

The Time of my Life.

I.

English has the expression “the time of my life” that you use when somebody asks you about an event which you’ve been to – “How was the party”, “How was your holiday in Aspen?”, “Did you enjoy yourself”, “What was it like to see ‘Faust’ without anybody having any clothes on?” and so on. And when it was really, truly, exceptionally good, then you can reply “Oh I had the time of my life”.

That’s the answer I will give you when you ask me “What has it been like to be doing the stuff you have been doing for a living all those many, many years?”. I’ll reply “I had the time of my life”, and – here comes a bit of semantics - what I will be referring to when I say that will be – well – my life. There are cases, in other words, where the “the time of x” function -the function which maps every suitable x to the time of x - has a fixed point, and where, moreover, that fixed point is y’s life, for some, unusually fortunate, individual y – where the life of y is the time of the life of y. For all I know it is quite rare for lives to be fixed points of the ‘time of’ function; and those whose lives are among those thinly spread fixed points are people who have been very lucky indeed.¹

That I have been among those lucky ones has had more than one reason. One is that the things I have been doing over the years have never ceased to fascinate me. But a second one has been just as important. All these years - from the time when I became a graduate student at the University of Amsterdam in 1961 to this very day - those that I have had the chance and pleasure and honour to work with have enhanced the scientific experience – that of discovering arguments and that of discovering truth – rather than detracting from it through the kinds of niggardliness and meanness of which there is supposed to be quite a bit in science and academia – no less so, people keep telling me, than there is in other walks of life. So to all of you - those of you who are here and those who are elsewhere: thank you for having done so much to make my life into the fixed point of which I just spoke.

I would like to include in these thanks most particularly the three persons that have been the principal guides in my early attempts to find myself as a scholar (or scientist, whichever) and the memory of whom has never left me: Evert Beth, Richard Montague and Arthur Prior. All three of them died well before their time.

¹ You might accuse me of having sacrificed analytic sense to dramatic effect here. There really is no need to talk about fixed points of any function. All that’s needed is to distinguish between two ways of making sense of the description “the time of my life”. As part of the standard use of “I had the time of my life” it means something like “the best, most memorable, most wonderful of all the times that have been part of my life”. In the sense

I mean it here it rather means “all the time of my (professional) life” (i.e. “the duration of my (professional) life”). Thus analysed what I am trying to say “the whole time covered by my life has been identical with the most wonderful part of it”, amounting to, more or less, “it was all wonderful, from beginning till now”. But I like the version of the text better, even if it isn’t altogether coherent.

And we can only speculate how much more they might have done for science, scholarship and philosophy had they lived longer. (And here I am not thinking of anything like living to the ripe age of seventy.) Even though their lives were unfairly limited, the contributions they were able to make have been pivotal. But for me it was not only the beauty and importance of their work that was crucial, but also their personalities and the generous help they gave me - when I was at that vulnerable age where you have to discover both yourself and a niche in the world where you want to be for much or ever after. If it hadn't been for them, it is virtually certain that I would not have been standing here today - in the midst of you, with your scientific interests and passions and your contributions to our common endeavour.

When I speak of those to whom I feel a special gratitude, there is one other person who I want to mention here - someone who has shaped the scientific discipline to which I have given the bulk of this life - the theory of meaning in natural language - more, I believe, than anyone else, and who, through a unique combination of scientific acumen, fairness and uncompromised commitment to finding the truth, has made and kept semantics the serious, constructive and friendly enterprise and community it has been from its beginnings back in the late sixties and early seventies. For those of you who have lived through all or most of this period it will be obvious that I am referring to Barbara Partee. For four decades Barbara has been a kind of role model for me, and she has been that in more than one way. For one thing there are her many own scientific contributions and her exceptional gift of penetrating and transparent exposition. But of those I do not wish to speak here. For most of all she has been for me a model for how one should behave on public occasions, such as lectures and conferences: always try to come up with questions or comments that *advance* the arguments and claims of the speaker, rather than trying to put him or her down for the cheap thrill of winning a point. I know I have never come even close to how well she has, time and again, been able to demonstrate how this should and can be done. But I do think I have understood the message and that I developed a sense of the standards she has set. In any case she has succeeded in impressing many in her field in the same way as she has impressed me, and in our joint efforts to follow suit we have helped her in making the field both a better and a very much more agreeable one.


II.

I have chosen the title for this talk not just for the sake of making a somewhat belaboured pun. The real reason is that a preoccupation with time has been a kind of red thread that runs through my intellectual concerns from the time when I was a Ph.D. student at UCLA in the second half of the sixties to the present. And I want to use the opportunity that this occasion affords me - perhaps 'abuse' would be a better word' - to tell you something about that thread. I am well aware that this is a bit of self-indulgence, but my excuse is that the ways in which my preoccupations with time have changed as time itself went on are a reflection of how our thinking

about logic, language and mind has changed over the course of these years; and as I see it, those changes have been dramatic and profound.

When I came to UCLA in the fall of 1965 I had the exceptional good fortune that my very first semester there coincided with a visiting professorship of Arthur Prior. I took the upper division course on tense logic that Prior offered that one semester he was at UCLA. The material he presented was what went in to his subsequently published book "Past, Present and Future", and a good part of the course dealt with versions of the temporal logic (or "Tense Logic") that I believe is still most widely associated with Prior's name, the system with the two sentence operators P (for "Past") and F (for "Future"). The name "Tense Logic" (which I believe Prior had coined himself, but which in any case he seemed to be happy to use) indicates at least one of the things that this logic was trying to be: a formal account of the logic and semantics of the tenses, as they are found in English and other natural languages like it. The "P,F calculus" was modelled on modal logic, with the possible worlds semantics it had been given in the preceding decade, mostly through the work of the prodigy Saul Kripke. Like modal operators such as "possibly" and "necessarily" P and F were treated as sentence operators – operators which, roughly speaking, turn present tense sentences into corresponding past and future tense sentences. (Thus when the operator P is applied to the sentence "It is raining", we get the sentence "It was raining"; and when we apply the future tense operator F to this sentence, we get "It will be raining".) Formally, the sentences to which P and F are applied are analysed as "generic propositions". Generic propositions, in the sense intended here, can be identified with sets of times – those times at which the given sentence is true. Thus the input that "It is raining" provides to the tense operators P and F is the set of times at which it is raining; and the result of applying P to this input gives us another generic proposition, that expressed by the sentence "It was raining". This is another set of times – the set of all those times at which "It was raining" is true, or, in other words, the set of all those times that are preceded by some time at which it was (or is or will be) raining. (Thus, if it keeps raining ever so often until the end of time, then this generic proposition will be the set of all times. If that seems counterintuitive to you, then do not think there is something you are missing. *you* shouldn't be blamed; there is something here that should raise our suspicion that the theory doesn't quite work the way it should.)


[Picture 1]



When it comes to formalising tensed sentences of English, there is quite a bit that you can do with the P,F-logic – surprising perhaps, given how simple the logic looks. But the formalisations that it allows you to carry out are not particularly natural. And moreover, there are also many sentences that you cannot formalise at all with this logic. Among these – or so it seemed to Prior when he gave his course at UCLA – are sentences in which there is some reference to the *metric* of time, e.g. sentences which say that it was a certain amount of time ago that something was the case, and not simply that it was the case at some time or other in the past. So it was natural for Prior to raise the question how metric notions might be introduced into the P,F-logic: How should that be done, what needed to be added, or did we need a completely different logical formalism? And for one of the students in his class that question became the theme of his doctoral dissertation.

The ‘answer’ to Prior’s question that is proposed in this dissertation takes the form of a more complicated and powerful tense logic than the P,F logic, one in which P and F are replaced by two *2-place* sentence operators: S (for “Since”) and U (for “Until”). For any two generic propositions q and r , $S(q,r)$ can be read as “It has been the case that r ever since some time when it was the case that q ”, and, symmetrically, $U(q,r)$ as “It will always be the case that r until some time when it will be the case that q ”.

[Picture 2]



In the S,U logic we can approximate something like a metric of time if we assume propositions that are true “at regular intervals” - propositions like “This clock says “tick””, “The sun rises over the horizon”, and so on. For instance, let m be the proposition that is true exactly at times when it is midnight and at no other times. (So m is roughly speaking the generic proposition expressed by “it is midnight”; but we need not worry over how this proposition might be best expressed in English.)

Then the S,U-logic allows us to express that some other proposition q was true today (viz. before now but after the last time before now when it was midnight) as follows: $S(q \ \& \ \neg m \ \& \ S(m, \ \neg m), \ \neg m)$. In a similar way we can express that q was true yesterday (i.e. between the last and the one but last time it was midnight), or that it was true the day before yesterday, and so on. But the formulas get more and more complicated, and at a rate that seems to outstrip the complexity of the corresponding English sentences by a large factor.

There is a lot you can say in the S,U logic, much more than in the P,F logic. In fact, there is an important and natural sense in which you can express any temporal relation whatever in this logic (assuming certain constraints on the structure of time; but these constraints cover what is arguably the conceptually most likely option for the structure of time, viz. that according to which it has the structure of the real number line; among other things that is the time structure assumed in mathematical physics).

In this sense the S,U logic is (close to) optimal. But in another sense it is not, at least not when considered as a logic that tells you about the way in which temporal relations are expressed in a language like English. For even the translation of quite simple sentences can lead to the strangest contortions – we just had a glimpse of that. The translations that the S,U logic enables us to give may preserve the truth conditions of the translated sentences, but they make a joke of the idea that logical forms should also capture something of how the truth conditions of the translated sentences are determined by their grammatical form. (Curiously, it is precisely *because* of these bizarre contortions to which the S,U logic condemns us when we try to use it to provide logical forms for sentences of ordinary language that this logic has found a home in quite a different place than the one for which it was originally intended, viz. as one of the ingredients to a certain approach to program verification in CS, and, more recently, as an ingredient in the specification languages that have been adopted by the computer hardware industry for the specification of new chip designs.)

But, as I said, as a tool to do semantics of natural language the S,U logic is really no good at all. My own doubts on this point came to plague me soon enough. But I needed some time (as well as considerable heartache; the S,U logic was so *pretty*, I thought) to come to terms with this. And it wasn't just realising that tense logic was not the way to go. There was the more urgent question what could take its place. Brooding over that turned into a fairly drawn-out process, which ended, after about ten years, with the adoption of a quite radically different way of dealing with the semantics of tense (and of a good many other things besides).

There are quite a number of reasons why the S,U logic is no good for a theory of tense, and in retrospect they seem all so obvious that one has difficulty putting oneself into the frame of mind of someone who could ever have failed to see them. (But then we should not forget - this is one instance of the more general moral I am trying to get across in this memoir - that those were early days, when Montague

Grammar, and with it the revolutionary insight that a precise formal semantics could be given for substantial fragments of natural language with the model-theoretic tools of mathematical logic, were only just beginning to sink in.) Perhaps the most obvious of those reasons is that in natural language the tenses often work in tandem with temporal phrases like “yesterday”, “next week”, “at five”, “on the first of November”, “three days after their wedding” and so on. Surely, if there are any phrases in natural language that can be said to *refer* to something, then these phrases are among them, and it is hard to think of them as referring to anything other than times. In fact, such adverbial phrases do not at all seem to work like temporal operators in the way that the tenses are assumed to work in tense logic (and that it was arguably not all that unreasonable to believe they did work). But if tense logic – be it the P,F-calculus, the S,U-calculus or any other such system – is to be the logical form formalism for temporal reference in natural language in general, then *all* devices that contribute to temporal reference must be cast in the mould of tense operators, including the apparently referring temporal adverbs. But for these the mould is nothing short of procrustean.

It is much more natural to treat these phrases as what they appear to be so unequivocally, viz. as bona fide referring phrases. But that has two consequences. First, the “temporal logic” that we get in this way is one with constants and variables that range over times and predicates which denote relations between them like the “earlier-later” relation or the relation of temporal inclusion. That kind of logic is essentially the kind that we have had since 1879, when Frege published his *Begriffsschrift*. In comparison with new-fangled systems like tense logics that is actually, you might say, a simplification (though for just that reason it may be not as much fun for the working logician).

The second consequence is that since times are now among the entities that are referred to by some of the terms of our logic, their existence and basic properties should be beyond suspicion – times ought to be among the “first class citizens of our ontology”, as it sometimes put. And that is something that many have seen as a problem (which is perhaps the main reason why Tense Logic seemed so attractive to some). For time is something quite abstract, and there has been no end of speculation within philosophy about what time “really is”, and what its distinctive properties are. These worries were first raised in antiquity and for many of us they are as much alive now as they were then.

Moreover, there are good reasons for thinking that when we want to give an intuitively plausible account of how tenses and other devices of temporal reference function in natural language, it isn’t just times that we should acknowledge as first class citizens. When we look into the workings of natural languages a little more closely yet, we can see that adopting just times as new first class citizens is not going to be enough. Events and states are needed as well.

That events and states must be part of a viable account of temporal reference in natural language is another one of those observations of which it is hard to

understand, once they have been made, that they could ever have been overlooked. It seems so obvious nowadays that the role of most verbs is to describe events of various kinds – actions that people perform – getting up, lifting their arms, eating buns, going to work - or things that happen to them - getting born, falling in love, getting married, falling into a trap, being hit by a falling branch, dying – or things that just happen, like explosions, landslides, the rising of the sun - and on and on and on; and the verbs that are not used to describe events – *know, love, own, stand, be*, and so forth - serve to describe states. (One way to see the inescapability of the conclusion that this is what verbs do is to consider deverbal nouns – “explosion”, derived from “explode”, “knowledge”, derived from “know”, “dying” or “death”, derived from “die” and all the nominal gerunds we get by attaching “-ing” to any verb whatever. All evidence suggests that these nouns denote events or states, and that the events and states they denote are the very same as the ones described by the verbs in which the nouns originate.

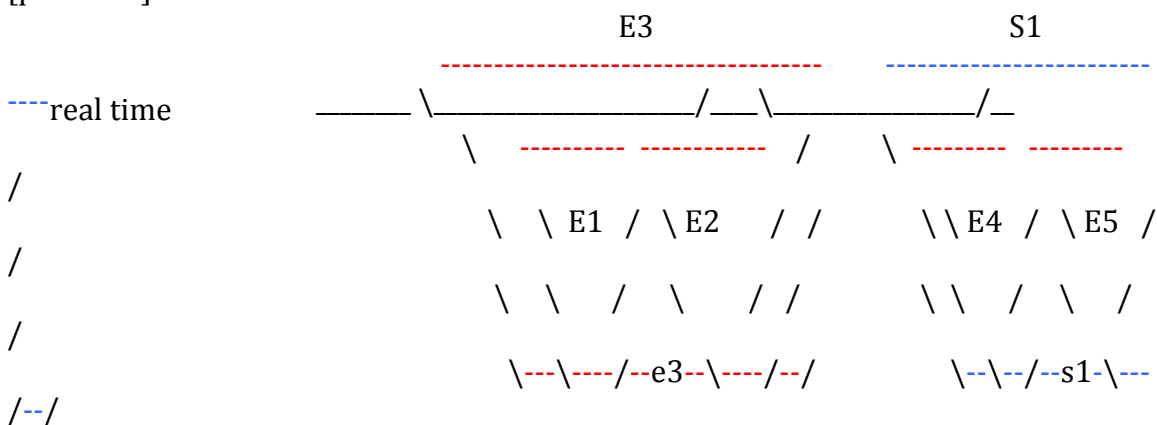
In many ways events and states seem more problematic entities than times. What kinds of event are there and how many? What fundamental relations are there between events, in particular: How are events ordered in time? Can one event be a part of another event? And in what sense (or in what different senses) can events be parts of other events? Do events (or at least some events) have a spatial location, and if yes how is that location determined? And these are just the first and most obvious questions. States raise similar questions and then there are also questions about the relations between states and events (which prove to be especially important for semantics) and about the relations in which events and states stand to times.

None of these questions have easy answers and some of them may well seem beyond the giving of definite answers altogether. And so they should make us wary of adopting events and states as entities that our logic can talk about. Indeed, it looks as if we ought to feel much more squeamish about accepting a logical formalism which talks about events and states than about one which talks about times. But it all depends on your philosophical point of view – or ‘*parti pris*’, one might be inclined to say. Bertrand Russell (and his one time student Norbert Wiener, who later became the founder of Cybernetics) saw things differently. Times, Russell observed, cannot be perceived as such. Time is a ‘construct’, which we humans have developed in order to think and talk and theorise about events in a systematic and coherent way, something without which there could be neither theoretical physics nor the way in which we ordinarily talk about the world as it presents itself to us: as a network of causally interlocking actions and happenings. And, moreover, if time is the kind of medium that it is usually assumed to be in physics, with the structure of the real line, and where the individual temporal instants correspond to the individual real numbers, then times are far too short to be registered by us in any direct way – and even that is only a way of speaking, for in fact the instants themselves have no duration at all.


Events, Russell thought, were different in this respect. Events *are* elements of our experience and so are certain relations between events, among them the temporal relations of precedence and overlap. (In fact, during the relevant part of his life Russell came near to assuming that events were the *only* elements of our experience.) Moreover, Russell and Wiener were able to prove that, starting from a family of events with the relations of precedence and overlap between them, you can construct a linearly ordered structure of temporal instants, in which each of the events from the event family from which the construction starts can then be embedded: it either occupies a single instant of the new time structure or else some interval made up of such instants.

The Russell-Wiener construction can be used to throw light on a problem that is quite different from the application Russell himself had in mind. Assume - as in fact we have already done - that verbs are used to describe events and states. Then a piece of non-fictional discourse, e.g. an eyewitness report of an actual event, but also bit of narrative fiction like some segment of a novel or short story, can be analysed as a description of a collection of temporally ordered events and states. When we subject this collection to the Russell construction we get a “discourse time structure”, just as we get a structure of “real time” when we apply the construction to the totality of all real events – those which make up the complete history of our world. This discourse time will be of course much smaller, and also much ‘coarser’, than the ‘real’ time structure that the construction extracts from the full history of our world, which includes all actual events whatever. More specifically, the discourse time induced by a truthful eyewitness report will be a coarse-grained image of some part of real time – that part that is spanned by the events reported in the discourse: There is a “homomorphic” map from this segment of real time onto the time structure of the discourse, together with a mapping that associates each discourse event and state with the real event or state that it represents. What we get in this way is a many-one correlation between real times and discourse times: single discourse instants correspond to intervals of real time; and discourse events that occupy a single instant of discourse time are the images of intervals of real events that occupy stretches of real time.

[picture 3]



discourse time



The appeal of this picture is that it gives a new sense to terms that had already been in use for many years (antedating the advent of formal semantics by several decades) in linguistic descriptions of the behaviour of tenses, especially of the past tenses of the Romance languages. For instance, in that literature one finds characterisations of the difference between the French Passé Simple and the Imparfait along the lines of:

The events described with the help of the P.S. are *punctual*, whereas the events or states described with the help of the Imp are *durative* (that is, they last through some extended period of time).

One cannot help feeling that there is a sense in which this description of the difference between PS and Imp gets at a distinction that is in some sense “right”. But what is that sense. In particular, what could be the sense in which a PS event can be “punctual”?

Indeed, how should the relevant opposition between “punctual” and “non-punctual” events be defined? One of those who had been exercised by this and similar questions was Christian Rohrer, then the professor for Romance linguistics at this University. Christian thought that the new methods of model-theoretic semantics might be the key to an answer here. And when I came over to Stuttgart in the summer of 1978 (he had invited me for a few weeks), he put these questions to me, hoping that someone who had been spoon-fed Montague Grammar (at the time when even the spoon was being invented in tandem with what it was intended to feed), might be in a good position to crack these questions.

However, model-theoretic semantics as such is not the key to this problem. The intended models of model-theoretic semantics are either real worlds, with real events and real time, or copies of those worlds, at a more abstract level but without leaving anything out from the worlds they are copies of. The events of such a model - the events that we are confronted with in ordinary life, and that we talk about when reporting on some real life episode - are typically *not* punctual in the sense of real time. Punctual events in that sense would be far too short for us to be able to register them directly, just as we cannot register directly the temporal instants that would be their ‘durations’. On the other hand, there is a legitimate sense in which the events that figure in our talk *can* be seen as punctual; that is the sense of the *discourse time* which is determined by the discourse that speaks of those events - a time structure that built from the (sparse) substratum made up by those comparatively few events that figure in this discourse.

Indeed, when we analyse bits of French narrative prose along these lines, then we find that, by and large, PS events (that is, events introduced by sentences whose verbs are in the PS) are the ones that come out as punctual in the sense of discourse time, whereas the events and states described by sentences in the Imp come out as non-punctual. Admittedly, the correspondence is only a rough and imperfect one. But it is systematic enough to suggest that this way of accounting for the supposed “punctuality” of PS events is on the right track. And if it is on the right track, then that has important implications. For it suggests that this is what human interpreters do when they process narrative bits of texts: they construct some kind of (small scale) *model* of the world as the bit of text describes it, a model in which (“model”) events and states play a central part and within which they create their own temporal substratum.

And the account suggests something else, and which is equally important. What explains why it is typically PS events that come out as punctual in the sense of discourse time, and not the events or states described by verbs in the Imp, has to do with how the signals these tenses provide for connecting the events or states described by the verbs to which they are attached to events or times that have already been introduced into the discourse model. And, more often than not, the events or times with which the new event or state is to be connected with introduced into the discourse model by an earlier sentence, and not by the one in which the given verb is contained. Thus the semantic contributions of PS and Imp reach beyond the limits of the single sentence and can be captured only within a framework in which the unit of analysis is not the sentence, but coherent discourses and texts. (And this is equally true of the other tenses of French and, more generally, of tenses in languages that with well developed systems of tense marking.

Establishing such connections requires two things. First, an event or time has to be selected from the discourse model as is to serve as the ‘anchor’ for the new event or state, i.e. as the element with which the new state or event is to be connected. And, secondly, the nature of the relation – is the new event or state located after its anchor in time, or is it before, or do the two coincide? – has to be determined as well. The details of how the different tenses accomplish this are quite complex, and in large part this is because the tenses themselves are not the only players in this. As a rule several other factors are involved as well. One important factor, and one that is notoriously hard to describe accurately with the tools that current linguistic theory provides, are the rhetorical relation or relations between the new sentence and the one or ones preceding it. The importance of the role that rhetorical structure plays in the determination of temporal relations was already prominent in work going back to the eighties, e.g. by Webber. It was given centre stage in the work of Asher and Lascardes in the early nineties, who extended their insights about rhetorical structure into a much more comprehensive theory of discourse interpretation than the one that I am about to introduce here (their Segmented Discourse Representation Theory, or SDRT, for short. But here I am running ahead of my story).

Another factor, more readily accessible to currently available tools of analysis, is *Aspect*. When the verb of the new sentence describes an *event*, then this event is often understood as following the event or time that has been reached so far in the narrative, though other relations (precedence, inclusion) are also possible, depending on the rhetorical relations between the new sentence and the one or ones immediately preceding it. When what the verb describes is a *state*, then this state is normally understood as holding at the anchoring event or time. This is one important way in which the aspectual distinction between eventive and stative aspect - or between perfective and imperfective aspect as it is also called - manifests itself. That is, aspect does not just show itself in the structure of the event complex that is described by a given sentence - whether the complex consists of just a single event, or a single state or some combination of events and/or states - but also the influence it has on *inter-sentential* temporal relations. Furthermore certain tense forms - the French Passé Simple and Imparfait are among them, but a similar distinction can be observed between for instance progressive and non-progressive tense forms in English - carry aspectual as well as temporal information, and thus contribute to the fabric of events, states and times that the discourse projects in two ways: through what they say about temporal location of the given event or state and through their telling us whether what is described is an event or a state.

A description of these matter that is as brief as this one could not possibly reveal the true complexities of all this. But it should suffice for the two points that I am most concerned to get across.

- One is that when we interpret a piece of discourse - or a single sentence in the context in which it is being used - we build something like a *model* of the episode or situation described; and an important part of that model are its event structure, and the time structure that can be derived from that event structure by means of Russell's construction.
- And the second point is that if we want to capture the semantic role of the tenses of the verb, we have to make explicit how they help us to connect the events or event complexes described by the verbs to which they are attached with events or times contributed by earlier sentences.

Once you have convinced yourself that this is how the tenses work, the next question you cannot help asking is: Is this discourse-related, sentence-transcending function just something special about tenses, or is it a more general fact about natural language? Do we find a similar capacity, for establishing inter-sentential semantic relations also with other, sentence constituents, and with constituents that do not have any obvious connection to time?

It was an attempt to show that that is in fact the case - that sentence-transcending effect are not an isolated feature that distinguishes tenses from other grammatical constituents - which led, in the summer of 1980, to the first explicit formulation of what has come to be known as *Discourse Representation Theory*. As a matter of fact,

time and tense played no part in this first formulation, which was only concerned with questions of quantification, reference and anaphora. (It focussed on the so-called *donkey pronoun phenomenon*, a puzzle that had been presented quite a few years before by Peter Geach (Geach, 1962) but that had just seen some kind of revival at that time. (The revival of Geach's donkey problem can perhaps be seen as an indication that natural language semantics had by then reached the point of maturity that enabled us to see how vexing a problem Geach's problem really was.) But a version of DRT that deals with donkey anaphora and tense at the same time appeared the very year that saw the appearance of the first temporality-free formulation. And the interaction between temporal and other aspects of interpretation have been a central part of the DRT-agenda from that time until the present, leading to a growing awareness of how closely temporal and non-temporal aspects of meaning are intertwined as our understanding of natural language semantics has progressed. (Though this is an aspect of natural language semantics that is still often underestimated. Even today there is a tendency among semanticists to ignore or play down the pervasive role that temporality plays in (at least) the languages of Europe and to thereby rid themselves of the additional complexities that we have to face when these interactions are given the attention they deserve.)

III

Those were the early eighties, and here ends the more dramatic part of the story I wanted to tell today. But it is not - by any means - the end of the story. For what concerns me personally, it is true that a good part of what I have been doing over the three decades that followed could be described as consolidation and extension - consolidation and extension of the basic assumptions about the nature of meaning, and about the relationship between language and logic, that emerged from the considerations which led to DRT in its original form. But that does not mean that from that time onwards there haven't been significant changes in perspective. First, there are new avenues that have been pursued by others, who have modified these basic assumptions or taken them further. In an objective survey this is the part that should have got the lion's share of what time and space remain. But today's is an essentially egocentric exercise, so I will - in full awareness of the imbalance - keep my remarks about general developments to an improper minimum and focus on the more closely autobiographical.

Of particular importance for the development of my own views on the relation between language and logic has been the development within a DR-theoretic framework and perspective of the theory of presuppositions. (Most of the main players in this development, Van Der Sandt, Geurts and Beaver, are here today.) Then there is SDRT, already mentioned above. SDRT is a theory that has taken on a life entirely its own, but one in which the DR-theoretic foundations are nonetheless still clearly visible, now as components of much more powerful and complex package, that, you might say, *really* deals with *discourse*. Other important

developments that can be traced back to the origins of formally explicit dynamic perspective are Reyle's Underspecified DRT (Reyle, 1993), the strongly computationally motivated lambda-DRT of Pinkal and the people around him (developed in part with an eye on underspecification as well) and, within the last few years, the "Layered DRT" developed by Geurts and Maier as a unified framework for dealing with a range of quite different phenomena, each of considerable complexity in its own right: attitude reports, but also aspects of language that for a long time had almost totally escaped the attention of formal semanticists: quotation, appositives and parentheticals. I wish I could say more about all and each of these. But I am afraid we must press on and leave these developments for another occasion.

But there is one development, also due entirely to others, that has affected my own work and thinking more than any other and to which I want to devote a few more words here. This is the integration of (ideas from) DRT with *Dynamic Semantics*. When using the term "dynamic Semantics" here, I am thinking in the first instance of Dynamic Semantics as it was developed by Groenendijk and Stokhof and then, in many variations, by others (most of them from a Dutch production line that for a number of years seems to have been operating around the clock). One of the original motivations behind Dynamic Semantics was to offer an alternative to DRT that made the same predictions about pronominal anaphora (including, in particular, donkey anaphora) but without the representational commitments that had been claimed to be indispensable to an adequate account of those phenomena in DRT's first published version. The anti-representational stance of the original 'Dynamic Semantics Program' never appealed to me all that much. Even if the predictions of the first version of DRT can be achieved in a non-representational way (or at any rate in a way that is less explicitly representational), that wasn't the *true* motivation behind the DRT enterprise, let alone the main or only one; I hope that is something I have been able to make clear in this talk. (It is true, though, that Dynamic Semantics in the style and spirit of Groenendijk and Stokhof has been in part responsible for the 'crypto-representational' attitude that for many years I have quite often adopted in public, downplaying DRT's essentially representational stance and stressing its usefulness as a toolbox for describing a growing range of linguistic phenomena.)

But there is also another side to Dynamic Semantics. To provide a kind of alternative to DRT's representations, which it considered unnecessary and to be dispensed with, Dynamic Semantics has, building on some of the ideas that are already clearly recognizable in the work of Heim and in Veltman's Update Semantics, bequeathed us with a cluster of interconnected new concepts of compelling beauty and great usefulness. Arguably the most important of these notions, and also probably the best known ones, are the concept of an *Information State*, that of a *Context Change Potential* as well as, in their wake, so to speak a dynamic definition of entailment. These concepts have proved powerful tools in formal semantics, even where the goal is just to come up with theories that assign natural language sentences the intuitively correct truth conditions. But they are also important for intrinsic conceptual reasons; it must be especially gratifying to those to whom we owe them

that at long last their importance is becoming widely recognised and put to good use within the philosophy of language and in philosophical logic.

The integration of Dynamic Semantics and DRT has been achieved through the work of a number of people. Arguably, the first and most decisive step was taken by Reinhard Muskens, but others – Jan van Eijck, Paul Dekker and (omnipresent as mentor, tireless as penetrating commentator and through his own work) Johan van Benthem have been among them – have played an important part as well. In a way this is an integration of an unusual kind. It is an integration of two systems which started out as competitors and which resembled each other in the predictions they made, while differing dramatically in their over-all architecture and underlying philosophy. Both systems are ‘dynamic’, but they are that in very different ways. In DRT the dynamics is a dynamics of interpretation: it is the incremental character of the construction of Discourse Representation Structures (DRT’s semantic representations). But the representation *formalism* - the formal language in which the semantic representations (or logical forms; there is no difference here) of sentences and texts and discourses are cast - has a classical, and thus static, model-theoretic semantics. In contrast, in Dynamic Semantics it is the model-theoretic semantics itself that is dynamic – it is here that notions like Information State and Context Change Potential belong and make their crucial impact.

In the integrated theories I was referring to, the “formulas” are DRSs (from some suitable “DRS language”, chosen to suit the analysis of some particular set of linguistic phenomena). As always in DRT, the formalism is used to build discourse representations incrementally, so the theory is dynamic in this ‘interpretational’ sense that is distinctive of DRT. However, in these integrated theories dynamics is found twice over; for in addition to the incremental construction of DRSs as text and discourse representations, the model-theoretic semantics of the formalism is dynamic as well, this time in the sense distinctive of Dynamic Semantics: The model theory tells us how formulas of the formalism, that is DRSs, transform given information states into new ones, intuitively by incorporating into them the discourse contents they represent.

To date the most sophisticated integrated systems of which I am aware are those developed and applied by Maria Bittner and Adrian Brasoveanu. Some of these systems are tailor-made to the particular languages whose syntax-semantic interfaces are described with their help, which constitutes another important development of the dynamic approach to semantics: both through their form (e.g. through the particular semantic operators that are included in them) and through the ways in which these forms are employed in accounting for sentence and utterance meaning, these systems afford us new insights into what is specific to the syntax-semantics interface of particular languages and, after stripping away the particular, into what different languages have in common. And by continuing along this path, we may ultimately come to see more clearly what is common to all. If there are such things as semantic universals, this is probably the most likely way we will discover what they are.)

But what about my own part in all of this? I do not want you to come away from this event with the impression that all those years that have passed since the early eighties I have just been sitting there, treading water. For one thing, though I can be accused of having been slow in realising the importance of new accomplishments by others in some cases (and that includes some of the ones mentioned above, I did realise quite quickly the importance of the work by Van Der Sandt et al. on presupposition that became available in the early nineties, and that has deeply influenced my understanding not just of the nature of presuppositional phenomena, but of the relation between form and meaning in natural languages generally and also of the role that formal logic can and should play in theories of natural language semantics. And I think that in emphasising the importance of *presupposition computation* – i.e. the problem how linguistic presuppositions can be identified, and must be identified before the question of their satisfaction can be answered (or even raised) - I have added a further dimension to this new perspective. Indeed, when trying to compile a little list of topics where I feel I have made some headway over the last two decades, presupposition is one of the topics that should be included. Here follows the list, with presupposition as first item.

- There are three issues in presupposition theory to which I believe I have helped to contribute something new. The historically first issue has to do with what it is for a presupposition (triggered by some word or construction in a given sentence) to be 'satisfied' in the context in which the sentence is used. For long the prevailing view had been that presuppositions are either satisfied wholesale or else must be accommodated wholesale. Against that we (Antje Rossdeutscher and I, see also the next item on this list) argued that the task of an interpreter who is confronted with a presupposition has to be able to *justify* it – a term chosen to capture intuitively what the interpreter is meant to do, but that the account we proposed must make formally explicit. In general presupposition justification takes the form of a mixture of 'verifying' and 'accommodating': some of the material needed to see the presupposition as entailed by the context is found in the context as is, but there may still be something missing to complete the entailment, and that is then the bit that needs to be accommodated. The case where a presupposition follows without any accommodation whatever and the one where all of it is to be accommodated are limiting cases of the more general process of presupposition justification, in which verification and accommodation complement each other.

The second point is the one already mentioned: before the question of presupposition satisfaction can arise for an interpreter he must know what the presupposition *is* whose satisfaction is at issue. In almost all of presupposition theory the solution to this problem is taken for granted – much as semantics, in the sense of a syntax-semantics interface that outputs a first semantic value or representation, is taken for granted in most of pragmatics. But the computation problem, I claim and have tried to show, is *not* all that trivial.

One of the reasons why it isn't trivial is directly related to the third issue that has been prominent in my thinking about presupposition phenomena. Presupposition is extraordinarily widespread in our uses of natural language; not only is it the exception rather than the rule the rule for natural language sentences to be without presuppositions, it is very common for a sentence to generate not just one presupposition, but two presuppositions or more. It is multiple presuppositions, often triggered by constituents of other presuppositions, that render presupposition computation a surprisingly intricate matter.

But it is not just the computation of presuppositions that is complicated by multiplicity; the problem of presupposition justification is made much more complicated as well. When a sentence comes with a bunch of presuppositions, generated in the same or in different parts of it, then presupposition resolution turns into something that is much like simultaneous constraint solving, with the different constraints often pulling in opposite directions. Investigations in this aspect of presupposition theory are still at an early stage. But from the little we know it appears that the resolution of such conflicts is often governed by strict systematic principles.

As far as I can see there is plethora of hard and interesting work to be done on both these aspects of the multiple presupposition phenomenon, the computation of multiple presuppositions and their simultaneous justification.

- A second area where I feel there has been progress in my own thinking is that of lexical semantics. My work in this domain has been largely in collaboration with Antje Rossdeutscher, a collaboration which started shortly after we both came to Stuttgart in the late eighties. In recent years that work has taken a new turn, through a systematic attempt – mostly Antje's – to uncover the internal syntactic and semantic structure of individual words, and the systematic connections between the structure and meanings of different words that are built from the same roots. I have come to see this work increasingly as promising and important, and while I am not the one in the driver's seat in this venture, I am proud to be part of it. The work is important, I believe, because it opens up a new territory that has been off limits for formal semantics far too long. There is a wealth of structure below ground, so to speak, that we never bothered to look for because we were so busy marvelling at all the wonderful structures that language allows us to erect from the ground upwards. True, to discover word structure requires digging, but there is much to be discovered by the digger who persists. That there is all this structure to be found there has long been evident to morphologists – in fact, that is to a large extent what morphology is about. And it shouldn't come as a surprise to anybody; for without the internal structure through which many words are systematically linked to each other learning the vocabulary of one's language would be an even greater challenge (and chore) than it is. Surely, the internal structure and organisation of words must make it easier to require and retain them, even if we still do not really seem to understand precisely how that helps.

The work we are doing on lexical structure also intrigues me because it promises an over-all architecture for the syntax-semantics interface that reaches from roots all the way to the riches of sentence and discourse meaning.

- Thirdly there is the work on the representation of propositional attitudes and, more generally, of complex mental states, and building on that on a semantics of “attitude reports”. The central intuition that informs this work is that when people think about, and talk about, the mental life and activities of others or their own, they are as often as not thinking or talking not just about single attitudes – single beliefs, single doubts, single desires, single intentions – but about several attitudes at once. These attitudes, moreover, are often connected with each other in various ways, and a particularly important type of connection is that where two or more attitudes are ‘internally coreferential’: they are targeted on the same entity or entities in a purely psychological sense, which is independent of whether a common target exists outside of the mind. The almost exclusive focus in both linguistics and philosophy on single sentences of the pattern ‘x believes/doubts/desires/ intends that φ ’ has detracted from this evident complexity of thought and mental representation, and from the incessant preoccupation with it on the part of us, ordinary talkers and thinkers. In particular it has been an obstacle to a proper appreciation of most of ordinary language about the mind.

My main aims in this area have thus been (i) to develop precise formalisms for the representation of complex, internally connected mental states; and (ii) to use these formalisms in a semantics of a much broader range of propositional attitude ascriptions than have usually been considered hitherto. And not just ascriptions in which combinations of connected beliefs, desires, intentions and the like are attributed to a given agent, but also ascriptions that trace how such complex attitudinal states develop over time – i.e. how people learn, forget or change their opinions, goals and plans – as well as those which attribute the same or related attitudes to more than one agent. A formal system in which such complex attitudinal states and descriptions of them can be represented was put in place to this end in the second half of the nineties. (On the face of it this looks like a quite simple extension of standard DRT, but its semantics and logic are quite intricate.) More recently this system has come to be used in a range of linguistic applications (among them in lexical semantics, where intentionality plays a much bigger role than I realised for many years). Also, during the past couple of years, in which I have returned, for part of my time, to the professional environment of an actual philosophy department, this formalism and extensions of it have become indispensable ingredients to most of what I have been trying to contribute to current issues in philosophy of language and mind.

These applications have led to a number of further assumptions about the structure of mental content. The single most important one of these is that the mind is populated not just with “thoughts” – beliefs, desires etc. – all of which have *propositional* content, but also, and no less importantly, with *entity representations*. Entity representations are just that – representations of particular entities: persons,

physical objects, places, times, events and so on. They represent the entities they represent as having certain identifying properties, or as standing in certain relations to each other, but often also as standing in certain relations to the representation and its representor themselves. They can serve as constituents in the representations of propositional contents, and by entering into several propositional contents at once they can create the kind of referentially connected network of thoughts I alluded to above. The assumption that a large part of mental content is constituted by 'libraries' made up from such entity representations has proved essential to accounts of how reference and related aspects of meaning are transferred from speakers and authors to their audiences. I am becoming more and more convinced that such communication-oriented accounts are a necessary complement to semantic theories which endeavour to analyse linguistic meaning as a user-independent commodity.

- Finally, returning to the subject of time, that too has remained a concern to this very day, but with a gradual shift in the direction of more specifically linguistic problems. In this regard my work has followed a general trend within the semantics of tense and aspect – that of looking closely into how the different constituents of language that contribute to the expression of temporal relations and event structure interact. One thing that has become increasingly clear to me – and the section on Friday that was devoted to tense and aspect² has reinforced this conviction – is that on the whole tenses are not the main target in this domain. There are many languages that make do without tenses, but it seems that aspectual structure plays an important part, and “aspect operators”, which serve to impose aspectual structure (or “event structure”) on the semantic contributions of lexical verbs, seem to be playing an important role in all languages, even if their importance varies. Somehow, and helped by the aspectual distinctions that their languages allow them to make, speakers of languages without tense systems seem to be able to handle temporal relations between simple or complex events remarkably well – about as well as the speakers of languages with morphologically explicit (and often elaborate) tense systems like English.

Much of this work has something to tell us about the different ways in which human languages cut the cloth of information that speakers want to convey to their audiences – into one part that gets explicitly expressed and another that is left to the audience to reconstruct, by combining what has been expressed explicitly with whatever relevant information that is already in its possession. If your language has a rich repertoire for expressing aspectual distinctions (or for expressing other factors that relate to temporal location of events, e.g. types of evidentiality), then you need few or no devices for directly and explicitly expressing temporal relations; if you are a language with an elaborate tense system, you can skip on explicit realisation of aspect and/or evidentiality. We are only beginning to see glimpses of what the general repertoire is of possible devices from which individual human languages make their respective selections and of the possible trade-offs that are

² Of the workshop (3-5 Sept. 2010) at which a version of this talk was delivered.

available to a language that must satisfy the basic requirements for useful and effective verbal communication between creatures like us - with our cognitive capacities and our practical and intellectual needs. But, echoing a remark I made earlier, it is here, I believe, that the true insights into semantic universality and diversity should be sought and can be found.

These newer developments and the perspectives on basic conceptual issues they open up form a striking contrast with the tense-logical perspective and technology that I started out with in this aperçu - those that stood at the cradle of the semantics of time and tense, but were then kicked out of the nursery room fairly soon after. A striking example of this difference can be seen from what linguists have been saying in recent years about the words *since* and *until*. Their analyses do not exactly *refute* what a tense-logical system like the S,U-logic tells us about the meanings of these words. But the message they carry is clear: what the S,U-logic can teach us about the ways in which these words make their semantic contributions to sentences of English is very little at best. (I say: 'at best' because, at least in my own experience, leaning on formalisms like the S,U-calculus in doing semantics of natural language is likely to get you off on the wrong foot.) Maybe there is some truth in the idea that without the support of formalisms like these we would never have got where we are now. But if that is so, then that only goes to show that sometimes progress really does require kicking away the ladder that got you where you are.

##

Most of what I have been saying will have been hard or impossible to understand for many of you who have come to this final meeting of the workshop. But I hope that at least one thing has come across: It isn't just that a great deal of progress has been made over the past forty years on matters of detail. More importantly, as part of the efforts we have made, and essential for those efforts to be sustained and go forward, there has been a radical change in point of view. Perhaps the most important aspect of that perspectival change - at least for me - is the change of the relation between semantics and formal logic. That change fits well with a more general change in our understanding of the role of logic in science, philosophy and cognition. During the period I have been talking about we have seen an abundance of new "logics - formal systems with a well-defined syntax and model- or proof-theoretical characterisations of entailment - and this multitude keeps growing almost by the day.

Confusing as this plurality is for anyone - and it must be that especially for someone who is coming to the field of logic now - it is not a sign that logicians do not know, or that they know less and less, what they really want. We are no longer looking for the single 'right' logic, the *Characteristica Universalis* that Leibniz saw as the ultimate goal of the science of logic; or at least, that is no longer our only, or even our principal endeavour. Most of the new logics are not intended as competitors for the Leibniz prize. Rather, they are intended as *models*, at some level of abstraction, of information and information processing of various sorts. In many cases they are

models of the workings of some part or aspect of human cognition, or, by much the same token, the workings of some part or aspect of one or more of the languages we speak. But not all that is modelled in this way is tied to human cognition. In computer science logics are used to model, again at some suitable abstraction level, aspects of computation and deduction (as a species of computation) not necessarily as they are performed by humans, but as they are, or ought to be, executed by real or virtual computing machines. One important implication of this general development is that the traditional animosity between formal logic and human psychology is a thing of the past. And instead of that we are now faced with a fundamental question: Exactly what are the relation or relations between logic and cognition? This is an immensely difficult, multi-faceted and difficult question. But it is no longer one that you will be dismissed (in any of its senses) for bringing up.

On a more local note I add that from this new point of view there is no fundamental difference between the formal systems developed within the logic community and representation languages developed for the purpose of doing formal semantics, such as, in particular, the DRS languages that have been defined within DRT. The only difference, if any, is in the sort of things these systems are trying to capture.

Indeed, we can look at the P,F-calculus and its cousin the S,U-calculus from such a perspective too. But then we have to include the question: How good are these systems for what they are used for, or are meant to be used for? And when it comes to the question what they were meant and designed to be used for originally - and for which we tried to use them for a little while - then, I have argued, the answer can only be negative: These logics just aren't not very good models of the behaviour of tenses or, for that matter, of any other devices for temporal reference that we find in natural languages. Nor do they seem to hold much promise as models for how people handle time in thought. But as I noted, there are other, non-cognitive uses for logic too, and if these particular logics are still useful in a non-linguistic and non-cognitive setting, then, well, luck has been with them.

IV.

I started out with a quite artificial use of the expression "I had the time of my life". I conclude with its standard use. If you ask me - later today, or later this week, or later this year, or at any time I will still be around for the asking of it - what I thought of this workshop, my answer will be - well, by now you can guess what it will be, and not just because you can recognise a rhetorical ploy as well as the next person:

Yes, I *did* have the time of my life. And for that the ones to be thanked are all of you who are *here* - all of you, but most of all of course the speakers who turned these three days into a genuine intellectual adventure of the first order, including my predecessor just now, who showed us again that you can say something truly important and - one hesitates to use the word in an address like the one I am giving, but there is no better - *deep*, while yet being entertaining and funny.

And most of most of all the three who thought of this event in the first place: Klaus and Arndt and Torgrim; and to Arndt and Torgrim thanks for doing all the work behind the scenes (some of it on the proscenium, but that was only a tiny fraction of it), to get us here and to look after us the way we have been.