\neg$ Narration(DRS(d),DRS(e)).

The punchline is to come now. Connected with the second possibility, that of attaching DRS(e) to DRS(d), is another instance of the principle Narration: (τ,DRS(d),DRS(e)) > Narration(DRS(d),DRS(e)). So now, L & A argue, we have two defeasible conditions with logically incomparable antecedents and contradictory conclusions, viz.

(i) (Elaboration(DRS(b),DRS(d)) & ¬ Elaboration(DRS(b),DRS(e))) > ¬ Narration(DRS(d),DRS(e))

(ii) $\langle\tau$,DRS(d),DRS(e)$\rangle$ > Narration(DRS(d),DRS(e))

Suppose that DRS(e) instantiates the antecedents of both these conditionals. Then we have what looks like another instance of the Nixon Diamond and as in the case of (172) this can be taken as an indication that the given assumption leads to the impasse of an incoherent interpretation. In this way the interpreter, L & A seem to assume, can verify that an attempt to attach DRS(e) to DRS(d) aborts and that another attachment site must be found.

Our considerations have already enabled us to see how problematic appeals to Nixon Diamond configurations as indicators of discourse incoherence can be. But in the present case such an appeal seems even more problematic (if that were possible) the the ones that have been encountered so far in these comments. It is more problematic in that some of the premises which make up the Nixon Diamond pattern are only hypothetical. This is so for the antecedent $\langle\tau$,DRS(d),DRS(e)$\rangle$ of the second of the conditionals (i) and (ii) above. That premise is *assumed* in the derivation we are looking at *in order to* prove that DRS(d) *cannot* be the attachment site for DRS(e) – in other words, in order to reduce it ad absurdum. If there were compelling reasons for believing that Nixon Diamond configurations are a trustworthy sign of discourse incoherence when they arise in situations of direct reasoning – i.e. when all four premises of a given Nixon Diamond configuration have been established free of derivation-internal assumptions – then we might contemplate whether this role of Nixon Diamonds could be extended also to hypothetical cases like the one that is before us now. But since we have seen that the non-hypothetical instances of the Nixon Diamond are unreliable incoherence indicators to begin with, the basis for such an extension to hypothetical cases is lacking.

That the NIxon Diamond could play such a role in connection with the attachment of DRS(e) as part of the interpretation of (176) can also be argued more directly. Suppose it was possible to make the argument work that

DRS(e) cannot be attached to DRS(d), by using principles of discourse interpretation that completely ignore the contents of the new DRS and the discourse representation into which it is to be integrated. Then this very same argument should now prevent DRS(e) from being attached by Narration to DRS(b) and by Elaboration to DRS(a). And by the same token, DRS(d) could never have been attached by Narration to DRS(c) and be Elaboration to DRS(b). Obviously (as expressed in the trivial observation above) no such procedure, which only makes use of purely structural principles of discourse interpretation and discourse structure could possibly succeed here.

Intuitively, the reason why DRS(e) should be attached to DRS(b) (rather than to DRS(d)) by Narration and to DRS(a) (rather than to DRS(b)) by Elaboration seems clear enough. What DRS(e) talks about isn't part of a meal, but rather an event that one would assume followed the meal referred to by DRS(b). These are conclusions we draw on the basis of world knowledge. World knowledge, in this case, is just as important here as it is to explain the difference between (152.a) and (152.b) or the incoherence of (172). From a perspective of theoretical elegance and conciseness the constant need to draw on world knowledge, and the 'Laws' (such has the 'Push causal Law' and the 'Win Law') is of course unappealing, especially since in the absence of a more explicit theory of the content and structure of 'Knowledge Boxes' any such appeal has the flavor of the ad hoc. So one can sympathize with L & A's desire to limit those appeals as much as possible. But in their account of discourse popping this desire appears to have got the better of their sound judgment.

We encounter the very same difficulties when we try to apply L & A's proposals to the problem of explaining why (176) cannot be felicitously continued with (176.f). Suppose that it were possible to account for the incoherence of (176.a-f) with no more than the content-neutral principles which L & A appeal to in their account of discourse popping. Then the very same conclusion could be reached for the following alternative (176.f') to (176.f).

(176) f'. He had a double of one of Scotland's finest Highland malts at the White Hart on his way home.

(176.f') of course is fine as a continuation of (176.a-e) and for the intuitively obvious reason that a stop at a pub on the way home can be naturally understood as the next event in the sequence that started with the meal and then moved to the dancing competition. The reason why (176.f) cannot be understood in this way is that it is hard not to take it as making reference to a further course of the meal which is the topic of (176.b,c,d). So, it is in virtue of its content that (176.f) ought to be attached by Narration to DRS(d) and by Elaboration to DRS(b). But that is impossible because at the point when it is (176.f)'s turn to be interpreted the Right Frontier Principle has already made these nodes inaccessible.

If the Right Frontier Constraint plays a crucial role in accounting for the incoherence of (176.a-f), it can do so only in conjunction with a mechanism that is capable of recognizing DRS(d) as the proper attachment point for (176.f). Incoherence then follows from the fact that DRS(d) has already been removed from the Right Frontier. Recognition of DRS(d) as the proper attachment point must be based on the recognition that (176.f) is about the consumption of food. But there is an additional, and even more compelling reason for identifying DRS(d) as the designated attachment point for DRS(f). The phrase *for desert* is anaphoric in a sense that is broadly similar to the one in which we found *the cup* to be anaphoric when we reflected on the content of the second sentence of (172) and its connection with the first sentence. The phrase *for desert* raises the question: 'Desert of which meal?' and the only answer that the discourse context enables us to come up with is that this meal must be the one spoken of in (176.b,c,d). As in the case of (172), therefore, anaphora resolution establishes a link to some earlier part of the discourse. But this time the link comes with the implication that the event described in (176.f) is a part of the event described in DRS(b). That identifies (176.f) as an elaboration of (176.b). But for an attachment to DRS(b) it is now too late.

I think there can be little doubt that this is the reason why continuing (176) with (176) strikes us as infelicitous. But it isn't at all obvious how a story along these lines is to be formalized within the framework L & A develop. The case suggests that substantial adjustments of that framework may be needed.

Case 3: Another Tense Form: The Pluperfect.

The large majority of the examples L & A discuss are discourses consisting of sentences in the Simple Past. In fact, the tense forms of the sentences that make up those discourses hardly play any part in what they have to say about those discourses, although it is of course perfectly plain that the tenses of the sentences do play a part. For instance, if in any of those examples the tense of one of the sentences is changed from a past into a present or future tense, then none of the things L & A have to say apply. It is easy enough to say (in a first approximation at least) what the Simple Pasts contribute to all these discourses: each Simple Past locates the main eventuality of its sentence in the past of the utterance time. But evidently that still leaves a lot of room and also any possible temporal order between the main eventualities of the different sentences that constitute the discourse. And L & A are of course right to emphasize, as others like Moens and Steedman and Webber had done before them, that the challenge is to account for how the main eventualities of such sentence sequences are related in each case and to articulate how their temporal relations are tied up with non-temporal relations, and how these non-temporal relations are recognized. Still, the Simple Pasts do play their parts in all these examples, parts that differ from those that would have been played by other tenses in the same position, and it is evident that a theory of discourse interpretation that is to deal with more tenses than just the Simple Past would have to make these differences explicit eventually.

There is only one tense form other than the Simple past which L & A do consider and that is the Pluperfect. But again, the semantic role of the tense form as such does not play much of a role in the account they propose. For the discourses they consider, in which Pluperfect sentences occur in conjunction with Simple Past ones, this is arguably somewhat more problematic than it is for the discourses that consist of Simple Past tense sentences only.

In fact, it isn't clear whether L & A consider the Pluperfect a distinct tense at all, with a semantics that is different from that of the Simple Past, in the way that is assumed in most of the literature on Tense and Aspect (including much of the literature on DRT, in which the semantics of tense played a central role from the very beginning). Although L & A use DRSs as building blocks for their discourse representations, these DRSs are always DRSs for 'sentence radicals', expressions that are like the complete sentences of which their discourses are made up, except that the finite tense has been removed, such as 'John greet Max', which describes an event but without any indication of its temporal location.) The following quotation from 'Temporal Interpretation, Discourse Relations and Common Sense Entailment', pp. 30,31 gives a flavor of their point of view: 'Our formalism reflects the intuition that the pluperfect acts as a syntactic discourse marker to indicate that only a restricted set of discourse relations is possible, thus yielding different inferences about discourse structure.' While there is an important point to this way of looking at the Pluperfect, the quote is striking in its complete silence about those temporal and/or aspectual properties of the Pluperfect that have been central to most other studies of the semantics of tense.

It will be useful in connection with the comments on L & A's treatment of the Pluperfect that will follow, to review, in a few lines, the two main analyses of the Pluperfect that can be found in the Tense and Aspect literature. The first of these is the one of Reichenbach, which we looked at in considerable detail earlier on. To repeat, Reichenbach treats the Pluperfect as a tense in the full semantic sense of the word. Reichenbach's semantic analysis of the tenses involves besides the speech time ST and the event time ET a third element which he calls 'reference time', or RT and each tense is characterized by two temporal relations, one between ST and RT and one between RT and ET. In particular, the Pluperfect is characterized by the relations 'RT before ST' and 'ET before RT' (whereas Reichenbach takes the Simple Past to be is characterized by: 'RT before ST', 'ET coincides with RT').

On the second analysis the Pluperfect is what is in the morphology of the name 'Pluperfect', and perhaps even more clearly in its other name 'Past Perfect', viz. that it is the past of a perfect. On this view the Perfect is an operator in its own right that is present in all sentences with 'perfect' tenses, Past Perfect, Present Perfect, Future Perfect, and the various infinitival Perfects. The Perfect is an aspectual operator, which transforms any eventuality description into the description of the corresponding result state: the state s specified by the transformed descriptions is that state which is the result of the eventuality specified by the original description, which serves as input to the Perfect operator. On this second view the Pluperfect is, qua tense, just a Simple Past, but of a sentence in which the Perfect operator has applied, and thus a Pluperfect sentence is a special kind of Simple Past tense stative sentence, for which the state described is a result state.

Although these two analyses are on the face of it quite different, they yield in general the same kinds of predictions about the contributions that Pluperfect sentences make to the discourses of which they are part. This is so in particular, for the simple discourses involving Pluperfect sentences to which L & A (with one exception) limit their attention, in which a Simple Past sentence is followed by a Pluperfect sentence. The reason for this has to do with the principles that govern the determination of reference times that is part of Reichenbach's account. For discourses in which a Simple Past sentence is preceded by a Simple Past sentence it is the eventuality introduced by the latter that identifies the Reference time for the former. The effect of this is that the eventuality of the Pluperfect sentence is located in the past of the eventuality of the Simple Past sentence. When the Pluperfect is analyzed as the Simple Past of a perfect, then the connection between it and the Simple Past sentence preceding it is much the way it is as in other cases where a Simple Past sentence is followed by another Simple Past sentence that describes a state: The described state is interpreted as temporally including the eventuality described by the first sentence. This entails that the eventuality of which the Pluperfect sentence describes the result state, which ends the moment that result state starts, must be in the past of the preceding Simple Past sentence. Thus the temporal relation between those two eventualities is the same as that predicted on the Reichenbachian account.

L & A's own account of the Pluperfect would seem to be closer to the 'past of a perfect' account than to Reichenbach's. (This, among other things, is suggested by their remark, quoted above, that the Pluperfect 'acts as a syntactic discourse marker': the discourse marker is an addition to the tense of the sentence, which is a Simple Past just like the tenses of the other sentences occurring in the discourses they discuss (those sentences whose tense form is the Simple Past in the traditional, morphological sense of the word). But they appear to be somewhat ambivalent on the matter, to the point of inconsistency in their own terms. One difference between Reichenbach's account and the 'past of a perfect' account of the Pluperfect is that they make different predictions about what the 'main eventuality' is of a Pluperfect sentence. On the Reichenbach account this is the eventuality that is located in the past of the Reference Time and thus, in L & A's 'Simple Past + Pluperfect' discourses in the past of the eventuality contributed by the first sentence. On the 'past of a perfect' analysis the 'main eventuality' is the result state, which in 'Simple Past + Pluperfect' discourses is simultaneous with the eventuality of the first sentence. By their own lights L & A do not seem to be quite consistent on this point. At one point of their discussion of 'Simple Past + Pluperfect' discourses they invoke the 'States Overlap' Principle, which they also make use of in their argument why (172) is incoherent, which entails that when the second of a pair of Simple Past sentences is stative, the state it contributes overlaps the eventuality from the first sentence. But elsewhere, and more dominantly, they seem to think of the main eventuality as the one that Reichenbach locates in the past of R. This is so, for instance when they argue that the rhetorical relation between the two sentences of such a discourse can be understood to be Explanation because it is possible to perceive some kind of causal connection between their main eventualities.

However, setting aside what appears to be an infringement of their own terminology, L & A seem to be right in allowing both the result state and the eventuality of which it is the result state to play their respective parts in their account of what contributions Pluperfect sentences make to the discourses in which they occur. But I shouldn't jump the gun. Let us turn finally to a look at a few of the examples involving Pluperfect sentences that L & A consider.

Our primary focus will be on discourses of the 'Simple Past + Pluperfect' form. The first observation L & A make, which is in line with their over-all concern for the rhetorical relations that must be put in place as part of any proper discourse interpretation, is that the discourse in (178) is infelicitous, while those in (179) and (180) are fine.

(178) Max poured himself a cup of coffee. He had entered the room.

(179) Max entered the room. He poured himself a cup of coffee.

(180) Max slipped. He had spilled a bucket of water.

When we compare (178) and (179) we see that they both appear to describe a succession of two events: (i) Max enters the room and (ii) Max pours a cup of coffee. This is the interpretation we get for (179), on the intuitively plausible assumption that there is nothing in the interpreter's Knowledge

Box to interfere with the default assumption that the two sentences stand in the relation of Narration. Likewise, the content of (178) would seem to be just such a succession of events: the result state of the second sentence temporally includes the event described in the first – i.e. that of Max pouring a cup of coffee – and that situates the event whose result state it is – i.e. the one of Max entering the room – before the coffee-pouring event. Why is (179) a felicitous way of presenting this succession of events while (178) is not? Clearly this is something about which an account that restricts itself to the purely temporal dimensions of the semantics of tense is unable to explain. But L & A have a convincing story about this. It is a story that finds support in the comparison of (178) with (180), which has the exact same form as (178), but is perfectly felicitous.

L & A's story makes use of the central insight on which their general account of temporal reference is based, viz that each new sentence or clause must be rhetorically related to what has come before. In particular, in each of (178), (179), and (180) the second sentence must be rhetorically related to the first. In the case of (179) there is nothing new to the matter: the relation is Narration, and so long as the event of the second sentence makes sense as the next event in the narrative after that of the first sentence (and no special considerations like the lone relevant to the interpretation of (152.b) intervene) Narration will be adopted, and with it the temporal order of the main events of the two sentences it implies. But for (178) and (180) this is different. Pluperfect sentences are rhetorically special (and this point puts L & A's account nearer that of Reichenbach than the 'past of a perfect' account) in that they must be related to their attachment anchors by a relation that is compatible with the fact that their eventualities are in the past of the main eventualities of their anchors. Not all of the rhetorical relations belonging to the repertory of L & A's theory are compatible in this sense. In particular, Narration is one of the relations that aren't. Explanation on the other hand is one of them and so is Elaboration. So it is from this restricted set of possible rhetorical relations that the interpreter of (178) or (180) must make his choice.

Intuitively, the relation that makes sense in both cases is Explanation. (L & A do not touch upon the question how members of the restricted set of rhetorical options are discarded, but let us follow them in taking it for granted that Explanation is the only possible candidate in the interpretations of these two discourses.) Why then does such an interpretation succeed in the case of (180) but not in that of (178)? Intuitively the reason seems clear, but for semanticists with the formal ambitions of L & A it poses yet another challenge. In the case of (180) assuming that first and second sentence stand in the relation of Explanation makes sense because it is possible to understand the spilling of the water as providing a causal explanation of the slipping: Max slipped on the water that he himself had spilled. But in the case of (178) justifying that the two sentences stand in the relation of Explanation doesn't seem to work. In what sense can Max's entering the room be understood as an explanation, causal or otherwise, of his pouring a cup of coffee? Somehow no good story seems to come to mind.

The matter is subtle, for what may seem to be only minor changes in content can turn the infelicitous (178) into a discourse that is acceptable, and on an interpretation in which the rhetorical relation between the two sentences is identified as Explanation. An example is (181).

(181) Max poured a cup of coffee. At last he had managed to find the door to the room with the coffee machine.

(181) suggests a scenario in which Max has been wanting to get himself some coffee, but being a large and new building he has trouble finding the room where he can get some coffee. The second sentence states the event that finally makes the direct realization of his desire possible.

It is a curious and remarkable fact that while we can interpret the second sentence as providing the kind of information that justifies understanding the two sentences as related by Explanation no such construal seems possible in the case of (178). Exactly what explains this difference is unclear to me, and L & A do not seem to have a clear answer either. They do note the problem, however, observing that the causal relations that we are willing to assume in order to justify an Explanation relation between a Simple Past sentence and a Pluperfect sentence may be much weaker than those that lead us to assume Explanation as the relation between two sentences both of which are in the Simple Past, as in (152.b). Thus compare (152.b), which we repeat with the 1Simple past + Simple Past' variant of (181).

(152.b) Max fell. John pushed him.

(182) Max slipped. He spilled a bucket of water.

The preferred reading of (152.b), L & A argue early on in their paper, is that where the falling is caused by the pushing. But without further context (182) seems poised between two equally plausible interpretations, that where Max slipped on the water he had spilled himself and that in which one of the concomitant effects of his slipping was spilling a bucket of water. In this case Narration is not overruled. It is an interesting question how on L & A's account both readings are in principle available in a case like this. Presumably this should be accounted for in terms of the interpreter's uncertainty whether or not there might be information that would overrule the Narration interpretation in a Penguin Principle configuration or his contemplating a derivation of the interpretation of the discourse from premises that include such information as well as one on the basis of premises from which such information is missing. But that is just one of the many further questions that a theory of the kind L & A propose would have to address eventually.

Summing up this last part of our discussion: The assumptions about possible causal connections between eventuality of different types that are part of our World Knowledge appear to differ in strength, which manifests itself in how we can, or in other cases refuse to, interpret discourses describing eventualities of those different types; and, it might now be added, there appears to be a remarkable degree of consistency between speakers as regards what the respective strengths of those hypotheses are, as something that manifests itself in the high degree of consistency with which speakers react to the discourses in question. I conclude by juxtaposing examples of such discourses; most of these have been reviewed in these comments, but not from exactly the perspective referred to in this paragraph.

(183) a. John pushed Max. Max fell.

b. Max fell. John pushed him.

c. Max fell. John had pushed him.

d. Max spilled a bucket of water. He slipped.

e. Max slipped. He spilled a bucket of water.

f. Max slipped. He had spilled a bucket of water.

g. John greeted Max. Max stood up.

h. Max stood up. John greeted him.

i. Max stood up. John had greeted him.

j. Max entered the room. He poured himself a cup of coffee

k. Max poured himself a cup of coffee. He entered the room. (?)

l. Max poured himself a cup of coffee. He had entered the room. (?)

(183.b) has been our paradigm of a 'Simple Past + Simple Past' configuration that a clear majority of speakers interpret as involving Explanation, overriding the default interpretation that would connect the two sentences as instantiating a Narration pair. The same prominence of a cause-effect relation between the event of the second sentence and that of the first renders interpreting the 'Simple Past + Pluperfect' configuration (183.c) as a pair involving Explanation entirely unproblematic. Moreover, the Simple Past + Simple Past' configuration (183.a), which we didn't considered so far, is interpreted straightforwardly as involving Narration. Understanding the second event as caused by the first isn't necessary for Narration, but it doesn't hurt.

The 'Simple Past + Simple Past' configuration (183.e), we noted above, is ambiguous between a Narration and an Explanation interpretation. Here the assumption that there could be a causal relation between the two main events is strong enough to allow for an Explanation interpretation, but not so strong as to clearly favor that interpretation. And a fortiori the assumption is strong enough to license an Explanation interpretation of the 'Simple Past + Pluperfect' configuration (183.f). An interesting twist to these examples is that a causal relation in the opposite direction, with the slipping being the cause of the spilling seems consistent with World Knowledge assumptions as well. For me personally (183.d) is ambiguous in much the same way as (183.e), and perhaps even with a preference for the Explanation interpretation.

For the Simple Past + Simple Past' configuration (183.h), L $ A's very first example, the Narration interpretation seems much preferred – that is their avowed judgment and I think they are unquestionably right in this. The same seems to be true for (183.g). Here too a Narration interpretation seems clearly preferred, at least so long as no special contextual information can interfere. But although the assumption of a possible causal relation between the two event types is so weak that Narration interpretations appear to be clearly preferred for both (183.g) and (183.h) nothing really stands in the way of such an assumption, so that an Explanation interpretation of (183.i) is possible.

In the progression represented by these four groups of three discourses each the last group represents the end of the line. We already reviewed L & A's observation that (183.i) is infelicitous and their reason for that: Here there is nothing in our World Knowledge Box that justifies assuming the kind of causal relation that is needed in support of Explanation, not even for the discourse (183.l) whose form earmarks that relation as the favored candidate. (183.j), we also saw, is fine, because Narration doesn't back-up from an independently motivated causal connection of the sort that is needed for Explanation. To be told in this form that Max entered the room and then poured himself a coffee is perfectly in order. A final twist is presented by (183.k), one of the discourses in (183) that we did not consider earlier. Here not even Narration is justifiable, and that in spite of the fact that it is comparatively undemanding. Apparently an event of a person pouring himself a cup of coffee doesn't even provide the kind of contingency basis for the event of the next sentence without which even Narration isn't a viable option. Exactly what that contingency relation is I am not able to say, but compare (183.k) with the seemingly quite felicitous 'Max poured himself a cup of coffee. He left the room.'

To deal properly with the general class of discourses exemplified in (183) is a task the importance of which L & A have recognized more clearly than perhaps anyone before or after them. But it is a task that even today,two decades after the appearance of L & A's 'Temporal Interpretation, Discourse Relations and Commonsense Entailment' remains largely unaccomplished.

4.8.2 Concluding Summary

As a final assessment of L & A's 'Temporal Interpretation, Discourse Relations and Common Sense Entailment' let me say this. There is no question that this paper, and the general approach of SDRT to the analysis of discourse it proposes have set themselves goals of formal precision and generalizability that are vastly higher than any approach before it, or, if I am right, any that has been presented as a serious competitor since. Setting yourself such ambitious goals has the great merit that it forces you to recognize complications and distinctions that otherwise would have been unlikely to emerge. And in this respect L & A's work has already shown its undeniable and very substantial fruitfulness.

This admirable sense of purpose, with the various empirical discoveries that it has helped us to make when attempting to apply the method L & A present in this paper should be clearly distinguished from what can be said about the actual solutions L & A present for the various puzzles they discuss and of which only some have been reviewed here. In my exploration of those solutions I have made no effort to tone down criticisms where it seems to me these solutions were incomplete, misguided or mistaken. In this latter respect the paper is not an unqualified success: There is much one can learn from it, but that is possible only by doing a great deal of the work oneself, and often that seems to include redoing the authors' work for them.

Arguably the most serious problem with the approach L & A advocate lies beyond the explorations of these comments: In the exact development of the system of non-monotonic reasoning that supports the inferential model of discourse interpretation L & A pursue, and in an exact articulation of how the system is to be applied in *complete* derivations of interpretations for simple and more complex discourses.

References

Asher, N. & Morreau, M. (1991), Commonsense entailment: A modal theory of nonmonotonic reasoning., *in* 'Proceedngs of Commonsense entailment: A modal theory of nonmonotonic reasoning.'.

Bach, E. (1986), Natural language metaphysics, *in* R. Barcan-Marcus & al, eds, ‘Proceedings of the Seventh International Congress on Logic, Methodology and Philosophy of Science’.

Beaver, D. (2004), ‘Linguistics and philosophy’, *Journal of Symbolic Logic* **27**, 1–53.

Clark, H. (1997), Bridging, *in* P. N. Johnson-Laird & P. C. Wason, eds, ‘Thinking: Readings in Cognitive Science’, Cambridge University Press, pp. 411–420.

et al, M. P. (2004), Centering: a parametric theory and its instantiations. NLE Technical Note TN-02-01, University of Essex.

Grosz, B., Joshi, A. K. & Weinstein, S. (1983), Providing a unified account of definite noun phrases in discoure, *in* ‘ACL 83 Proceedings’.

Grosz, B. & Sidner, C. (1986), ‘Attention, intentions and the structure of discourse’, *Computational Linguistics* **12**, 175–204.

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

Kamp, H. (1981*a*), ‘Evénements, représentations discursives et référence temporelle’, *Langages* **64**, 39–64.

Kamp, H. (1981*b*), A theory of truth and semantic representation, *in* J. Groenendijk et al., eds, ‘Formal Methods in the Study of Language’, Mathematisch Centrum, Amsterdam.

Kamp, H. & Reyle, U. (1993), *From Discourse to Logic*, Kluwer.

Lascarides, A. & Asher, N. (1993), ‘Temporal interpretation, discourse relations, and common sense entailment’, *Linguistics and Philosophy* **16**, 437–49.

Link, G. (1983), The logical analysis of plurals and mass terms: A lattice-theoretical approach, *in* R. Bäuerle, C. Schwarze & A. von Stechow, eds, ‘Meaning, Use and Interpretation of Language’, Walter de Gruyter, Berlin, pp. 303–323.

Mann, W. & Thompson, S. (1988), ‘Rhetorical structure theory: Toward a functional theory of text organization’, *Text* **8(3)**, 243–281.

Morreau, R. (1997), 'Fainthearted conditionals', *The Journal of Philosophy* **94**, 187–211.

Partee, B. (1973), 'Some structural analogies between tenses and pronouns in english', *The Journal of Philosophy* **70**, 601–609.

Sidner, C. (1979), Towards a Computation Theory of Definite Anaphora Comprehension in English Discourse, PhD thesis, MIT, Cambridge.