

# Meaning in Micro-Communities

Henk Zeevat  
University of Amsterdam\*

## 1 Meaning

One way of describing the enterprise of natural language semantics is by analogy with interpreted formal languages, e.g. the language of arithmetic interpreted on the natural numbers. English is then the formal language consisting of the well-formed sentences of English interpreted on the structure we happen to find around us. The business of the natural language semanticist is to describe the correlation of the formal expressions with that structure. This paper is directed against this particular interpretation<sup>1</sup> of the enterprise of logical natural language semantics.

Formal semantics of natural languages is based on the claim that there is no significant difference between the interpreted formal languages that are studied in logic and natural languages. An important part of this claim is that it makes sense to say that words of natural languages have a determinate meaning, which can be articulated as a model-theoretic object.

Is this claim reasonable? It would appear that the literature contains only one account that makes the claim plausible for some words: the causal reference accounts (Kripke 1980, Kaplan 1971) for proper names and natural kind terms. These accounts come with a presupposition: the (ontological) reidentifiability of the bearers of names and the natural kinds through time. If we accept this presupposition, we have an account that predicts that a current use of a natural kind term or a proper name refers to the kind or object that was baptised by the term or name by the people who came up with the term or the name. If one is rather broad-minded, it is possible to extend the account to other categories of words: names of professions, activities, etc. But while we are reasonably confident that objects and natural kinds can be reidentified, this is more problematic for these other categories.

Extending the account to the functional words and syntactic constructions is even more problematic. In this paper however, I will take the main results of logical semantics for natural languages for granted and thereby the possibility of adequately dealing with those function words that have a logical analysis. It

---

\*I want to thank Anna Pilatova for commenting on a draft of this paper

<sup>1</sup>See e.g. R. Montague, *English as a Formal Language*.

is after all the functional words on which the semantic tradition has most to say, though it should be obvious that asking for a full account of function words from logic seems misguided: there is serious scope for other disciplines as well, e.g. psychology. It should be clear however that a justification of the claim that function words have a meaning which can be articulated as a modeltheoretic object is not at all easy to find.

The problem with the causal account is that in its justification of the meanings for words in natural language, it clearly rules out that these meanings can change. The causal account has other problems: it works for successfully baptised entities, not in cases where no regular baptism has occurred or where there are unclarities as to what was baptised. The account also has no place for the information subjects may associate with a given word.

The causal theory makes it possible to justify meaning in the community of users of a language without having to attribute full knowledge of that meaning to each individual user. It is a social fact about the community that a name or a substantive word is associated with an object, a kind or a property, it is not a fact about the individual user. The attempts to conceive meaning as a property of the individual user run foul of the observation that competent users do not always have sufficient information to single out the object or the kind or property. And if there is sufficient information, it may be different information for different users.

My aim in this paper is to reinterpret formal semantics of natural language as part of a theory of the information flow that is typical for interpersonal communication. The theory starts from the intuition that in conversation there is an intention-governed information flow from one party to the other which is fully explainable in terms of properties of the individual language users involved in the conversation. The social nature of natural language is part of the explanation of why this information flow is possible, but there is no need for social meanings beyond the general rules of reference. The task of natural language semantics is to help explain the possibility of common grounds of linguistic knowledge for conversational partners that allow them to transfer information from one to the other according to the speaker's plan.

This leads to a more limited ambition for semantics. One is tempted as a semanticist to start from the concept of an ideal speaker that has full knowledge of the semantics of the language. In my proposal, the ideal speaker has a full knowledge not of social meanings but of the mechanisms of reference and of the functional part of meaning, i.e. of functional words and the meaning of syntactic combination. It is not required that she know the social meanings of each individual word.

My main thesis is the following: the analogy with an interpreted logical language is fruitful — though inaccurate— when it comes to considering micro-communication: two people engaging in linguistic communication and much less fruitful when it comes to accounting for the full community of English speakers. For the micro-community, we —and the conversational partners—

assume a common ground between the two speakers which is the basis for an explanation of why information flows.

Let us assume that *A* says to *B*:

(1) John is happy

For this to work as a communication, it must be common ground between the conversation partners that “John” refers to some entity, “happy” to some property and that “is” means that the property is attributed to the entity. It is not required that the proper (or an effective) way of identifying John or of defining or showing the property happiness is common ground between them. They can each be ignorant of either or both and it can be common ground between them that one or both is ignorant of either or both.

Even in the worst case, information would still flow, information that would only be usable for *B* when *B* acquires more information about the person called “John” or about the property of happiness. This can be spelled out in a precise way: the sentence eliminates *B*’s epistemic alternatives where the person called “John” does not have the property called happiness. Models of this common ground are models in which the following presupposition is fulfilled: where “John” denotes an object and “happy” is related to application criteria and *A* and *B* have this belief *de re* about the objects and the criteria. Each of the worlds meeting the common ground is a world in which natural language semantics applies in the sense that it meets the requirement that the particular words involved have a meaning.

Rather than assuming that words have a meaning and explaining why this should be the case, communication presupposes that it is common ground between the conversational partners that the words have a meaning. Communication does not presuppose common ground meanings as such, though there is no reason to rule out communalities in the ideas of the two participants about the meaning of the words or even a common ground about aspects of the content.

We find the minimum degree of communality<sup>2</sup> in the communication between Johnny and his sister if Johnny has to pass on the message that Mr Smith is delayed to his father at the request of his sister who is equally unaware of the identity of Mr Smith or of the meaning of delayed. For successful transfer of real information between *A* and *B* something more is required.

What we should explain is how it comes about that interlocutors share—or appear to themselves to share—such a common ground with each other. The explanation of this fact is not very difficult: *A* and *B* are speakers within a larger community of speakers in which language is frequently used. *A* uses the word “John” and maybe even knows who John is in virtue of his being a participant in the use of the word “John”. This is transparent to *B*, etc.

---

<sup>2</sup>In terms of Haas-Spohn (1994) the two kids convey exactly the formal meaning of the sentence to each other.

In virtue of the use of words in the larger community, it may be a correct belief of  $A$  or  $B$  that  $x$  is called “John”, or that “happy” is correctly applied to  $x_1$  at  $t_1$ , to  $x_2$  at  $t_2$  etc. and that  $x$  at the time of speaking is sufficiently like their private samples in the relevant respects. In fact, normally  $A$  and  $B$  will have a sample of correct uses for “happy” and will know of various proper names which indeed are used as the names of people acquainted to them.

So it being common ground between  $A$  and  $B$  that each of them is a normal language user is sufficient for it being common ground between them that various semantic facts about the words obtain. It is also sufficient for assuming that further experience with language use may increase the information contained in the exchange. And of course, it is often (“normally”) the case that the membership of the community gives a common ground of semantic facts relevant to the communication in question that goes well beyond the minimum requirement.

The inaccuracy of the claim that a natural language is like an interpreted logical language should now be clear. Rather than to a single model, the language used in the exchange is tied to the class of models of the common ground assumed in the conversation. Each of those can be thought of as containing a fuller version of English as an interpreted language. But they are all different and none of them can be described as the “true” English, if indeed there is such a thing. The speakers of English accept —as it were— the idea of English as a formal language, but they do not know which formal language it is.

From the perspective of the micro-community formed by speaker and hearer the causal theory does not find an anchoring point. What is called what depends on what others call what, not on facts of baptism that are not within the perspective of the two speakers. The success of the information exchange within the micro-community depends on the communalities between the speakers, not on their conforming to the norms of the larger language community.

## 2 Common Ground and Partiality

Semantics and semantic common grounds only are basis for or explaining information flow in conversation. That  $B$  accepts the information in  $A$ 's message is dependent on  $B$ 's assumption that  $A$  makes these sounds because  $A$  wants to convey a content to  $B$ , a content determined by the form of the message. And  $A$  produces the sounds because of  $A$ 's plan to convey that message to  $B$ .

So in my requirement of a common ground within the micro-community, I am just following Grice. In Grice's conception the speaker intends us to realise a message on the basis of us recognising her intention to convey that message. Both the speaker and the hearer must reason about their common ground with the other in order to reach the conclusions that the hearer will recognise the intention and that the speaker intended the intention. In this reasoning, syntactic and semantic beliefs (“facts”) play a crucial role.

The existence of such common grounds is not problematic, given the sort of facts that we want to have inside and the fact that they are partially based on the general use of language. Assuming a common ground is not necessary if one only wants to explain that information can flow. Some sort of decoding mechanism is then all that needs to be assumed. But common grounds are necessary if we want to explain successful communication and the fact that communicative behaviour has purposes. The speaker cannot think that her action is rational without assuming that there is a sufficiently large common ground, the hearer cannot take the word of the speaker unless she also assumes that the common ground between her and the speaker is sufficient.

The construal of semantics as above allows partiality of meaning. In particular, it allows ignorance of the meaning of words for language users. But it also leaves open the possibility that there is no complete meaning to be had: even if we increase the number of possible semantic facts to the maximal extent possible, we will not be able to reach something that could be considered by the Fregean or by the Montagovian as a proper meaning. It is however not my view that this is the general case. On the contrary: a proper name of a real person would be a counterexample and, also, a natural kind term for a properly identifiable natural kind.

It is useful to try to characterise what is involved in these cases. There is a name, e.g. “Bill Gates” which has the syntactical properties of a proper name and which is used as a name. E.g. we can infer from a person  $x$  being called “Bill Gates” by somebody that “Bill Gates” is the name of that person  $x$ . Once we have found out that “Bill Gates” is the name of  $x$ , we can use “Bill Gates” to refer to  $x$  ourselves. The existence of a name is the fact that a certain sound pattern has the properties of a name, i.e. that it is used as a name. Next to the existence of the name, we need the existence of the referent. Now in the case of Bill Gates there does not seem to be any serious doubt that he exists. And finally we need evidence that “Bill Gates” is used for Bill Gates. Most of us share evidence for the facts that “Bill Gates” is a name, that Bill Gates exists and that “Bill Gates” is a name for Bill Gates.

The same three facts obtain around the word “gold”. There are semantic and syntactic rules that make “gold” an existing natural kind term (or possibly something less specific, like a mass term), science tells us that gold really is an element and there is abundant use of “gold” as a name for gold.

While facts about the existence of a word in a category are not difficult to ascertain, the existence of the entities to which a word seem to refer is often far less clear than in the two examples we considered. And evidence as to a word applying to an entity is conditional on the robustness of the existence of the entity.

While we may feel that there is no problem with the existence of a substance like bread, there is still a much harder problem in answering the question what precisely makes a piece of bread bread. Is it the ingredients, the way of preparation, the general look and feel, its qualities as a food? It is simply false that

in the use of bread in English an implicit answer to these questions is hidden. Moreover, from a linguistic point of view, there is little that distinguishes “gold” from “butter” or “bread”.

This may be thought of as vagueness rather than as partiality. But the two are related: the lack of articulation of the concept of bread, or perhaps the impossibility of a precise articulation, is responsible for its vagueness and for the partiality of the meaning that an individual speaker has at his disposal, however well informed he is about the use of the term in question. Vagueness is partiality that cannot be overcome.

Let’s consider another word, the adjective “groovy”. It is not clear what the relevant attributes are that support the application of the word and it is equally clear that like “bread” it is a vague predicate. It is also clear that there may be serious differences of opinion as to who or what is groovy. And this happens not only with words from the realm of fashion and manners, but is pervasive with words from politics, religion and philosophy. The claim that the truth about the application of terms from these areas lies hidden inside their use in ordinary language should be denied, even while one should admit that testing a particular theory about the meaning of a word against ordinary usage or common sense can be very helpful. So I take it that partiality of meaning is an essential property of speakers of a normal natural language. Linguistic methodology cannot idealise away from it without distorting its object.

Is it possible to extend this account of information flow to the language community as such? Can we discover a common semantic ground for all the speakers of a natural language? This does not work. A language would have to be a generalisation over all the speakers. A common ground between them does not exist for the simple fact that they do not know each other. Can we assume a substantial amount of shared semantic knowledge? It seems that the communalities are rather restricted: even in the words that are important in every-day life, divergences are frequent.

We have nowadays systems of language education, dictionaries and grammars. They increase shared knowledge, but do not principally change the situation. The importance of these developments for our culture is hard to overestimate, yet that is exactly what happens if one claims that they make a difference of principle. What they do is highly praiseworthy: they tend to increase the common ground over larger sections of the population and increase its content. Thereby they enormously increase the possibilities of linguistic communication and the quality of the information that we can obtain from it. But they do not recreate our natural languages as platonic objects for the semanticist to investigate. Small communities and fragments should be the domain for our reasoning about meaning as that what makes information go from one person to the other in linguistic communication.

One last point. I have talked about rules of use for names and natural kind terms and also about general mechanisms of reference. Let me try to make that clearer. My view is that substance words can be categorised in categories which

determine both how they can be used to form larger expressions and to what they can refer given experience of their use. Syntax can in principle be taken as our guide in determining what categories there are, especially if we take the perspective of the comparative syntactician and use syntactic distinctions in one language to distinguish categories everywhere. The question to what they refer must in principle be tackled per category. Let me as an illustration try to say something about the use of the words “bread” and “sleeping”. Instances of bread are the sample which is acquired through experience of the use of “bread” and the rule for mass terms seems to say that something must be the same stuff as in the sample if “bread” is to be applied to it. If something is the same stuff, it must fit in with the sample in an unproblematic way: it must look, feel, taste, be made and used in a way which is more like the elements in the sample than things which are not bread (the negative sample). This is —I would claim— not a heuristics for identifying bread, it is the rule that governs how we use the word “bread”. The category of “sleeping” identifies it as classifying a state of humans or animals. The sample of sleeping humans or animals (including oneself at certain times) determines a way of identifying new instances of sleeping, by judging appearance (closed eyes, regular breathing), inaccessibility to external stimuli, and suspension of waking activities. It is again a question of being more like the elements in the sample than like the elements in the negative sample in a number of relevant dimensions, the dimensions that hold the sample together. And this is again how we use “sleeping”, not a heuristics, even though scientific concepts of sleeping have been given.

Sometimes we have more than similarity with a sample. This is a question of luck, deriving from the ontological status of its referent. We have this luck with the category of names for people and with mass terms that refer to proper elements. And in some other cases.

### 3 Language Change

An important aspect of natural language is the possibility to introduce new words, to make (temporary) agreements about the use of certain terms, to clarify one’s meaning and to ask for clarification. There are special devices for doing so.

- (2) Let  $n$  be a natural number such that  $P(n)$ .  
 Let us call this microbe Bertie.  
 Some models will exhibit the properties  $P$  and  $Q$ . We will call these models nice models.  
 I mean with partiality, the property of having an indeterminacy in the application of a certain word.  
 What do you mean with “causal reference”?

It may be thought that these devices are typical for mathematics and scientific discourse in general and indeed they are. But several researchers have found

that the establishment of a terminology by metaphor and local convention is typical for situations in which the conversationalists move into a new area. E.g. work on the HCRC Maze Game Corpus (Thompson et al. 1993) (a task that was specifically designed to gain more insight in referential devices) shows that speakers —after a success— employ the same device over and over again and also explicitly introduce local conventions for reference.

Discourse Representation Theory (DRT) (Kamp & Reyle 1990) provides the means for a formalisation of these relations almost without further ado. We need a way of stating that “the N” is locally a name for the discourse referent  $x$ . Clarifications of relations, or of uses of relations can be encoded in a similar way. The local character of these conventions and determinations is guaranteed since the object we are building in DRT is a representation of the current common ground. Exportation to general knowledge, or to some database representing general semantic knowledge is thereby avoided. Like proper names, other substantive words should be thought of as presupposition triggers that resolve to objects and properties that are so-called in the current common ground or otherwise in general semantic knowledge or —in the last resort— by assuming an appropriate object or property (accommodation). The spreading of local conventions —when they are not obliterated by new local conventions using the same word— to the community at large is an important source of new words coming into the language and for giving new meanings to old words.

Another important source of language change is error. My great-greatgrandfather became a victim of error involving his own name —possibly without his knowing so— as he was born a Sefath and died a Zeevat. Ever since the error was made, we and the registrars have been using the mistaken spelling. There is no sense in which the old name still applies to us and the error seems to have created a new name. The error is made possible precisely because of the partiality of the knowledge of the name by the community at large: the clerk in question was not familiar with it or possibly confused it with another name.

There are also more natural processes that lead to change. The latin word for mistress *domina* has come to mean woman (*donna*) in modern italian without any of the associations of grandeur that are still associated with the french *dame* from the same root. Probably this is a question of politeness which causes overapplication. At the same time the word *femmina* (woman in latin) has become depreciative (like the word *wijf* in Dutch which also used to mean just woman, with the current *vrouw* denoting upper class adherence in the older days). Politeness leads to overapplication of the higher forms with automatic loss of meaning, while the continuing opposition with the neutral form makes the latter become depreciative.

Another subtle effect is illustrated by the following example.

- (3) Saturday promises to become a nice day



This is a standard dutch way of indicating that the speaker has evidence for the fact that Saturday might well be a nice day (it typically like real promises allows for error). *Promise* is here a modal auxiliary and seems to derive from the possibility of underpinning the likelihood of an action of  $x$  by the information that  $x$  has made a promise to do it. See e.g. Verhagen 1997.

In my opinion, error underlies all these changes. Overapplication is an error, though a deliberate one. It constitutes a change of use which changes the common grounds in particular situations. Apart from the factors which lead to a change in use (politeness, jumping to the wrong conclusion, the tendency to misunderstand the real message in the word) there is nothing to explain.

The concept of error presupposes a norm. The norm for correct language use can be given very simply: use language as everybody else does. This is the rational kernel of the causal theory. It gives special case of the norm, for the cases where it is possible to do so. In this way, the correct way to use a name (or a correct way) requires an event in history where the name was assigned to its bearer and an assurance that our current use goes back to that event. Typically, a similar criterion cannot be given where the original use does not by itself provide a good criterion of use. The medieval butcher had quite a different profession from contemporary butchers. An application of “hip” is not impaired by the fact that none of the original users would have supported the application.

It seems that names and natural kinds are just the cases where by nature there is not much change and contemporary use reflects the intentions of the originators fairly accurately, as far as can be ascertained. It is however a *non sequitur* to infer that in this case the norm is given by the origin and faithful causal transfer. The norm can also in this case be stated as: use language like everybody else does. Conforming to the origin and faithful transfer is just what everybody else does in special cases.

Can semanticists carry out their tasks in the presence of semantic change and uncertainty on the part of language users? I can see no clear problems arising. Categorical knowledge of semantic properties of words and the meaning of certain functional words is what classical natural language semantics is good at.

There is no problematic idealisation involved in the assumption that a language user can have full categorical knowledge, since there is no reason for assuming a large set of categories. Moreover, many of the conceptual relations between words can also be reduced to categorical knowledge. So at least part of lexical semantics is unaffected by my reinterpretation.

Can natural language semantics can make a contribution to the study of semantic change? Change is not part of the traditional concerns of natural language semantics.

Some areas come to mind. The first is the area of transfer processes by which new meanings for old words come about. Here special meanings (derived from transfer processes) are sometimes able to take the place of the original meanings.

Formal theories of the transfer process are helpful in studying this change.

Second, one can expect a contribution for the study of stability. A word expressing a proper natural kind like “beech” or “man” has a far greater stability than a word expressing a profession like “grocer” or “major” or of a color word like “grey”, which in turn is much more stable than fashion words like “hip” or “groovy”. Whereas “human”, “grey” and “hip” are not unlike in their categorial characterisation and in the way they are learnt, explanations are necessary for the different degree of stability one expects. It would seem however that this is more a question of the constitution of the world than of the constitution of language.

## Conclusion

The incompatibility between natural language semantics and the phenomenon of semantic change is due to the attempt of defining meanings for the community of speakers of a natural language as a whole. I have tried to start at the other side: what does one person try to convey to another in a single communication. The meaning (in a successful communication) depends in that case on the semantic beliefs in the common ground between the two speakers.

It seems absurd to me to try to develop theories that will predict that where speakers are successful, they can yet have conveyed completely different things to each other because their use of words transgressed rules of meaning obtaining the community as a whole.

The use of language in the larger community is important as a source for the common ground in a communication, but it does not go any further than that. The source of semantic knowledge is not formed by the mythical baptisms and conventions, but by the use of language to which the communicators have been exposed.

Change of meaning is here any social process that changes the use of words, and thereby the common grounds that communicators will share. This will automatically lead to a change in the information that is normally expressed by the same sentence before and after the change.

## References

- Grice, H. Paul, 1957. Meaning. *Philosophical Review*. 66 p.377-388.
- Haas-Spohn, U. *Versteckte Indexikalität und subjektive Bedeutung*. PhD Universität Tübingen 1994.
- Kamp, H. & U. Reyle. *From Discourse to Logic*. Kluwer, Dordrecht 1990.
- Kaplan, D. Quantifying in. In: Linsky (ed). *Reference and Modality*. Oxford 1971.
- Kripke, S. *Naming and Necessity*. Oxford 1980.
- Montague, R. English as a Formal Language. In: Thomason, R.H. (ed.) *Formal*

Philosophy: Selected papers of Richard Montague . Yale UP, New Haven 1974.

Thompson, H, A. Anderson, E. Bard, G. Doherty-Sneddon, A. Newlands & C. Sotillo. The HCRC Map Task Corpus. Natural Dialogue for Speech Recognition. In: F. Bates (ed) Proceedings of the DARPA Human Technology Workshop I. ARPA/ISTO 1993.

Verhagen A. "The girl that promised to become something" A exploration into diachronic subjectification in Dutch. In: T.F Shannon and J.P. Snapper(eds.). The Berkeley Conference on Dutch Linguistics. Lanham, 1998.