
MEANING, INTERPRETATION AND SEMANTICS

MARTIN STOKHOF
ILLC/DEPARTMENT OF PHILOSOPHY
UNIVERSITEIT VAN AMSTERDAM

1 Introduction

This paper¹ explores, quite tentatively, possible consequences for the concept of semantics of two phenomena concerning meaning and interpretation, viz., radical interpretation and normativity of meaning. Both, it will be argued, challenge the way in which meaning is conceived of in semantics and thereby the status of the discipline itself.

For several reasons it seems opportune to explore these issues. If one reviews the developments in semantics over the past two decades, one observes that quite a bit has changed, and one may well wonder how to assess these changes. This relates directly to the status of semantics. If semantics is an empirical discipline, one might expect that most changes are informed by empirical considerations. However, one may also note that the core notion of semantics, meaning, today is conceived of quite differently than in, say, the seventies. How can that be? How can that what semantics is about, be different now from what it was back then? Or is this perhaps an indication that semantics is not as empirical as it is often thought to be?

Moreover, it seems that in some deep sense meaning as explicated in semantics and interpretation as studied in various philosophical approaches are strangely at odds. Meaning is what interpretation is concerned with: meaning is, at least so it seems, what in the process of interpretation language users try to recover (or analogously, what they try to convey in production). Yet, the way meaning is conceived of in semantics seems not to square all that neatly with how the process of interpretation is supposed to proceed. In particular it seems to lack some of the intrinsic features that various approaches to interpretation assume it to have. Given these discrepancies, one wonders how the two can be incorporated within a single theory. And that such a theory is desirable goes, it may be presumed, without saying

These are the reasons that figure in this paper. At the background there are some others. Semantics started out as interdisciplinary undertaking, drawing on philosophy, linguistics, and logic. Nowadays, however, there is little interaction between philosophers and semanticists. This is deplorable since some of the issues which philosophy deals with seem highly relevant for semantics, while in semantics theories have been developed which are not being taken into account by philosophers. Then there is also a broader issue: the split at the beginning of the twentieth century within philosophy itself between the analytical and the continental traditions, in particular

¹I would like to thank Johan van Benthem, Jan van Eijck, Thomas Hofweber, David Israel, Frank Veltman, Ed Zalta and two anonymous referees for their critical comments on an earlier version.

phenomenology and hermeneutics. Where the analytical tradition has by and large focussed on the development of theories of meaning, the continental approaches have concentrated on the problem of interpretation. For an overall theory of natural language semantics both are relevant, but there is an obvious problem if one wants to make the connection. And then again there is a consideration concerning the analytical tradition itself. This has been largely though not exclusively informed by formal logic, a trait inherited by semantics. Bearing in mind the distinction between logic as a tool and logic as a set of specific philosophical assumptions, the question is whether semantics uses logic merely as a tool or whether it indeed embodies some of these assumptions. The answer to that question obviously has repercussions for what we (can) take semantics to be.

These are broad issues, however, too broad to be addressed here. The aim of this paper is merely to raise a few questions concerning the status of semantics and identify some relevant considerations. In section 2 three views on the disciplinary nature of semantics are outlined and some philosophical assumptions of the enterprise are identified. Then follows a summary discussion of radical interpretation (section 3.1) and normativity of meaning (section 3.2), two phenomena which challenge these assumptions. In section 4 some recent developments in semantics are reviewed and evaluated in the light of the foregoing.

2 The status of semantics

The most common initial reaction to the question ‘What is the nature of semantics?’ is that, of course, semantics is an *empirical science*. Natural language meaning is an empirical phenomenon and constitutes its object of inquiry. However, this is a position that is more easily taken than defended. What exactly are the data and how hard and independent are they? How does one gather them? And to what extent do they provide one with an independent testbed for the descriptions and theories based on them? Here it is important to note that with very few exceptions semanticists use a particular methodology to approach their subject matter. The empirical domain is said to consist of intuitions, of individual speakers of the language, about meanings of expressions, or classes thereof, and about meaning relations, such as entailment and synonymy.² At the background of this view stands,

²The idea that the empirical status of natural language semantics simply follows from the fact that natural language meaning is an empirical phenomenon, though intuitive at first sight, becomes problematic if one tries the following analogy. Consider the development of possible worlds semantics for modal logic. Modal logic already existed, so in that

it seems, the Chomskyeian revolution in linguistics. Two main characteristics of this revolution can be found in semantics as well. First of all, there is the assumption that language exists as an (infinite) object, i.e., as an entity in its own right, with its own structural properties and content, which can be studied independently of its actual use and its historical development.³ Secondly, the idea of an individual having complete knowledge of language is concomitant to the idea of intuitions constituting the primary data of the linguist. These two features of the Chomskyeian approach are then wedded to the descriptive kind of semantics that dates back to (at least) Fregean days. The result is an account of meaning that is individualistic, platonistic, and ahistorical.⁴ Usually this is justified by an appeal to the kind of

sense it was an empirical phenomenon. Then Kripke (or Hintikka, or Kaplan, . . .) ‘discovered’ (?) its semantics. Does that make possible worlds semantics an empirical discipline? The answer is (obviously?) ‘No’, so something else has to be brought to bear on this issue. Presumably, one would point out that although modal logic existed, and in that sense was an empirical phenomenon, it came into existence only as a human construction and that is what distinguishes it from natural language. But this defense is weak for at least the following reasons. First of all, it seems to imply that there is a principled difference between those sciences which deal with natural phenomena and those which are concerned with human constructions, inventions, artifacts and the like, and that only the former are truly empirical. Secondly, it seems to assume that there is a clear-cut distinction between what is natural and what is a human construction (in this sense). That, however, is doubtful: are human societies, for example, natural or constructed? There seems to be no easy way to decide this issue, and presumably the same holds for human languages.

³This assumption is not unique to the Chomskyeian approach, of course, but also characterizes various forms of structuralism. In modern semantics, however, it makes its appearance via the Chomskyeian revolution. Taking his lead from developments in modern logic, Chomsky applied the notion of a mathematical characterization to natural language syntax, thereby firmly establishing the idea of natural language as an infinite object in linguistics. In Montague’s work the infinite nature of the semantics and the syntax of natural languages come together. (And are treated mathematically in an elegant fashion, which still inspires much of the work in semantics today.)

⁴The coming about of this methodology is actually quite complicated. The appeal to intuitions as the empirical data which semantics seeks to explain, has more than one source, historically speaking. One of the most important of these is the rise of the subject as the core notion in philosophy as part of the epistemological turn which characterizes modern philosophy. Actually, Chomsky regarded Descartes, who was, of course, one of the main initiators of this development, as a direct predecessor of his ideas on language, meaning, and knowledge. Another source, which is also quite relevant for modern semantics since it has its roots not just in Chomskyeian linguistics but also in philosophical logic, is Frege’s platonism with respect to meanings, which is not so much a straightforward ontological (metaphysical) doctrine, but rather an epistemological one, which is intended to guarantee intersubjectivity of meanings by postulating their objective status. So Cartesian rationalism and Fregean platonism are two different ways of vindicating a particular methodology in semantics, viz., the one which regards intuitions of individual language users about meanings and meaning relations as the primary data of semantics.

abstraction that is supposed to be characteristic for every empirical science. Physics studies movement on a frictionless plane, or the properties of a perfect vacuum; analogously, semantics deals with meaning and its properties on an equally abstract, yet empirical level. Or so it is assumed. But whereas physics, abstraction notwithstanding, produces empirically testable predictions, it is far from clear that the same holds for semantics. Do we predict intuitions? What would that mean? Or are our formal theories intended as descriptions of the actual processes that occur in natural language users? If so, how are we to test them? And given the fact that these theories deal mainly with structural aspects of meaning and not with lexical content,⁵ are they not hopelessly incomplete?

A less naive reaction might be dubbed the *engineering view*: here semantics is taken to be a much more modest enterprise. Instead of viewing it as a theory (in a strict sense) about what meaning is and how language users operate with it, one refrains from such far-reaching claims and merely requires that semantics provide a formally and empirically adequate specification of the input–output conditions. Whatever the actual underlying processes and meaning itself (which could be some kind of cognitive structure, or a system of neural networks), turn out to be, an account of them is to be in accordance with the results of a semantic theory, but the latter

Note finally that once people started questioning the central role of the subject different views on the nature of language arose along with it. A prominent example is provided by the later works of Wittgenstein.

⁵This distinction often figures in a defense of the empirical status of semantics. The idea is that a clear distinction can be made between structural semantic operations and lexical content, where the structural operations are invariants over the varying contents of judgments. Such invariants can then be studied at the level of a single language, or at that of all natural languages. Or one can look even more generally at all forms of symbolic reasoning devices, which leads to one definition of what logic is all about. However, vis à vis natural language semantics the following observations may shed some doubts on the viability of the distinction. First of all, it is framework-relative, i.e., it depends (partially) on the formal languages one has decided in advance to use. For example, if one uses a sortal first-order logic different aspects will be regarded as structural semantic operations than if one uses a type-theoretic framework. This does not mean that the choice is a matter of mere convenience, but it does show that what structural semantic operations are, is not a purely empirical issue.

Secondly, at a higher level, the distinction is interest-relative. (Cf., its origins: logic) It depends on an independent, theoretical decision which features of meaning one is interested in in the first place. In logic one is concerned with reasoning and hence one studies those aspects of meaning which are invariants with respect to that. But how could such a decision be made in advance for natural language? Again, this is not to say that the distinction is not useful or even necessary to make. The point is that it is not empirical, which means that it can not be appealed to in a defense of the empirical status of semantics.

is supposed to be conceptually neutral with respect to the former. This is not unlike one of the ways in which semantics plays a role in practical applications, such as in natural language query systems, translation systems, and the like, where semantics itself is not at stake in the actual systems, but functions on the meta-level as a kind of correctness criterion.⁶ Obviously, this is a more modest stance, and it saves semantics from some of the more nagging questions. But it does so at a cost, since this point of view seems to rob semantics of quite a lot of its empirical content. If natural language meaning is in some sense of the word a systematic phenomenon, it can most probably be captured by formal (logical, mathematical) means: that is not a very exciting claim. (This is not to imply, of course, that actually giving such an account is easy or straightforward.) Moreover, this line of defense betrays the actual practice of semanticists. In reality there is a lot of debate about the empirical aspects of semantics,⁷ so presumably the goals of the people debating these issues are less modest than this view of its status can account for.

The engineering view leads quite naturally to an even more radical position, which is that semantics is a *deductive science*. On this view the apparent empirical contents of semantics are more like the intuitions (about space, for example, or numbers) that put some constraints on a mathematical theory. This way of conceiving semantics has a certain attractiveness. It lends it a kind of elegance and puts it aloof of some of the difficult questions raised by the empirical point of view. However, it seems that it will not do in the end. The question of justification can be avoided for some time, and it may even be quite reasonable and justifiable to ignore it in a certain stage of the development of a discipline, but in the end it will simply not go away. Even mathematics struggles with this problem, albeit not in the same way and to the same extent. On the deductive view semantics is concerned

⁶Within such systems all one is concerned with is with matching expressions of one kind with expressions of another kind. As such this matching of expressions is not semantics in and of itself. (This is, it seems, the gist of Searle's famous Chinese Room Argument.) Semantics, of course, does play a role in specifying and checking the expressions involved and the matching between them. Only given a semantics are the expressions in fact representations, and is the matching a translation (or some other kind of semantic relation). However, do note that semantics plays a role here only at a meta-level, and is not really present inside the system.

⁷A case in point being the debate about compositionality and representationalism, which obviously rests on the (implicit) assumption that the inner apparatus of a semantic theory embodies empirical predictions. One question that immediately rises is: Which features of theories is it that are considered to have a 'real' counterpart? The concepts used? Just the mechanisms and rules proposed? Clearly, this is an issue that needs to be addressed in a systematic fashion, but it is one which is discussed only rarely.

with the discovery and study of abstract structures underlying an empirical domain. However, there is structure everywhere, and not all of it is equally significant. So even on this detached view there remains a question to be answered, or, equivalently, a demand to be met: How does one decide which structures are important, significant, or, as one, however reluctantly, is inclined to say, real? ‘Reality bites’: significance as a criterion of adequacy is necessarily transcendent to a theory, so the empirical question, it seems, can be forestalled, but not avoided.

But if that is true, as it seems to be, neither can one forego the task of examining the philosophical assumptions on which semantics rests. This involves not just identifying such assumptions, but also asking the question of their justification. Above, two such assumptions were identified, viz., that language is an (infinite) entity and that linguistic (semantic) competence is individual. These assumptions are under heavy attack, and this means that one also needs to investigate whether they can be modified. Can semantics be embedded in a different, more justified set of assumptions about the nature of language and competence? Or can such assumptions be dropped altogether, i.e., can semantics be done in a philosophical void?

Two issues which seem highly relevant in this respect will be discussed in some more detail in what follows: radical interpretation and the normativity of meaning.

3 Meaning and interpretation

As was indicated above, meaning is what semantics is supposed to deal with and a theory of interpretation is traditionally assumed to be drawing upon its results. Interpretation, according to this view, is the act of determining what the meaning of an utterance is.⁸ (Indexicals present a complication, which necessitates a distinction between meaning and content.) So meaning and interpretation are thought of as standing in a quite definite relationship: meaning somehow must be specified independently of, and prior to, interpretation.

One place where this assumption comes to the fore quite clearly is in the traditional division of labor between semantics and pragmatics. According to the Gricean scheme, which has been adopted for a long time by many people working in the field, semantics deals with meaning, basically

⁸Analogously, in the production of an utterance, the sentence uttered is assumed to have a meaning attached to it which is independent of the particular act of utterance. In what follows only interpretation is discussed, but most of the discussion applies to production as well.

conceived of as truthconditional content, and pragmatics deals with (certain aspects of) interpretation.⁹ The case of conversational implicatures provides a nice illustration. The utterance is supposed to have a specific meaning, which somehow does not fit into the present conversation. Then pragmatics takes over and decodes the intended message, in terms of conversational implicatures.

But traces of this idea of a hierarchical relation between meaning and interpretation can be found elsewhere, too. For example, in speech act theory the distinction between direct and indirect speech act embodies a quite similar idea: the direct speech act performed by a certain utterance constitutes its normal (conventional) meaning. And interpretation, in this case the act of figuring out the indirect speech act that is performed, comes into play only after that: when there is some anomaly between the meaning and the situation in which it is produced. And most approaches to the meaning of literary figures, such as metaphors, and literary texts in general, likewise assume that (literary) interpretation proceeds on the basis of an independent and pre-given literal meaning.

These are all cases of situations in which interpretation comes into play when there is a meaning, but this meaning somehow does not fit. But a similar idea underlies the conception of interpretation as it operates in situations when no initial meaning can be assigned, for example when an expression is used that we do not know the meaning of. Here, too, it is assumed that such a meaning exists, and that what interpretation does is recover it (from the context, or by other means).

The picture that emerges sets meaning and interpretation clearly against each other. Meaning is an entity, some kind of abstract or mental object, which is somehow contained in an utterance, which is produced in a mechanistic fashion, and which requires for its production no conscious processing.¹⁰ Interpretation, on the other hand, is a process which operates on utterances and in which one intends to recover a meaning in a more or less conscious fashion. Note that according to this picture, reasoning plays a crucial role in interpretation, and no role, or hardly any, in semantics.

It would be wrong to suggest that there is no process of interpretation that has these characteristics, or that no distinction can or should be made between meaning and interpretation. However, it does seem that the tradi-

⁹Of course, some people do not agree with Grice that this is where the boundary is to be drawn. But, nevertheless, they share the idea of an underlying distinction and a hierarchical relation.

¹⁰Other features such as the idea of atomism and the concomitant adoption of compositionality as the main underlying principle are intimately related to this view.

tional picture is distorted and misleading, in particular where it postulates a hierarchical relationship between the two. In order to substantiate this point two phenomena will now be discussed which challenge this picture.

3.1 Radical interpretation

Radical interpretation poses a simple but real problem: How does one understand the utterances of a speaker in a tongue that is completely foreign? One knows nothing about the language, the convictions and beliefs, customs and other ways of behavior. At first one does not even know how to segment the sounds that are uttered, how to determine the boundaries between the units of speech. Is this an exclamation, a word, one word or two? A sentence, or a series of such? Is this a statement being made, a question being asked? Does this express bewilderment or firm conviction? A lot needs to be established and anyone who has actually been confronted with speech in a completely foreign tongue knows how insurmountable the difficulties seem at first, how impossible the task at the outset appears to be.

But radical interpretation does not stop here. Sure enough, when conversing in one's own tongue with one's fellow speakers one takes it for granted that all mean the same things with the same words and sentences, that thoughts and conjectures, desires and queries, are communicated by the use of a common tool: a shared language with shared meanings. But on what assumptions is this conviction based, and how justified are they? Is there any reason to think that meanings are shared in this way? And what exactly is it that is shared? What kind of object is a meaning that one can share it, within a language and across languages? Or is it not objects that are shared, and is meaning to be located on an entirely different ontological plane? And how exactly is meaning tied to linguistic entities, to words and phrases? Are they merely conventional signs for something that is in essence not of a linguistic nature? Does an interior intentional act bestow meaning upon the words and phrases used? And how is it that some inner mental process or action is able to link two such utterly different entities as a meaning and an expression? For expressions are physical entities, written signs or spoken sounds, and meanings are not like that at all. Or is the meaning somehow located in the mind and fused into the sign as it is produced, the way a letter is slipped into an envelope?

It is questions such as these that have been the central concern of much of twentieth century philosophy of language, and the various answers that have been proposed to them, have shaped, in one way or the other, the semantic theories that make up the confusing multiplicity that the field

exhibits, even today. The fact that the problem does not seem to have been solved (and are philosophical problems ever?) should not blind one to the progress that has been made. Of course various positions are still held today, and positions which are quite incompatible at that. But the ways in which they are defended have changed quite a lot, and in that sense it seems that a better understanding of the issues involved has been reached.

The most celebrated form in which the problem has been presented, is probably that of Quine's thought experiment of radical translation.¹¹ Here one is invited to imagine oneself trying to come up with a translation of a language of an unknown people whose culture is completely unfamiliar and whose tongue bears no affinity to one's own. Starting from premises as parsimonious as possible Quine allows information about observable happenings in the immediate environment, the verbal behavior of this people displayed in reaction to these events and a little introspection, into one's own verbal reactions. Only a particular kind of utterances, those immediately tied to sensory input and more or less commonly agreed upon, are up for translations that can be verified. The rest, Quine says, is theory: one constructs the remainder of the translation manual on the basis of so-called analytical hypotheses, of which there will be many, mutually incompatible, yet all in accordance with whatever data one can gather.

Although Quine originally formulated the idea of radical translation as a philosophical thought experiment, one does well to note that it has a real, empirical counterpart. For example, if one considers the case of first language acquisition without any pre-conceived ideas about the underlying mechanisms, it seems to fit Quine's bill quite accurately. All the child has to go on is its perceptual input from the environment, the verbal behavior of the adults surrounding it, and, of course, its own mental make-up. Now, it is important to note that it would be begging the question if it were assumed that the latter already contains so much content and structure that the problem is solved: that is exactly the issue that is at stake. Of course if innate mechanisms, introspective access to language independent meanings and the like are postulated, the problem becomes solvable — in principle. But precisely such postulates are not philosophically harmless, to say the least, and create problems of their own. Likewise, if one were to say that, of course, semantics can not account for the situation of radical translation, or language acquisition, but has to be looked upon as a theory that applies to the end stage, so to speak, one does indeed set aside these worries. But given that radical translation, language acquisition (and, as shall be argued

¹¹The classic reference is [12, chapter 2].

shortly, interpretation in general) has priority, at least chronologically, how should one evaluate a theory about the result that has no way of accounting of how it came about in the first place?

But there is another and perhaps brighter way to look at this predicament. For why exactly is it that the task seems so formidable and the result so hopelessly inadequate? It seems that at the bottom of this reaction lie the same premises which were identified above as characteristic for semantics, viz., that language is an entity that is somehow surveyable in its entirety and that competence is an individual matter.¹² Given these assumptions the results that Quine holds are all one can get, certainly fail to make for an adequate theory of translation and ipso facto of meaning for the language translated. But one might turn the table and take the initiative by asking oneself the following question: If that is all one can get, perhaps that is because that is all there is to it? Indeterminacy of translation, in Quinean terminology, may simply be due to indeterminacy of meaning, of their language, and of one's own.

This point can be reinforced by considering the case, not of radical translation of one vernacular into another, but of radical interpretation within the same vernacular.¹³ If one can bring ourselves to refrain from the assumption that every individual language user enters the situation of interpretation equipped with a complete theory of the structure and content of the language used, it seems quite natural to conclude that whatever meaning one attaches to utterances of others and thus whatever theory of meaning one comes to possess, is the result of (acts of) interpretation, and not a tool that one brings along right from the start. This is, it seems, what Davidson's idea of prior and passing theories (in [4]) is about. Language users enter a situation of interpretation with some partial theory of meaning, but rather than the norm to which subsequent utterances have to conform and which determines their interpretation, this is a theory that is changed by the very act of interpretation. Of course, some more or less stable generalizations may emerge over time, but in principle these are also up for modification and they do not constitute anything like a theory of meaning in the traditional sense. There are some obvious parallels between this view and the hermeneutical approach to interpretation, but for reasons of space that connection can not be investigated here.

¹²Another pertinent assumption seems to be that of a particular kind of fit between language and world. This assumption is equally dubitable but not equally relevant for the present argument. It is, however, clearly one that is at work in the view of semantics outlined at the beginning of section 2.

¹³As Davidson urged us to do, e.g., in [3].

3.2 Consciousness and normativity

A different, but arguably related issue which is relevant for an appraisal of the status of semantics concerns the inherent normativity of meaning and the role of consciousness.

That it is hard to establish semantics as a well-defined empirical discipline (with a clearly distinguished subject matter, goals and tools) might be partially explained by pointing out that meaning is a peculiar object indeed. Meaning is many-faceted: it derives from various sources, plays different roles, interacts in various ways with other cognitive and non-cognitive functions. This is clear from various phenomena, such as radical translation and direct reference, and it is what recent developments in semantics are trying to come to grips with. Another, and much more fundamental, reason is that meaning displays both ‘mechanic’ and ‘conscious’ features. Most of the time it is produced and processed completely mechanically, without any conscious effort or awareness. Yet in other cases it is used and scrutinized deliberately and consciously. In view of that, some have constructed an opposition between linguistic semantics and hermeneutic interpretation¹⁴ the former dealing with meaning at the mechanistic level and the latter being concerned with what happens in conscious interpretation. This can be regarded as a way to implement the meaning–interpretation distinction. However, it can be argued that this distinction lacks a basis in reality: there are no corresponding distinct realms of investigation, nor do linguistics and hermeneutics constitute rival approaches of the same realm. Meaning is something which is a natural phenomenon in this sense that it is located in human neurophysiology, something which often runs its course without any conscious effort or even awareness. (One could substitute ‘cognitive’ for ‘neurophysiological’ here, that does not alter the point.) But meaning is also something that is used consciously, which is manipulated, created, worried and wondered about. And it seems both individual and social. The individual aspect comes to the fore, for example, in misunderstandings: ‘my meaning’ does not get across. And the social is prominent for example in successful negotiations about meaning, or quite distinctively, in learning: understanding is established cooperatively.

The tension between compositionality and contextuality that has shaped a lot of the philosophical discussion about meaning in the twentieth century is related to this. Compositionality goes naturally with the mechanistic picture whereas contextuality seems to capture the essence of the role meaning plays in awareness and conscious use of language (and not only in, say, lit-

¹⁴Cf., e.g., [1]; cf., also [6], [7, chapter III.1].

erature, but also in everyday situations.) That this tension arose (at least in philosophy) in the twentieth century is no coincidence. With the disappearance of a ready made world as that to which language applies, the mechanistic picture, and compositionality along with it, lost one of its main underlying motivations. (Holding on to the purely individualistic point of view still saved this picture, at least until recently.) The world to which language applies is only partly ready made: it is also a social construction (which does not mean that it is not objective in some meaningful sense of the word). This social construction informs the meaning of language as much as does the physical, ready made aspect of the world. And this means that an exclusively mechanistic account of language as an unconsciously used tool must give way to one which acknowledges also the role of conscious use and the symmetric relation between language and its users involved: people create language and are being created by it.

This dual, janus-faced nature of meaning is apt to generate a curious, but significant problem if, that is, one sticks to the traditional view on how meaning should be accounted for. For this view runs straight into a problem which is in fact characteristic for many of the theories that have dominated linguistics, semantics and cognitive science since the sixties. It is what Jackendoff ([11]) has called ‘the mind–mind problem’. Where Descartes created the problem of how to account for the relation between the material world, including the body, and the disembodied mind, the cognitive revolution has doubled this worry. Now there is also the relation between the computational mind and the conscious mind to account for, a gulf to be bridged between the automated, non-introspectable mechanisms that are supposed to constitute human cognitive abilities and the phenomenological awareness of the world and of the mind.

It should be noted that as such this observation does not constitute an argument against the possibility of giving a thoroughly reductionistic account of the phenomena involved. One might hold that the mere fact that one can be aware of some cognitive activity or ability in itself does not show that the actual phenomenon can not (or should not) be explained in reductionistic terms, leaving the awareness as such as a kind of epiphenomenon that can be explained in other terms. In the case of meaning this would amount to the idea that the empirical object of semantics actually consists in the underlying neurophysiology of linguistic information processing (perhaps under some suitable abstractions). The phenomenon of awareness of meaning on such a view should be regarded as an epiphenomenon of the unconscious, low-level processes, that perhaps may perform some function,

but, if any, one which does not interfere with meaning as such.¹⁵ But, at least as far as language is concerned, it can be argued that such an approach would be thoroughly mistaken.

The crucial observation, which strongly suggests that the reductionist stance will fall short of supplying us with a satisfactory account, is simply this: language is not only embodied, it is also essentially normative.¹⁶ The rules, of whatever nature they turn out to be, which govern the use of language do not simply describe actual linguistic practices, they also constrain them in a normative fashion. Any theory of meaning that does not account for this (and semantics in its present state certainly would seem to have a hard time doing so) is incomplete at the very least. Normativity shows that language and meaning are the result of a subtle and very complex interplay between the physical, psychological and social nature of humans. Moreover, this normativity seems to be intimately tied to awareness of meaning and to conscious interpretation and production. If that is true, then a reductionistic approach must fail.¹⁷ Some elaboration on these points is in order.

First of all, one should note that the fact that certain processes (even some cognitive ones) are not introspectable (vision might be an example, respiration certainly is) is unproblematic. With respect to meaning, however, things are different. Here is a realm of processes, ‘objects’ and the like which do have conscious counterparts. Users reflect on meaning, they worry and argue about it, they try to figure out what someone’s utterances mean, they anticipate on how others will interpret their own utterances, and so on. How is all this to be accounted for if meaning were completely mechanistic and hence completely sealed off from the conscious mind?

There are certainly reasons to think that the kind of awareness users have of language and meaning is different from awareness that occurs, or can be made to occur, with other kinds of neurophysiological processes, such as movement, breathing or vision. The latter can be attended to, i.e., watched and influenced to some extent. With meaning much more is at stake. If one attends, e.g., to visual processes, paying attention to the way in which one is visually aware of a certain object, to the difference that various lighting

¹⁵Note that such a move would be quite in line with some of the fundamental tenets of the Chomskyeian view.

¹⁶This is, of course, one of the main points of the so-called ‘rule following considerations’ in Wittgenstein’s *Philosophical Investigations*.

¹⁷This is not to deny that a lot of actual linguistic behavior is in fact by and large the outcome of a mechanical and unconscious process. The point is that a truly reductionistic approach claims something much more far-reaching, viz., that this is all there is to it. And that seems one step to far.

conditions make or the position of one's head, etc., what one observes is not ipso facto subject to normative judgments. Take the familiar example of a two-dimensional drawing of a cube. The object can be seen 'bottom up' or 'top down' and one can even force oneself to switch between the two views, but there is not a correct way of seeing it. With meaning, however, this is different. Suppose one is not certain how to interpret a certain utterance, since a particular word occurring in it can have either one of two meanings. Here, one does not merely attend to what is going on, one is puzzled or worried. There are not two possibilities on equal footing, which one can switch between at will. One has to choose one and the choice makes a difference: only one of them is right.

An important observation which is closely connected with the previous one and which underscores the peculiar status of the object of semantics, is that interpretation (and, analogously, production) stands in need of justification. That is to say, both the process and the results of interpretation have to be accounted for in terms of reasons, not causes. Of course, when prompted for such justifications one may find oneself saying in the end: 'Well, that is what it means'. That, however, is not an ultimate justification, a reference to some a priori, foundational base.¹⁸ Often it is rather a pointer to an accepted account (or one that is assumed to be such) and if pressed, one is able to extend the account, at least in principle. (Not necessarily by oneself: sometimes one has to refer to others.)

To avoid confusion it is important to distinguish here between two different questions that may be raised. The first is the question 'What is the meaning of this expression?' ('Why do you interpret it in this way?', 'Why do you use it this way?'), and the second is the question 'Why does this expression mean that?'. The first one is a request for justification, the second asks for a historical, causal explanation. And an answer to the second question, even when factually correct, never constitutes an answer to the first since it lacks the required normativity.

But do not 'justifications have to come to an end' (as Wittgenstein reminds us)? Yes, and indeed sometimes the exclamation 'That is what it means!' may signify that, too. But note that the coming to an end of justifications is not the reaching of some ultimate justification, nor of some ultimate, true causal explanation. It points to the fact that justifications can be given only within a certain framework. It even may point beyond the framework, to some natural phenomena which constrain the framework

¹⁸The assumption of such a base is a distinctive feature of the atomistic view on meaning and one that is intimately linked to the special status the subject is assumed to have.

from the outside. However, it is important to realize that such a pointer, although it refers to something that is actually causally active, is not the same as a causal explanation. For that too can be given only within a framework.

So the argument is something like this. Consciousness is a source, or starting point, of normativity. Very roughly, normative practices arise from congruent causal patterns in behavior through consciousness of and reflection on those patterns. A shared nature (physiological and cognitive) is responsible for sufficient similarity in the causal mechanisms, which are probably calibrated further by means of common ways of adapting to a shared physical environment. Such congruences define a group of agents. Such group turns into a community if these congruences are noted and their effects are reflected upon. The behavioral patterns humans actually display are neither completely arbitrary nor are they completely determined by external constraints imposed by the environment, their physiology and so on. The various patterns are contingent actualizations within a wider range of possibilities. It is awareness of this fact and the (cognitive and non-cognitive) exploration of other, non-actual possibilities that lead to normative practices. Once agents are aware of the fact that they behave in a certain way but could also behave differently, the need for normative determination of behavior becomes imperative. This is where regularities turn into rules, behavioral patterns change into institutional practices.

This is, of course, only a rough sketch. It indicates, however, that normativity arises from an interaction between consciousness and causal mechanisms which takes place in a temporally extended, social setting. Also, it implies that normativity holds a central position in language, both with regard to meaning as well as to interpretation. And this has far-reaching consequences for the theory of meaning. The familiar architecture of a semantic theory which deals with the unconscious mechanistic realm of meaning and a pragmatic theory which is concerned with interpretation can no longer be maintained. For semantic phenomena exist which are intrinsically related with normativity and conscious use, such as mistakes and inconsistencies, lies and insincerity, which an adequate theory of meaning should be able to account for. A concrete example is provided by the Sorites-paradox:¹⁹ Under certain circumstances one can be led astray in the application of certain predicates (e.g., perceptual predicates such as color terms), be forced to apply them in what from a logical point of view is an inconsistent manner. It seems that this is a feature that is inherent to what meaning is and to how it is used, and therefore one that should not be explained away, but

¹⁹This was suggested by Frank Veltman.

accounted for.²⁰ Other relevant phenomena are connected with anticipation of referential interpretation, and so on.

That awareness is the source of normativity also explains why normativity is distributed in a language community, i.e., is social. First of all, the required awareness always also involves the linguistic behavior and the judgments of others. Secondly, the required awareness can not be located in one single individual. It has to be distributed over time, and over a community.

The consequence for the nature of meaning is straightforward: if a notion of normativity informs interpretation and interpretation informs meaning, in the sense that meaning is the outcome of interpretation rather than its starting point, then meaning itself has to be normative and distributed, i.e., social.

Time to take stock. In the above it was argued that the usual conception of what semantics does and what it is concerned with rests on assumptions which usually go unnoticed but which in fact are philosophically quite strong. Language and meaning are assumed to be (abstract or mental) entities which are infinite in nature and hence depend on an atomistic conception and require some form of compositionality. Meanings, moreover, are conceived of as individually graspable objects and interpretation and production are assumed to be processes which are individual-based and which use meanings as prefabricated elements in a toolkit.

Of course there is nothing incoherent in these assumptions as such. In fact they have a venerable philosophical ancestry and one may note that in particular cases, notably when in the case of formal languages, they actually lead to elegant and powerful theories. However, there seems to be ample reason to hold that as far as natural language is concerned these assumptions are unwarranted. They give rise to a distorted conception of meaning and, as a consequence of that, to theories which are conceptually not very satisfactory and which in the end will turn out to be of limited empirical use. Inasmuch as these assumptions go counter to some basic features of natural language meaning, two of which were discussed in the foregoing, the conceptual problem is fairly well-established. That does not mean, of course, that there is necessarily also an empirical problem. After all, conceptually deviant theories may survive from an empirical point of view for a long time. However, if one really takes on the task of modeling real, situated language, the empirical limitations will become clear rather quickly.

²⁰I.e., any form of logical 'stratification' is misguided here. A similar point can be made concerning semantic paradoxes.

4 Dynamic semantics

The analysis provided thus far assumes that semantics does indeed have the features identified earlier on, but one might well ask whether more recent developments may have perhaps distanced themselves from them, and, by doing so, have resulted in a more adequate explication of what meaning is. It appears that this is true only to a limited extent and that some crucial steps still have to be taken in order to arrive at a conception of semantics which will be able to deal with such issues (interpretation, consciousness and normativity) as were discussed above. Nevertheless, if one considers the way in which semantics has developed since its inception at the beginning of the 1970s, some definite trends can be discerned which seem to push it in the right direction.

Orthodox, Montague-style semantics is a perfect illustration of the traditional picture of the relation between meaning and interpretation sketched at the beginning of this paper. More precisely, it offers a theory of one element of that picture. It is concerned just with meanings and their structure and contains nothing about who knows them nor about the way in which they are known. This means that it does not say anything about interpretation as such. But inasmuch as it lacks a notion of interpretation, it also says nothing about what meanings are, in reality, i.e., for the users of the language.²¹ It clearly embodies the assumptions that were identified above: language and the realm of meaning are infinite and object-like, and competence, one may presume, is individual. It separates meaning from interpretation, and, concomitantly, semantics from pragmatics, and it clearly assumes that the former is independent from the latter.

Although the framework of Montague-grammar in the strict sense, i.e., as theory encompassing both semantics and syntax and embodying a particular view on their relationship, lost much of its appeal by the end of the 1970s, the style of semantics it propagates is still quite influential. Its scope was extended, not only descriptively, but also by enriching the logical apparatus in certain ways. However, by and large these extensions remain faithful to the original set-up and hence subscribe to the same assumptions. Real challenges based on a different conception of meaning started to emerge

²¹This becomes quite obvious when one considers what this style of semantics has to say about a particular kind of construction in which the ‘inside’ and the ‘outside’ of meaning meet, viz., propositional attitude ascriptions. When giving a semantics for constructions such as ‘s believes that ϕ ’ one in fact analyzes the point where meanings, conceived of as language internal objects, and language users, as intentional agents, meet. Small wonder that this style of semantics typically fails here.

only in the early 1980s, with the advent of situation semantics and discourse representation theory. Space does not permit a discussion of the ins and outs of these approaches, although that would be quite important in order to get a good picture of how semantics has developed. What follows is concerned with one approach which is of slightly later date and which seems to be a definite step in the right direction, viz., dynamic semantics.

Dynamic semantics changes the very notion of meaning and in doing so constitutes a real break with the Montague-tradition.²² In dynamic semantics meaning is considered as ‘context change potential’: the meaning of an expression is analyzed in terms of the way in which it changes the context in which that expression is used. Context is, of course, a ‘catch all’ phrase and one usually one concentrates on a particular aspect of it. In most formulations it is the information that speech participants have in a situation of linguistic information exchange. So contexts are identified with information states and meanings become operations on such states.

From the perspective of the issues discussed earlier this transition from truth conditional content to context change potential as the core notion of semantics evidently constitutes a step in the right direction. It emphasizes that language and meaning are situated and that interaction between language users and not an independent word-to-world relation is what meaning is primarily concerned with. By thus concentrating on the procedural features of meaning it takes one step away from the object-like notion of meaning that dominated the tradition. Also, by connecting meaning and information an epistemic perspective is brought to bear on meaning which invites a non-reductionistic, non-mechanistic view on how language users operate.

Still, it seems that in its present state dynamic semantics does not constitute as radical a departure from the old paradigm as is needed. The point here is not that dynamic semantics focusses on just one factor of context. There seem to be no principled reasons why the framework could not be enriched so as to take into account more than just the information states of speech participants and cover also, say, desires, goals, plans, and so on. (Of course this is not to say that such extensions are straightforward.) Rather, the claim is that dynamic semantics is still too much bound by the individualistic assumption of the Chomskyeian tradition. This should be apparent from the fact that in dynamic semantics meanings themselves are defined as entities which operate on states. And this implies that in some sense the states are given in advance, independently from the meanings.

²²Cf., [10] for one particular formulation.

An utterance, it is assumed, constitutes a transition from attitudinal states (or rather, a pair of them, one for the speaker and one for the hearer) to attitudinal states (or pairs of such). Notice first of all that this transition is not deterministic, it is not a function. Rather, it constitutes a constraint on possible outputs.²³ It follows that if that is the case, utterances not only do not deliver (pairs of) states, they also do not actually operate on them. Rather, both output and input are sets of constraints on states.²⁴

But this is perhaps just a side issue. The important observation to make is the following. This description of what an utterance does, i.e., constraining states in a non-deterministic way, assumes that states exist and can be specified independently from the discourse, i.e., from meaning and its use. (Notice that this holds irrespective of whether the constraints are formulated pointwise or globally.) The process can only be described in this way if it is assumed that states come into existence and thus can be identified and analyzed independently from any linguistic communication. By implication, it seems that this assumes that it is individual language users who are in such states and hence that competence is individual after all. Here, again, the Chomskyeian assumption of individualism and introspection can be discerned.²⁵ The reconstruction of meanings as updates on information states assumes that it is individuals that are in such states and that these individuals have privileged access to the states they are in.²⁶ Inasmuch as these Chomskyeian assumptions are questionable, it follows that although dynamic semantics is proceeding in the right direction, it still needs to take a basic step.

It more or less follows from the observations concerning radical interpretation, consciousness and normativity that attitudinal states themselves

²³Good examples of why the transitions are not deterministic are provided by all kinds of identifiers. Cf., e.g., [5].

²⁴Possible exception: the first utterance in a discourse. But even that is questionable.

²⁵This assumption is also present in one of the few attempts to give a formal reconstruction of social meaning, that of Peter Gärdenfors. (See [8].) An example that shows this is the following. A principle that plays a central role in Gärdenfors' theory is that of contingency preservation: 'if something is contingent for all individuals, the social meaning is contingent as well' (p. 295). Note that this takes the bite out of some of the claims of direct reference theorists: for example, for rigid designators a and b it holds that $a = b$ is necessary, even if for some (all) individuals it is contingent (e.g., because the reference of one of the terms is unknown). Note also that Gärdenfors needs to prevent precisely this result for his proof of the existence of a semantic arbiter (theorem 2, p. 297).

²⁶Notice also that individualism and mentalism (and neurophysiological reductionism) are intimately related (though not the same). Unfortunately space does not permit a discussion of the attempts of 'broad vs. narrow content' theories of meaning to save individualistic semantics in the face of these and similar challenges.

are as much the product of linguistic capabilities (of interpretation and production) as they are the input on which such capabilities operate. Meaning is the outcome of interpretation and interpretation itself is something that interacts with and in a sense derives from attitudinal states. A good example of the interrelatedness of linguistic capabilities and attitudinal states is provided by Wittgenstein's question about the dog.²⁷ Why can one say that a dog believes that its master is at the door, but can it not be said that it believes that he will return in two weeks from his business trip to Brussels? Obviously, because the dog can not display the appropriate kind of behavior, whereas humans can. But this behavior has certain distinctive linguistic components: expecting a friend to return in two weeks is something humans are able to do precisely because they have language. (The relevant behavior is linguistic in a wide sense: answering 'In two weeks' to the question 'When does he return?'; marking the date in a calendar; counting the days; and so on.) So the attitudes humans can ascribe to each other, and, more fundamentally, which they can actually have, depend to a large extent on their linguistic abilities

But this means that attitudes and attitudinal states and meaning can not be separated in the way in which dynamic semantics assumes they can. At best then dynamic semantics models an ideal, completely stabilized endgame. But that is not a very satisfactory view. (Cf., above where regarding radical interpretation a similar point was made.) And in the light of the fact that meanings are not only acquired, but also do change, and are often adjusted to fit a particular situation of linguistic information exchange,²⁸ one may seriously doubt whether dynamic semantics describes any real situations of linguistic information exchange at all.

Of course some of these issues are acknowledged, in some way or other, and attempts are made to come to grips with them. Promising developments in this respect are, e.g., analyses of information and information exchange which allow information to arise at the level of a group of language users without being derived from information of its individual members.²⁹ Another example is provided by attempts to incorporate elements of game theory into dynamic semantics. Game theory provides insights into particular kinds of cooperative behavior from which a semantics which takes interaction

²⁷*Philosophical Investigations*, Part 2, section i.

²⁸A point that is convincingly made, e.g., by Davidson ([4]). It is also one that is acknowledged in dynamic semantics. For an attempt to deal with some of the consequences seen the already mentioned [5].

²⁹Cf, e.g., [9]. It would be interesting to see how this compares with work on collective intentionality.

as one of its starting points may benefit. Work on evolutionary games seems especially relevant since it may lead to an analysis of how the normativity of meaning arises from linguistic and non-linguistic interaction.³⁰

5 Conclusion

The main argument of this paper is that if one takes conceptual considerations concerning meaning, such as radical interpretation and normativity, seriously, the traditional conception of what semantics is about, will have to change.

The idea of language and its meaning as an entity which consists of a potentially infinite number of sentences each of which has a definite pre-established meaning, will have to be abandoned. Consequently, the principle of compositionality loses its pivotal status, as does the idea of a strict dichotomy between structural semantic operations and lexical content. After all, the principle of compositionality and the concomitant distinction were the primary (though perhaps not the only) conceptual tools for solving the problem of giving a complete description of an a priori given, infinite language. But if language and meaning are not entities like that at all, much of the intuitive appeal of these notions is lost.

Likewise, the assumption of the conceptual priority of a priori, non-contextual meanings and the related hierarchical division of labor between semantics and pragmatics seems no longer justified. Interpretation is essentially a situated process and if meaning is the result of interpretation this situatedness will be all-pervading. Dynamic semantics arguably goes some way to reverse the order of conceptual priority here, but a more fundamental switch seems to be needed. A simple example. It can be shown that in dynamic modal semantics some questions of identification can not be solved unless the language contains demonstratives.³¹ The traditional reaction would be to postulate a difference between linguistic and non-linguistic sources of information and to confine semantics to the former. But if one take the situatedness of meaning to heart one takes this dependence on the environment right into the very notion of meaning.

Only if that is achieved is there any hope to get an account of the seamless way in which humans integrate information from different sources. For

³⁰This is only concerned with developments in formal semantics. Various issues, among which normativity, are center stage in recent informal philosophy of language. Exemplary here is the work of Brandom ([2]). However this type of work usually does not provide any direct clues as to how to develop formal semantics in the right direction.

³¹Cf., [10].

language is essentially a human phenomenon, and as such an integrated part of human information processing. The way humans obtain, process and transmit information is highly marked by their biological, psychological and sociological make-up. In particular, the embodied nature of the various means of dealing with information seems to be severely underrated in current approaches to semantics. Taking this feature seriously would, lead to semantic theories that will be much easier to integrate with other modules of human information processing, such as vision.

As was indicated above, one consequence of taking the idea of radical interpretation seriously is that meaning is the result and not the toolkit of interpretation. This reversal of the traditional order at least provides a handle on tackling the ‘mind–mind problem’. First of all, interpretation is a conscious process, at least in principle. That much of the everyday use of language takes place at the back of the mind is not an issue. Various abilities are employed consciously only if there is a reason to do so: if a misunderstanding arises, if there is a need to mislead. Secondly, the need to postulate a complete description of an infinite language is dropped and along with it the need to appeal to a mysterious mechanistic realm of the mind that completely accounts for our linguistic abilities.

Finally, as semantics develops along these lines the question as to its status as a scientific discipline will probably become more clear. In the end, of course, semantics is, or at any rate should be, about something that is real. And in that sense it is an empirical discipline. But the reality it tries to describe and understand is very complex. Its fact of the matter cannot be relegated to one realm, be this that of psychology, biology, sociology or some other discipline. But neither does it form a platonic realm of its own. Language, meaning, interpretation are complex phenomena precisely because they exhibit a variety of aspects and also because they do double duty. They are subject to conscious use and at the same time form the consciousness that uses them. A proper methodology for a comprehensive study of phenomena such as these is still in the making. (That and the complexity itself justify abstraction, of course, and the resulting approximations can be put to practical use. In those cases the engineering view may apply as well.) But the empirical and methodological complexity is balanced by significance. After all, few phenomena are as central to human existence as these.

References

- [1] J. Bouveresse. *Hermeneutique et Linguistique*. L'Éclat, Combas, 1991.
- [2] R. Brandom. *Making It Explicit*. Harvard University Press, Cambridge, Mass., 1994.
- [3] D. Davidson. Radical interpretation. *Dialectica*, 27:313–328, 1973.
- [4] D. Davidson. A nice derangement of epitaphs. In Ernest LePore, editor, *Truth and Interpretation*, pages 433–446. Blackwell, Oxford, 1986.
- [5] P. Dekker. Speaker's reference, descriptions and information structure. *Journal of Semantics*, 15:301–330, 1998.
- [6] H.-G. Gadamer. Semantics and hermeneutics. In *Philosophical Hermeneutics*, pages 82–94. University of California Press, Berkeley, 1976.
- [7] H.-G. Gadamer. *Truth and Method*. Crossroad, New York, 1989.
- [8] P. Gärdenfors. The emergence of meaning. *Linguistics and Philosophy*, 16:285–310, 1993.
- [9] J. Gerbrandy and W. Groeneveld. Reasoning about information change. *Journal of Logic, Language and Information*, 6:147–169, 1997.
- [10] J. Groenendijk, M. Stokhof, and F. Veltman. Coreference and modality. In S. Lappin, editor, *Handbook of Contemporary Semantic Theory*, pages 179–213. Blackwell, Oxford, 1996.
- [11] R. Jackendoff. *Consciousness and the Computational Mind*. MIT Press, Cambridge, Mass., 1987.
- [12] W. V. O. Quine. *Word and Object*. MIT Press, Cambridge, Mass., 1960.