Chapter 28
Arguing About Dynamic Meaning

Martin Stokhof

Abstract Whether, and if so in what sense, dynamic semantics establishes the need to move away from standard truth-conditional semantics, is a question that has been discussed in the literature on and off. This paper does not attempt to answer it, it merely wants to draw attention to an aspect that has hitherto received little attention in the discussion, viz., the question what role we assign to the use of formal systems in doing natural language semantics.

28.1 What Dynamics?

The question whether dynamic semantics constitutes a move away from static semantics, and if so, what that move involves and how it should be justified, has been discussed on and off since the first dynamic approaches appeared in the early 1980s.

Some have argued that the introduction of dynamic notions in semantics is superfluous and that whatever is treated by dynamic theories using their characteristic conceptual apparatus can be treated with equal descriptive adequacy by static semantics, and with greater explanatory success because it draws on more standard conceptual resources.

But ever since the advent of theories of dynamic interpretation and dynamic meaning, their proponents have tried to make the case that this development does constitute a legitimate, even necessary move beyond the truth conditional conception of meaning of the logical and philosophical traditions that had been one of the sources of inspiration of formal semantics in the late sixties, early seventies of the twentieth

Martin Stokhof would like to thank Johan van Benthem and the editors for their comments and their patience.

M. Stokhof (✉)
ILLC/Department of Philosophy, Universiteit van Amsterdam, Amsterdam, The Netherlands
e-mail: M.J.B.Stokhof@uva.nl

A. Baltag and S. Smets (eds.), *Johan van Benthem on Logic and Information Dynamics*, Outstanding Contributions to Logic 5,
DOI: 10.1007/978-3-319-06025-5_28,

century. Let us briefly review a couple of examples of the arguments that were employed.

Hans Kamp, in his seminal paper 'A Theory of Truth and Semantic Representation' [20] which introduced the framework of discourse representation theory (DRT), claimed that in addition to the truth-conditional concept of meaning, which focusses on the reference relation between expressions and external entities, there is a second conception, stemming from psychology, linguistics, and artificial intelligence, that is concerned with the articulation of 'the structure of the representations which speakers construct in response to verbal inputs.' Noting that these two conceptions, stemming as they do from quite diverse disciplinary traditions, had been separated for quite some time, Kamp states that:

> This separation has become an obstacle to the development of semantic theory, impeding progress on either side of the line of division it has created. The theory presented here is an attempt to remove this obstacle.

So the aim of DRT is to combine referential and representational notions of meaning. In DRT this is done basically by making semantic interpretation a two-step process. When processing a sequence of utterances or a text, the listener/reader creates a formal representation of its content by building a so-called 'discourse representation structure' (DRS), in which score is kept of what is being talked about, what is being said about the various entities, and how various pieces of information are related to each other. This complex representation then is assigned a truth conditional interpretation by defining an embedding of the DRS in a model theoretic structure, i.e., into the kind of model that is familiar from standard truth conditional semantics.

In this way both the representational and the referential aspects of meaning are being accounted for within DRT. What is important to note in the context of the questions that are central to this paper is that there is not one, unified concept of meaning that accounts for both aspects. Rather there are two different processes that are integrated, not conceptually, but at the level of the overall theory. One is the incremental build-up of a representation, and the other process is the non-incremental, 'in one fell swoop' embedding of the resulting representation in a model. The former is, sensu strictu, not a process of semantic interpretation. The latter is, but it is quite a standard one, at least in terms of the semantic concepts it employs. So, it seems appropriate to call DRT 'a dynamic theory of interpretation', rather than a semantic theory that incorporates a dynamic concept of meaning.

In a similar way, one of the main sources of inspiration and motivation of Heim's file change semantics (FCS; cf. [18, 19]), which was developed to give a non-quantificational account of definite and indefinite expressions, is a procedural take on how such expressions function: not by referring to something in the world, but by making available so-called 'discourse referents' (a notion that goes back to work by Karttunen in the late sixties of the last century). And Heim, like Kamp, ends up with a combination of a dynamic component that builds representations, so-called 'files', and a static, truth conditional semantics that interprets them:

> Roughly, the model of semantics that I am going to present will embody the following assumptions. The grammar of a language generates sentences with representations on various

> levels of analysis, among them a level of 'logical form'. Each logical form is assigned a 'file change potential', i.e., a function from files into files. […] The system moreover includes an assignment of truth conditions to files. Note that logical forms themselves are not assigned truth conditions, only files are. Only in an indirect way, i.e., via the files they affect, will logical forms be associated with truth conditions.

Thus FCS, like DRT, postulates a division of labour in semantics: semantic interpretation is (minimally) a two-step process. Assuming that the assignment of logical forms is part of the syntactic component of the grammar, the semantic component consists of two parts. First, it associates with each logical form an operation on a specific type of representation, the files. Thus a logical form itself is associated with a dynamic construct, viz., a file change potential, which, however, is not directly semantic in nature: it constructs a new file from an input file, but these files are themselves not semantic objects. They need to be interpreted and that is what the second component does: it assigns static truth conditions to files. So the dynamic aspects are accounted for in an indirect manner. It is not meaning as such that is dynamic, but rather the process of interpretation.

This indirect approach sets such theories as DRT and FCS apart from approaches that attempt to deal with dynamic aspects directly, by building them right into the concept of meaning itself. An example of a theory that introduces a concept of meaning that is different from the traditional, truth-conditional one, is update semantics, as developed by Veltman [31–33]. In the opening paragraph of 'Defaults in Update Semantics', Veltman refers to the standard definition of validity in terms of truth preservation, and then goes on to characterise his own approach, that of update semantics, as follows:

> The slogan 'You know the meaning of a sentence if you know the conditions under which it is true' is replaced by this one: 'You know the meaning of a sentence if you know the change it brings about in the information state of anyone who accepts the news conveyed by it.' Thus, meaning becomes a dynamic notion: the meaning of a sentence is an operation on information states.

Similar quotes can be culled from other papers, e.g., [13, 15, 17]. In the latter paper, it is emphasised that the dynamics of meaning may affect various aspects of a situation, not just the information states of the speech participants, thus proposing that meaning be analysed in term of 'context change potentials'. This is of some importance as it signals that these later systems progress beyond the initial conception in crucial respects. And it also (partly) explains why in a footnote that occurs in the passage just quoted Veltman relates the dynamic conception of meaning, not just to the work of Kamp and Heim, but also to earlier work by Stalnaker, which focusses very much on information update.

As for Heim and Kamp's work, as we just saw there is a subtle, yet principled distinction between the concept of dynamic interpretation of DRT and FCS, and the dynamic semantics of which Veltman's update semantics is a specimen. The difference, as we indicated, resides in the concept of meaning itself. As for Stalnaker's work of the 1970s [25, 26], that is motivated in yet other ways, it seems. Stalnaker is primarily concerned with the analysis of assertion and presupposition, and focusses on

the role these play in conversations of a particular kind, viz., information exchanges. Cf. the following quote from 'Assertion':

> […] acts of assertion affect, and are intended to affect, the context, in particular the attitudes of the participants in the situation

Thus, assertion as a speech act is a dynamic entity. But the concept of meaning, in particular, the concept of the content of an assertion, remains a static entity.[1] For Stalnaker, what an assertion contributes in terms of content is a proposition, i.e., a set of possible worlds, and the context changing effect is accounted for, not at the level of meaning, but at the level of speech acts.

Thus what we have here is a third kind of dynamic theory. The work of Kamp and Heim can be characterised as concerned with dynamic interpretation of linguistic structures. The approach of Veltman c.s. is involved with the development of a dynamic conception of meaning. And Stalnaker employs a static notion of meaning in what is basically a speech act level account of dynamics. Thus what we have here are three different notions of what 'dynamics' in the context of natural language meaning might mean: dynamic assignment of static meanings, in the Heim and Kamp case; dynamic meaning as such, as in Veltman's update semantics; and dynamic employment of static meanings, as exemplified by Stalnaker's approach. If one would look closer at the literature, one would presumably find even more variations than these three, but in order to set the stage for the main question, this should suffice.

That main question is whether, and if so in what sense, these are really rival theories. There are empirical issues involved here, of course, and conceptual ones, but first let's look at the issue from a theoretical perspective.

It would appear that the differences between theories of dynamic interpretation and theories of dynamic semantics are centred around methodological questions concerning the internal organisation of grammar. The issue of representationalism and compositionality is a good example of such a methodological consideration. DRT and FCS adopt a level of representation in grammar that is different from both syntactical structure and meaning proper, and that mediates between form and meaning. Having such an intermediate level of representation in the grammar implies that a strong form of compositionality (often referred to as 'surface compositionality') no longer applies: it is not (structured) expressions that are interpreted directly, i.e., 'as is', but representations that are built from them in an incremental way.

The differences between dynamic semantics and the Stalnakerian approach are not concerned with the organisation of grammar, but rather seem to focus on the concept of meaning as such. Consider the difference between a dynamic approach, such as Veltman's, and Stalnaker's account of information change. The difference is subtle, but real, nevertheless. It basically comes down to this: is information change something that is brought about using an expression that has a static meaning, or does it reside in the meaning of the expression itself? Dynamic semantics takes the latter route, and it does so unequivocally. Cf. the following quote from [15]:

[1] Of course, the content of an assertion itself is a context-dependent entity, in many cases, but that does not turn it into a *dynamic* one.

> The general starting point of the kind of semantics that dynamic predicate logic is an instance of, is that the meaning of a sentence does not lie in its truth conditions, but rather in the way it changes (the representation of) the information of the interpreter. The utterance of a sentence brings us from a certain state of information to another one. The meaning of a sentence lies in the way it brings about such a transition.

The difference with Stalnaker's view is that the latter regards information change as an external effect of the use of expressions that have meanings that themselves are perfectly static. Information change is a pragmatic effect brought about by static semantic means. In that respect the Stalnakerian view builds on a standard hierarchical view on the relation between semantics and pragmatics that goes straight back to their traditional semiotic characterisations. The dynamic view departs from that view in two ways. It no longer subscribes to semantics as the study of the relation between language and the world, with the associated referential and truth-conditional conception of meaning. And consequently, it draws the line between semantics and pragmatics differently.

From this perspective DRT, FCS, and their kin are something of a mixed bag. They 'side', so to speak, with dynamic theories in regarding dynamic aspects as part of semantics, but they locate them in the process of building representations. These are then interpreted in a standard, truth conditional way, and in that respect these approaches are more Stalnakerian than Veltmanian. Dynamic semantics can be regarded as a kind of straightening out of these issues: by redefining semantics (and implicitly redefining pragmatics) it eliminates the need for the kind of representations that are characteristic of DRT and sundry systems. One of the original motivations for the development of dynamic predicate logic was exactly this: the elimination of what was regarded as an unnecessary and insufficiently motivated complication in the grammar.[2] This centred essentially around the wish to maintain a particular, strong form of compositionality. But of course one might argue that there are independent reasons for having those kinds of representations as part of the semantics (and thus for turning strong compositionality from a methodological principle into an empirical issue).

Be that as it may, one core issue now appears to be whether natural language meaning is better modelled in the standard way, i.e., in terms of a static truth-conditional concept of meaning combined with a pragmatic theory that accounts for dynamic effects such as information change, or in the dynamic way, by constructing meaning in terms of context change potential.

One may try to answer that question in two ways: by an appeal to empirical considerations, and by conceptual arguments. From the empirical perspective, one might reason that the standard picture is standard for good reasons, and one would need to change sides only if there are there empirical phenomena concerning natural language that really can only be accounted by embracing a dynamic picture. From the conceptual perspective, things appear to be less constrained: when one tries to determine what counts as a convincing conceptual consideration, empirical adequacy is an obvious necessary condition, but it does leave room for other considerations.

[2] Cf. Groenendijk and Stokhof [15, Sect. 5.2].

Perhaps it leaves too much room, as it is not obvious that if from the empirical perspective all things are equal when it comes to the static–dynamic choice, there are indeed any decisive conceptual considerations to base a choice on.

28.2 A Step Back

But before we enter into considerations concerning a choice, we'd better ask a preliminary question: does it really matter? Is there a difference between these two conceptions that is worth investigating to begin with?

One of the decisive steps forward in the development of dynamic semantics for natural language was made from a logical perspective. In a number of papers,[3] van Benthem established dynamic semantics as a subject matter in its own right by identifying a proper meta-theoretical framework for studying its properties. Where Boolean algebra provides the general mathematical framework in which standard semantics can be formulated and important meta-properties can be studied, van Benthem showed that relational algebra plays a similar role for dynamic theories. It provides a general framework in which concrete dynamic theories can be studied and compared. As a general meta-theoretical framework it is not confined to dynamic systems used in natural language semantics but also provides the tools to study similar approaches in logic itself, in cognition, artificial intelligence and the like. This kind of inquiry into the formal, meta-logical properties that are characteristic of various dynamic systems was taken up by a number of authors.[4]

One example to illustrate the kind of concerns that are at stake here. A central question is what exactly distinguishes static and dynamic systems from each other. One way to go about answering that question is by providing a formal characterisation of what makes a system static. Usually this is done in terms of formal properties of the updates, i.e., the operations that take states into states, that the system makes available. It turns out that there are several ways to do so, with slightly different consequences for what counts as static and what not. In a recent study [23] Rothschild and Yalcin have traced the history of these attempts in great detail, so what follows is just a very brief illustration, and the reader is urged to consult the Rothschild and Yalcin paper for further details.

van Benthem provided a first definition of staticness.[5] According to this characterisation a system is static if its updates are eliminative and (finitely) distributive. If states are sets of some kind, these properties come down to the following: an eliminative update results in a state that is a subset of the state to which it is applied, and a distributive one is an update that works 'point-wise', i.e., its effect can be defined

[3] Among others, [2, 3], and the papers collected in [4].

[4] For natural language semantics we should mention, among others, Vermeulen [34], Visser [35]. Cf. also [10] for a computational perspective on deduction in dynamic semantics, and [7] for a more recent overview.

[5] In van Benthem [1].

in terms of its effects on the singleton elements of the state on which it operates. Both properties are needed to maintain staticness. An illustration is provided in [14] where it is shown that Veltman's update semantics is dynamic because its updates are not distributive, while maintaining eliminativity, where for DPL it is the other way around: the updates of that system are distributive, but lack eliminativity.

A generalisation was provided by Veltman [33]: where the van Benthem characterisation has Boolean algebra as its backdrop, Veltman uses the more general concept of an information lattice. Staticness is the defined in terms of the properties of the lattice: if a system's information lattice satisfies idempotence, persistence, strengthening, and monotonicity, it is static.

Rothschild and Yalcin take this meta-theoretical approach another step further. They focus on what they call 'conversation systems', which abstract away as much as possible from particular features of the language under consideration and its associated semantics, and talk only about states and the operations on states that the semantics induces. A general characterisation of staticness is then given as follows: a conversation system is static if and only if the associated state system satisfies idempotence and commutativity. They show that this characterisation encompasses both that of van Benthem and that of Veltman, and that it can be used to prove the non-staticness of various systems, such as FCS, DPL, and update semantics.

That different characterisations can be found in the literature is explained by how close one stays to a specific system or set of systems, with less general characterisations allowing for more fine-grained analyses. A case in point: as we just saw, the original van Benthem characterisation of staticness in terms of eliminativity and distributivity allows us to not only classify both DPL and update semantics as non-static, but also to differentiate them in an informative way. Using the more general approach of Rothschild and Yalcin that possibility disappears: from their perspective both systems are non-static because they are both neither idempotent nor commutative. But this is, of course, the usual trade-off between generality and specificity.

Be that as it may, for our present purposes what is important is that as far as formal systems and their properties are concerned there is substance to the distinction between static and dynamic systems. However, that still leaves the question open whether from the point of view of the semantics of natural language the distinction makes sense as well. That we *can* describe natural language meaning both in a static as well as in a dynamic manner, using the appropriate formal systems, and that such descriptions differ in the meta-logical properties of the systems, employed, does not imply that we *need* to do so. So the question remains, but we can be certain that it concerns a substantial distinction.

28.3 Fact or Fiction?

As should be clear from the way in which various theories in the broad realm of 'dynamics' are introduced and motivated, there is general agreement that when we observe language in its actual use it is abundantly clear that this has dynamic effects.

This is not something that anybody would want to deny. The 'dynamic wave' that started in the 1980s then marks, minimally, an increased attention for such dynamic aspects. In that sense all the various approaches discussed above are part of the same development. But what is under discussion is, first of all, to what extent these dynamic effects need to be accounted for in a linguistic theory, and, second, if so, in what manner. The first question is answered positively by many,[6] on the basis of quite comparable argumentation. It is to the second one that answers starts to diverge.

So it is particularly the latter question, i.e., the question where and how in an overall account of language it is that we should account for the dynamic effects of language use, that is central to many of the discussions in the literature. Obviously, this is closely connected to the question where we are to draw the line between semantics and pragmatics. Now suppose we start with the traditional semiotic characterisation of semantics and pragmatics, with semantics being concerned with the relation between language and the world, and pragmatics with the use of language. One reason that this is a good starting point is that it is a relatively theory-independent description, one that is stated in terms that are neutral and descriptive. Now, by phrasing the phenomena in question as 'effects of language use' we might seem to have settled already on an answer. For if we go by the semiotic characterisation, it would seem obvious that dynamic effects, described as effects of language use, should be accounted for in pragmatics. This seems indeed to be one way of settling the matter.

Of course, it's not always that straightforward, if only because there are many alternative ways of describing what semantics and pragmatics are concerned with. By way of illustration, let us look at the following passage from a recent paper [22] by Karen Lewis:

> Dynamic contents encode (some of) the effects of an utterance on an arbitrary input context. By contrast, static contents do not encode any updates to the context. On a static view, the effect(s) of content on the context has to be explained pragmatically. *These are fundamentally different sorts of explanations. Semantics describes facts about natural language.* Pragmatics, on the other hand, describes facts about rational agents who engage in co-operative activities. [emphasis in original, ms]

Clearly, Lewis has in mind a particular conception of semantics and pragmatics, and of the associated the division of labour between the two, that is different from the traditional one. First pragmatics. On the one hand the description given here is much broader than the traditional one, as it does not mention language or language use and language users, but talks about agents in general. On the other hand, if we narrow down 'agents' to language users, we get a conception that is much narrower as these language users are now restricted to rational ones that engage in cooperative activities. Of course, language users do, at least sometimes, act as rational agents, and they do, again at least in some situations, engage in cooperative activities. But it will hardly do to try to force any aspect of language use into that restricted mould.

---

[6] It is a testimony to the impact of the generative tradition, though, that even today many authors would seem to work with a more or less principled distinction between 'language-as-a-system' and 'language-as-use', which echoes the competence—performance distinction that Chomsky used to define the proper domain of linguistics as a scientific endeavour. Cf. [30] for further discussion.

Rather, the conception of pragmatics that is at stake here is one that goes back to Grice, who used it in the execution of a quite specific, philosophically motivated program. Thus this conception of pragmatics is not one that we can appeal to in order to classify a certain set of phenomena, in this case the effects of language use that dynamic theories are concerned with, without an independent motivation of why pragmatics should be (only) this.

Then semantics. The crucial question here is whether the characterisation of semantics that Lewis uses, viz., 'semantics describes facts about natural language', is sufficiently specific to rule out a dynamic conception of natural language meaning. Taken literally, it does not seem to do so. If we analyse the meaning of a certain expression in natural language as consisting in a context-change potential, in what way do we go beyond describing 'facts about natural language'? As was already noted, it is not that there is some theory-independent way of identifying what those facts are, that we can appeal to in order to answer this question.

It appears that we need additional considerations if we are to conclude, as Lewis intends to, that the dynamic effects that we are concerned with here can not be part of semantics, but must be accounted for in pragmatics. Adopting a classical, truth-conditional and static semantics as that which 'describes facts about natural language' will do the job. And it is a way of looking at things that has a venerable ancestry, given that it is the most prominent account of the 'language–world' relationship that the semiotic conception claims semantics is concerned with. But, and this is crucial, it can hardly be appealed to as an argument. It is a stipulation, one that may be justified in a number of ways to be sure, but it is not in and by itself a move that has argumentative force.[7]

So, it seems we're stuck: an appeal to prior characterisations of what semantics and pragmatics are is unlikely to be both sufficiently restrictive and theory-neutral to allow us to reach a decision as to whether the relevant effects of language use should be accounted in one or the other. And per implication that means that along these lines we will not be able to adjudicate the question whether dynamic semantics is a bona fide theory of natural language semantics. So what are we to do?

One obvious suggestion would be that we have been barking up the wrong tree all along in looking for conceptual-methodological arguments to decide the issue, and that we rather should go back 'zu den Sachen selbst'. After all, natural language is an empirical phenomenon, and so is its semantics. Therefore, shouldn't we be able to conclude on empirical grounds that dynamic semantics is on the right track, or that static semantics is the empirically adequate characterisation of what natural language meaning is?

To be sure, in the conceptual-methodological considerations that one can find in the literature, empirical arguments are deemed relevant as well. And that appears only natural if only because the theoretically motivated preference for dealing with

[7] And we would do well to note that it is not that even if we accept the semiotic characterisation of semantics as a neutral starting point: what 'the world' is, is left underspecified in that characterisation, and there seems to be no a priori way of ruling out that information states of language users are part of 'the world'.

certain phenomena in semantics, or rather in pragmatics, comes with the obligation to show that it actually can be done in that way. And that needs to be shown, then. But do note that there is a matter of fact here only if we have fixed the format of a semantic and a pragmatic theory. Which means that we turn around in circles: we can appeal to empirical considerations only if we assume we have reached consensus on the conceptual issues. And the latter, it seems, can not be obtained solely by conceptual considerations but time and again steers us towards the empirical.

The question of empirical adequacy has been discussed in the literature in some detail.[8] But what is important to note is that apparently it is not easy to come up with 'hard evidence', i.e., with a phenomenon of which we can show that it can, or can not, be accounted for in a particular way. In fact, if we look back at, e.g., the discussion between Kamp and Groenendijk and Stokhof in the late 1980s, we note that the argument was never about empirical coverage *per se*, but about accounting for a set of empirical phenomena in a particular way.

Another illustration of the fact that empirical and theoretical motivations come as a mixture, is provided by Cresswell in [9]. He takes DPL as his point of departure and then develops an alternative, in the sense of a theory that has the same empirical coverage as DPL, that is static, i.e., truth-conditional. The details need not concern us here, what is relevant for our discussion is the way in which Cresswell characterises his own enterprise:

> My purpose in these last two sections has not been to adjudicate between the use of double assignments, as in Groenendijk and Stokhof, and 'namely' variables [which are the new logical tool that Cresswell introduces, MS] , but simply to point out that the translation scheme shows that there is no empirical difference between the two approaches.

This is quite representative of a lot of work that has been done in this area. One takes a fixed set of phenomena, a given account of them (static or dynamic), and then develops an alternative account (dynamic or static) that has the same empirical coverage. That is interesting and revealing, but the key question in the present context is: what does it tell us about the choice between dynamic and static accounts of meaning? Does it tell us anything at all?

According to Cresswell it does. In the passage just quoted, he continues to draw the following conclusion:

> And since the use of free variables does not constitute a departure from the standard truth-conditional account of meanings, then neither do the empirically equivalent dynamic theories of semantics.

But that seems a non sequitur. As was mentioned above, one of the main elements in the development of dynamic predicate logic was to show that DPL and DRT account for exactly the same facts. Yet, neither proponents of DRT nor those of DPL subsequently claimed that therefore the two theories are somehow the same. On the contrary, the very fact that both theories were able to account for the same phenomena focussed the discussion on their conceptual differences, and on the justification of

[8] Beside the older literature that has been referred to above, cf. e.g. the more recent [8, 11, 21, 24]; and [22] (already mentioned).

their respective conceptual and methodological assumptions.[9] In other words, when comparing theories we can not look *just* at their empirical coverage, we also need to take into account their conceptual apparatus.

But then it seems we are indeed back at square one. Conceptual considerations seem to be unable adjudicate because the relevant concepts—of semantics, meaning, pragmatics, and the like—are intrinsically theory-dependent. And empirical considerations don't really help us out because obviously empirical equivalence (or empirical difference for that matter) simply does not say enough. So we are in a conundrum.

In such a case, it is best to take a step back, and to take a closer look at what kind of enterprise we are dealing with here. That involves a lot of issues and a great number of considerations, The issues are complex and we can not hope to do justice to all of them. So in what follows we just want to point to a particular aspect that hitherto has received little or no attention in the discussion: the role of formal systems.

28.4 Another Take

What appears to be a central, though not too often explicitly thematised, factor in how one adjudicates the issue, concerns how one views the role of formal systems in natural language semantics,[10] By 'formal system' we mean here minimally a formal language plus an explicit model-theoretic and/or a proof-theoretic account of its logic. The 'logic' part can be more or less extensive, depending on what the system is meant to capture.[11] It should be noted that especially in linguistic applications, the specification of the logic is often subdued, which explains the tendency to phrase these issues in terms of the role of 'formal languages', rather than 'formal systems'.

Let us start with natural language semantics. As we have argued elsewhere [29], there are (minimally) two main perspectives on the role that formal systems have to play in natural language semantics. One perspective is that a formal system is primarily a tool, something that is used in formulating and evaluating a theory. But there is an earlier, and arguably still dominant perspective on which a formal system is viewed as a model for a natural language with its semantics. Here the central features of the formal language employed are supposed to model similar features of the natural language that is being described. On this view, the task of the semanticist is twofold: first to find a formal language which has the required properties; and second,

[9] As was already mentioned earlier, compositionality played a key role in that discussion, and it is interesting to note that there are indeed good arguments that compositionality is not an empirical issue, but a methodological principle. Cf. [16] for more discussion.

[10] Similar considerations apply to formal systems in other domains, e.g., in the kind of naturalistic philosophical analysis that is exemplified in dynamic epistemic logic. We can not go into these matters here, but cf. e.g. [6] for discussion.

[11] Another explanation is that in the early days of generative grammar, natural languages were primarily studied from a syntactic point of view, often in terms of structural properties familiar from the theory of formal languages.

to devise a systematic relation between natural language and formal language so that, in the relevant respects, the latter can be seen as going proxy for the former.

It is interesting to note that Stalnaker, in [27], construes a major difference between DRT and DPL precisely in these terms: the latter, but not the former, models natural language. Now, contrary to the suggestion in [27, p. 4], it is obvious that the proponents of dynamic semantics do intend these to be part of a semantic theory of natural language, i.e., to be combined with a systematic account of how syntactic structures are mapped onto representations in a formal language. To give just one example, Dynamic Montague Grammar [13] closely follows the lead of Montague's original set-up and gives an explicit definition of the mapping that takes natural language expressions (or rather, their derivation trees) into expressions of intensional logic, which then is interpreted in a dynamic way. From that perspective, Stalnaker's claim (*loc. cit*) that:

> Dynamic predicate logic, on the other hand [i.e., in contrast to DRT, MS], is only indirectly relevant to any natural language. It defines an artificial language with new kinds of dynamic variable binding operations, obviously different from anything in natural language, but presumably intended to model, approximately, some of the devices used in natural language.

is decidedly odd. First of all, as the references just given show, dynamic semantics is explicitly intended to be part of a systematic theory. And secondly, the use of an artificial language with elements that lack direct counterparts in natural language applies to *all* formal semantic theories that implement indirect interpretation, be they static or dynamic. It definitely applies to DRT as much as it does to DPL, but it also applies to static theories, including those that implement Stalnakerian ideas.

But the quoted passage is also interesting for another reason, viz., because of the view on the role of formal systems it appears to assume, or, rather, the ambivalence with regard to this role. On the one hand, Stalnaker maintains that the extended variable binding of the existential quantifier in DPL is 'obviously different from anything in natural language'. This is true, but in a rather trivial sense. On the other hand, the introduction of these devices is 'presumably intended to model, approximately, some of the devices used in natural language'. Forget about the 'approximately' for the moment, what is interesting is that apparently, for Stalnaker there is a tension in using a concept or structural property of a formal system that has no direct counterpart in natural language to model some aspect of that very same natural language. But this is confusing, as it appears to put unrealistic constraints on the relationship between a formal system and what it models, in this case a natural language.

Different views on to the role of formal systems are connected with different views on what a semantic theory should do, on what the relation is between a semantic theory and what it is a theory of, viz., the semantics of natural language. The modelling approach constructs the explanatory force of a semantic theory in terms of the successful modelling of (part of) natural language semantics by means of a formal system. The other main perspective, which views formal systems as tools, defines the task of a semantic theory as *describing* relevant aspects of the semantics of a natural language, and on the basis of such descriptions providing explanations of various regularities. On this view, the use of formal systems is akin to their role

in other sciences, such as biology, or physics.[12] Here the primary criteria are expedience, in addition to economy, elegance and simplicity, but *not* similarity, in terms of structure and concepts. If we view semantic theory as a descriptive-explanatory enterprise, we would allow the theorist in principle the use of any formal system that gets the job done. Descriptive adequacy is the first and most important criterion. Of course, when two or more equivalent descriptions are available, other considerations become relevant, but similarity, it seems, is never a pre-condition, and hence not a principled consideration to judge semantic theories.

This appears uncontroversial. But there is a danger lurking here as well, for it would seem to follow, on this view of what the task of a semantic theory is, that we should be able to identify the facts and features that we intend to describe independently of the means that we bring to the task at hand. As we have seen above, that is a dangerous assumption: there may be facts in the sense of their being systematic patterns in judgements about entailment, or synonymy, or analyticity, but there certainly are no facts that can be classified as 'semantic' or 'pragmatic' independent from conceptual and methodological assumptions. And that implies that the key issue is whether the choice of a formal system can be made in terms that are not informed, one way or another, by such assumptions.

Of course, the alternative modelling approach is not better on this score, on the contrary. If we assume that the formal system that we use in our semantic or pragmatic analyses somehow models the relevant aspects of natural language because they share core features, we risk to loose any explanatory force. For this type of modelling can be considered adequate only if we have independent access to the relevant features: we can only judge whether some formal system accurately models aspects of natural language if we can compare them in the relevant respects. And that means that we have to be able to access the features of the formal system and those of natural language independently from each other.

With a formal system that is, of course, not a problem. We can investigate such a system and ascertain its properties. And we can, of course, simply design a system in such a way that it has the properties we want. In the case of a natural language, however, the situation is completely different. If a natural language is an empirical phenomenon then it is what it is, and has the properties that it has, quite independent from our access of them, and even quite independent of their accessibility.

Given this asymmetry, the use of formal systems in an descriptive–explanatory role in natural language analysis is puzzling. To put it bluntly: what is the point? If we need to have independent access to the relevant features of natural language that we want to model with a formal system in order to be able to decide whether the formal system is an adequate model, then what do we stand to gain? The answer here is not 'Nothing', since we can use formal systems as models in useful ways: to provide concise overviews of features that we are interested in, to come up with 'perspicuous presentations' of them. But one thing such models can not be, and that is descriptions with explanatory power.

[12] Of course, there are many differences as well, but these are not relevant for the main point that is at stake here.

To summarise: it seems that a formal system is regarded either as a tool or as a model, i.e., it is used either to describe or to express[13] features of natural language. In the first case the internal conceptual structure of the tool is of secondary importance (which is not to say that it is of no importance at all): what counts is whether it gets the job done, and it is first and foremost on those grounds that particular conceptions are justified (with other considerations coming in only as secondary). That is a pragmatic justification, one that comes without much ontological implications: that we can successfully describe a certain range of phenomena using a certain concept does not commit us to the existence of anything corresponding to the concept, not even after dutiful application of considerations of economy and simplicity. On the second view, however, the modelling one, this is essentially different: there the conceptual structure of the formal system is supposed to align with that of the natural language. But, and this is the crucial point, it can do so only by design, so to speak. We already need to have a good grip on the nature of what we want to model in order to be able to find a system that 'fits'. But explanatory value such a fit does not have, at least not as long as there are no independently, empirically motivated constraints on the formal systems that we can regard as candidate models.

So, if natural language semantics is an empirical discipline, one which has descriptive-explanatory goals, the proper perspective on the role that formal systems might play is a thoroughly pragmatic one. Formal systems are tools, selected first and foremost for their ability to get the job done, i.e., for the descriptive power that they provide. The tools in and of themselves don't need to have any essential characteristics in common with what they are applied to.

28.5 So What?

The consequence of these considerations for the issue that is at stake in this paper, viz., how to decide whether dynamic semantics is correct in claiming that natural language meaning is a dynamic concept, is straightforward: we can't decide the issue in a remotely theory-independent way. This is illustrated by the simple, but significant observation that the description of dynamic effects can be done in static terms, but the modelling of something dynamic has to be done in terms of something that is itself dynamic. These two perspectives are mutually exclusive since, as we argued above, they do not depend on any 'fact-of-the-matter', but represent choices to do things one way rather than another.

This does not make all discussion pointless, of course. Not all formal systems have the same descriptive power, so if we want to describe certain dynamic effects in static terms we still need to find the right system to do that. Analogously, if we want to model these effects we need a dynamic system that has the right dynamic features. Within each of these two settings there is ample room for discussion, as there can

[13] Or 'show'; cf. [28] for an extensive analysis of how the universalism of Wittgenstein's early work, with its associated distinction between 'saying' and 'showing', is connected with the two conceptions of the role of formal systems in natural language analysis outlined here.

be better and worse tools, and we need to go into the empirical details in order to be able to decide which is which. It is between these settings that discussion loses its point. There is no answer to the question whether natural language meaning 'is' dynamic, or not: we simply lack a theory-independent concept of natural language meaning that we can refer to in our attempts to decide the issue one way or another. If we ignore this, conceptual muddles and misguided discussions will be the result.

Such considerations as the above apply in a wider context that the question of dynamic meaning. They are not confined to natural language semantics, but also apply in other contexts, such as the application of logic in the analysis of philosophical concepts, in cognitive science, and so on. A classical example is provided by the study of modal concepts. Here we can use modal logics to model the properties of modal concepts,, but many of their features can also be described using a non-modal tool, such as first order logic. In this case too, basically the same considerations as outlined above apply. This is not an empirical issue, or one that can be decided on conceptual grounds. Ultimately is a a matter of choice, If we fail to see that, we end up in fruitless debates, or, to use van Benthem's poignant phrase, 'system imprisonment'[14]:

> Nevertheless, I am worried by what I call the 'system imprisonment' of modern logic. It clutters up the philosophy of logic and mathematics, replacing real issues by system-generated ones, and it isolates us from the surrounding world. I do think that formal languages and formal systems are important, and at some extreme level, they are also useful, e.g., in using computers for theorem proving or natural language processing. But I think there is a whole further area that we need to understand, viz., the interaction between formal systems and natural practice.

I read van Benthem here as arguing for a pragmatic stance that is akin in spirit to the one outlined in this paper. Indeed, there is a lot that we still need to understand about the role that formal systems can, and cannot, play in an adequate account of our 'natural practices', i.e., in coming to an understanding of the ways in which we reason, use language, and so on. But if we can agree that a pragmatic attitude provides a better starting point than the essentialistic perspective that has informed too much of our discussions thus far, we have gained something.

References

1. van Benthem J (1986) Essays in logical semantics. Springer, New York
2. van Benthem J (1989) Semantic parallels in natural language and computation. In: Ebbinghaus H-D et al (eds) Logic colloquium '87, studies in logic and the philosophy of mathematics, vol 129. North-Holland, Amsterdam, pp 331–375
3. van Benthem J (1991) General dynamics. Theor Linguist 17:159–201
4. van Benthem J (1996) Exploring logical dynamics. CSLI, Stanford
5. van Benthem J (1999) Wider still and wider: resetting the bounds of logic. In: The nature of logic, European review of philosophy, vol 4. The University of Chicago Press, Chicago

[14] In van Benthem [5].

6. van Benthem J (2013) Implicit and explicit stances in logic
7. van Benthem J, Muskens R, Visser A (1997) Dynamics. In: van Benthem JFAK, ter Meulen, Alice GB (eds) Handbook of logic and linguistics. Elsevier, Amsterdam, pp 587–648
8. Breheny R (2003) On the dynamic turn in the study of meaning and interpretation. In: Peregrin J (ed) Meaning: the dynamic turn. Elsevier, Dordrecht, pp 69–89
9. Cresswell MJ (2002) Static semantics for dynamic discourse. Linguist Philos 25:545–571
10. van Eijck J, de Vries F-J (1992) Dynamic interpretation and Hoare deduction. J Logic Lang Inf 1(1):1–44
11. Gauker C (2007) Comments on dynamic semantics. APA Central Division, Chicago
12. Groenendijk J, Janssen T, Stokhof M (eds) (1984) Truth, interpretation and information, Grass-series, vol 2. Foris, Dordrecht
13. Groenendijk J, Stokhof M (1990a) Dynamic Montague grammar. In: Kálmán L, Pólos L (eds) Papers from the second symposium on logic and language. Akadémiai Kiadó, Budapest, pp 3–48
14. Groenendijk J, Stokhof M (1990b) Two theories of dynamic semantics. In: van Eijck J (ed) Logics in AI. Springer, Berlin, pp 55–64
15. Groenendijk J, Stokhof M (1991) Dynamic predicate logic. Linguist Philos 14(1):39–100
16. Groenendijk J, Stokhof M (2005) Why compositionality? In: Carlson G, Pelletier J (eds) Reference and quantification: the Partee effect. CSLI, Stanford, pp 83–106
17. Groenendijk J, Stokhof M, Veltman F (1996) Coreference and modality. In: Lappin S (ed) Handbook of contemporary semantic theory. Blackwell, Oxford, pp 179–213
18. Heim I (1982) The semantics of definite and indefinite noun phrases. Ph.D. thesis, University of Massachusetts, Amherst (Published in 1989 by Garland, New York)
19. Heim I (1983) File change semantics and the familiarity theory of definiteness. In: Bäuerle R, Schwarze C, von Stechow A (eds) Meaning, use, and interpretation of language. De Gruyter, Berlin
20. Kamp H (1981) A theory of truth and semantic representation. In: Groenendijk J, Janssen T, Stokhof M (eds) Formal methods in the study of language, MC tracts, vol 135. Mathematical Centre, Amsterdam (Reprinted in Groenendijk et al. 1984, p 1–41)
21. Lewis KS (2011) Understanding dynamic discourse. Ph.D. thesis, Rutgers, New Brunswick, New Jersey
22. Lewis KS (2012) Discourse dynamics, pragmatics, and indefinites. Philos Stud 158(2):313–342
23. Rothschild D, Yalcin S (2012) On the dynamics of conversation
24. Schlenker P (2007) Anti-dynamics: presupposition projection without dynamic semantics. J Logic Lang Inf 16:325–356
25. Stalnaker R (1974) Pragmatic presuppositions. In: Munitz M, Unger P (eds) Semantics and philosophy. New York University Press, New York
26. Stalnaker R (1979) Assertion. In: Cole P (ed) Syntax and semantics 9—pragmatics. Academic Press, New York
27. Stalnaker R (1998) On the representation of context. J Logic Lang Inf 7:3–19
28. Stokhof M (2008) The architecture of meaning: Wittgenstein's Tractatus and formal semantics. In: Levy D, Zamuner E (eds) Wittgenstein's enduring arguments. Routledge, London, pp 211–244
29. Stokhof M (2012) The role of artificial languages. In: Russell G, Fara DG (eds) The Routledge companion to the philosophy of language. Routledge, London/New York, pp 544–553
30. Stokhof M, van Lambalgen M (2011) Abstraction and idealisation: the construction of modern linguistics. Theor Linguist 37(1–2):1–26
31. Veltman F (1984) Data semantics. In: Groenendijk J, Janssen TMV, Stokhof M (eds) Truth, interpretation and information. Foris, Dordrecht, pp 43–62
32. Veltman F (1986) Data semantics and the pragmatics of indicative conditionals. In: Traugott E et al (eds) On conditionals. Cambridge University Press, Cambridge
33. Veltman F (1996) Defaults in update semantics. J Philos Logic 25:221–261
34. Vermeulen CFM (1994) Incremental semantics for propositional texts. Notre Dame J Formal Logic 35(2):243–271
35. Visser A (1998) Contexts in dynamic predicate logic. J Logic Lang Inf 7:21–52