

H030
P853

What is Meaning?

Fundamentals of Formal Semantics

Paul H. Portner

 **Blackwell**
Publishing

© 2005 by Paul H. Portner

BLACKWELL PUBLISHING

350 Main Street, Malden, MA 02148-5020, USA

108 Cowley Road, Oxford OX4 1JF, UK

550 Swanston Street, Carlton, Victoria 3053, Australia

The right of Paul H. Portner to be identified as the Author of this Work has been asserted in accordance with the UK Copyright, Designs, and Patents Act 1988.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, except as permitted by the UK Copyright, Designs, and Patents Act 1988, without the prior permission of the publisher.

First published 2005 by Blackwell Publishing Ltd

Library of Congress Cataloging-in-Publication Data

Portner, Paul.

What is meaning? : fundamentals of formal semantics / Paul H. Portner.

p. cm. – (Fundamentals of linguistics)

Includes bibliographical references and index.

ISBN 1-4051-0917-3 (hardcover : alk. paper) – ISBN 1-4051-0918-1

(pbk.: alk. paper)

1. Semantics. 2. Semantics (Philosophy) I. Title. II. Series.

P325.P635 2005

401'.43–dc22

2004015923

A catalogue record for this title is available from the British Library.

Set in 10/12¹/₂pt Palatino

by Graphicraft Limited, Hong Kong

Printed and bound in the United Kingdom

by MPG Books Division, Bodmin, Cornwall

The publisher's policy is to use permanent paper from mills that operate a sustainable forestry policy, and which has been manufactured from pulp processed using acid-free and elementary chlorine-free practices. Furthermore, the publisher ensures that the text paper and cover board used have met acceptable environmental accreditation standards.

For further information on

Blackwell Publishing, visit our website:

www.blackwellpublishing.com

1 The Fundamental Question

Semantics is the study of meaning. Those who make this field the center of their academic lives need to enjoy being seen as unimportant, because many people they meet during daily life believe that the study of meaning cannot be a serious, much less a scientific, pursuit. For the general public, an issue is “just semantics” if it has to do with mere matters of form, and accusing someone of focusing on semantics is a way of saying they want to avoid the heart of the matter. An interest in semantics sounds like a deep concern for details of etiquette: perhaps useful in certain contexts, like the diplomatic party or ivory tower, but something which tough-minded realists will brush right aside.

The reputation of semantics within its broader field of linguistics is the opposite of this popular one in many ways. Semanticists tend to use a lot of tools drawn from logic and even mathematics as they go about their jobs as university professors (and almost all of them are university professors; the rest mostly work for computer companies). Because of this, they tend to write down their ideas using all sorts of funny symbols (such as λ , \exists , and \forall). We call this approach to language *formal*, meaning that it is couched in these logical/mathematical languages. The formalism of semantics makes the field virtually impossible for the uninitiated to understand. It also makes it very difficult for those who don't have a native talent for abstraction and logic to become initiated. In other words, semantics is thought to be really technical and hard. (You can imagine what this means for the image of semanticists.)

Though semantics as practiced by the specialist is formal and abstract, at least it is this way for a reason: Formalization allows the construction of very precise theories, and precise theories are better because they don't allow the theorists to fudge the data quite so easily as less precise theories do. Semanticists have to be especially cautious in this respect, I believe,

because discussion of meaning can very quickly turn into pointless contests between vague but strongly held opinions – exactly the sort of situation which the popular view of “it’s just semantics” rightly mocks.

Let me give an example: suppose we wish to understand the difference in meaning between a simple past tense sentence like (1) and a present perfect sentence like (2):

(1) Mary received the most votes in the election.

(2) Mary has received the most votes in the election.

After thinking about the two examples for a while, an intuition about the difference may emerge. One common intuition is that (1) simply reports a past event, while (2) reports both a past event and a current result of that event, such as that Mary will be the next president. Armed with this intuition, one can then start looking at other present perfect sentences, and the idea that they report the present results of past events may seem better and better (*Shelby has finished his dinner* – so he’s no longer hungry, *I haven’t slept for days* – so I’m very tired, . . .). Then suppose we encounter the conversation in (3):

- (3) Speaker A: Will Mary be able to finish Dos Passos’ *USA* trilogy by the next book club meeting? It’s so long!
 Speaker B: Well, she has read *Remembrance of Things Past*, and it’s even longer.

There is a current fact which speaker B is pointing out: that Mary might indeed finish the *USA* trilogy. Our intuition seems to be confirmed. However, we think about it some more, and there’s a problem. The fact that Mary might finish the book isn’t really a result of her reading *Remembrance of Things Past*. The fact that she read *Remembrance of Things Past* is merely evidence for what she might be able to read in the future. Is there something else which speaker B’s sentence is trying to report, some other situation which really is a result of her reading *Remembrance of Things Past*? One suggestion: the result is that she has read a really long book before. But notice that this alleged result is itself reported with a present perfect sentence (*She has read a really long book before*), suggesting that our explanation is getting circular, and in any case this “result” is rather ephemeral. Any other possibilities? There are plenty of concrete results of her reading *Remembrance of Things Past*: she has confidence in her ability to read long books; she knows who Albertine is; and infinitely more.

But this raises another problem: How does speaker A know which of these speaker B has in mind? This debate isn't going to get anywhere unless somebody provides a precise explanation of what they mean by the term "result." A formal theory of events and results and time might well be of help here.

Returning to the main thread of discussion, formal theories have some other advantages as well. They are good for implementation in computational settings and for exploring the relations between semantics and the other sub-disciplines of linguistics which have a formal flavor, like syntax, as well the interdisciplinary field known as cognitive science, in particular computer science, psychology, philosophy, and neuroscience. While indirect, these are also good reasons to study semantics in a formal way.

Nevertheless, even if you grant that abstraction and formalism are excellent things for those working at the frontiers of research, you may still want to think about the nature of meaning from a scientific perspective while not taking on the task of learning a lot of formal logic. The goal of this book is to help you do that. It is also designed for those of you who may be undertaking to learn the technical side of semantics, and feel you need a bit more guidance as to what those formulas are really all about.

In order to study meaning, we have to begin with some basic understanding of what sort of thing a meaning is. Trying to answer the question, "What is a meaning?" in its most general sense is a scary task mostly undertaken by philosophers! But we need to begin this book with at least a few basic considerations.

1.1 What is a Meaning?

In ordinary life, we sometimes find ourselves talking about the meanings of words and sentences. For example, I am at home reading, but find myself confused. I don't know what something means. If I ask my wife the meaning of a word, she will give me an answer:

"What does 'kakapo' mean?"

"It's a kind of parrot."

Or,

"What does this character mean?" → 狗

"It means 'dog'."

When I ask what a word means, I typically get more words – perhaps in the language I used in my question, perhaps in another, but nonetheless more words. Can meanings be words? The answer is obviously “no,” if we want to approach meaning as scientists. Because, supposing that the meaning of “kakapo” is “a kind of parrot,” what about the meanings of “kind” and “parrot”? More words. Eventually identifying the meaning of a piece of language with more language is bound to become circular, as a word is defined in terms of some of the very words which it helps define.

Pick a simple word and look up its definition in the dictionary. Then look up the crucial words in that definition, and so forth. Do the definitions become circular?

A more sophisticated view similar to this one is known as *meaning holism*. Most famously supported by the philosopher Quine,¹ the theory of holism claims that the meaning of a word or phrase or sentence depends on its relationships with other words, phrases, and sentences. For example, it might be argued that part of what makes up the meaning of *tall* is that it's opposed to the meaning of *short* (something that seems quite plausible). More precisely, holist theories tend to be *functional* in the sense that it is some aspect of the use of a piece of language which makes for its meaning, so we should really say that part of the meaning of *tall* is that if you call something *tall*, you should not at the same time call it *short*, and if you call something *tall* you should be willing to also call it *not short*. The big issue for holism is to find a way to say which of the relations among words, phrases, and sentences are important to semantics. Radical holism takes the position that there is no line to draw between connections of this plausible sort (*tall* with *short*) and all of the connections among words, phrases, and sentences (*tall* with *I like beans and so I plan to make red bean soup for dessert*). In that case, the semantic system of a language will be a complex, interconnected network, and all meaning will be relative to the whole system. It's difficult to see how meaning can be studied in a scientific way from this perspective. (I should say that it's hard to assess whether holism is in general incompatible with scientific linguistic semantics because over the years a wide variety of theories have been labeled as “holistic.”²)

Linguists who study meaning don't believe that the study of meaning should be unscientific. They feel this way in the first instance (I believe) not because they are better philosophers than Quine and his followers, but because their experience with language shows them that the way languages

express thought is not as arbitrary as the holist's way of looking at matters would lead you to expect. Instead, they find overwhelming evidence for deep and consistent patterns in how languages express meaning, patterns which are in need of scientific explanation. We'll see some of these patterns in the chapters to come. Thus, though the initial intuition that we started with, that the meanings of words involve their relations with other words, is probably correct in some sense, it does not by itself provide a basis for the scientific study of linguistic meaning. Therefore it has not been incorporated much into the thinking of formal semanticists. It is more relevant to the practice of the field of semiotics, the study of symbolic systems generally (including language to the extent that it has something in common with such things as the "meaning" of foods and clothes).

If meanings aren't words, our next guess might be that meanings are something in the mind: concepts, thoughts, or ideas. When you understand the meaning of the word *dog*, your mind (and brain) change in certain ways. At some point you form a concept of dogs (let's indicate the concept with capital letters: DOG). Then, you associate this concept with the English word *dog*, and from then on you have an ability to use the word *dog* whenever the concept DOG is active in your thoughts. From here it's a short step to saying that DOG is the meaning of *dog*. This theory would explain the patterns in how languages express meaning in terms of the nature of concepts, and perhaps ultimately in terms of the way the brain is structured. Let's call this perspective the idea theory of meaning.³

One obvious challenge for the idea theory is to come up with a sound psychological theory of what concepts and ideas are. This psychological theory needs to provide a concept or idea for every meaningful piece of language. Thus, there will need to be ideas and concepts associated with each of the following (at least in any situation in which they are meaningful):

Dogs and cats
The picture of my wife
Three
Whatever
The president lives in Washington, DC.
Had been sleeping
Why
Who said that we had to be at the airport so early?
 -ed (the past tense marker)

The idea theory needs to say what idea is associated with *whatever*, *why*, or *three*, and this doesn't seem as easy a project as explaining what idea

is associated with *dog*. At least, the idea theory provides no quick and easy path to a complete theory of meaning. But even if it's not going to be easy, the idea theory may work. Certainly, something is going on in our minds when we use words and phrases, so in some sense there are ideas associated with all meaningful language. Don't we just have to discover what they are (hire more psychologists!) and use them to explain meaning?

The question which the scientist of meanings needs to ask is not simply whether our concepts and ideas play a role in how we use language in a meaningful way – of course they do. The real issue is whether those concepts and ideas have the right properties to explain everything we need to explain about meaning. In other words, we can consider what we know about meaning already, and then check out whether the idea theory is consistent with that knowledge. Well, what do we know about meaning? Here are some basic points:

- i Sometimes pieces of language have the same meaning – they are SYNONYMOUS.

Dog

Canis familiaris

Mary kissed John.

John was kissed by Mary.

- ii Sometimes pieces of language conflict with each other in terms of their meanings – they are CONTRADICTIONARY.

The pig is on top of the turtle.

The turtle is on top of the pig.

- iii Sometimes the fact that one piece of language is an accurate description of a thing or state of affairs automatically guarantees that another is an accurate description of it too – the first ENTAILS the second.

Robin

Bird

The circle is inside the square.

The circle is smaller than the square.

The idea theory can say a bit about what it is for *Mary kissed John* and *John was kissed by Mary* to be synonymous. Suppose that in my mind I have ideas, or concepts, of Mary, of John, of kissing, and of "pastness." These ideas are combined into some kind of aggregate idea, the one associated with the sentence *Mary kissed John*. The idea theory would then want to say that *John was kissed by Mary* is associated with same aggregate idea,

and this is why the two sentences are synonymous. The second sentence has all the same pieces as the first, put together in a different order, plus an additional one, the passive voice (the fact that the sentence takes the form . . . *was kissed by* . . . instead of . . . *kissed* . . .). In some way, the idea associated with the passive voice exactly undoes the effect of putting *Mary* as the subject of the sentence in one case, and *John* as the subject of the sentence in the other. In other words, the sentence without passive voice, *John kissed Mary*, is not synonymous with *Mary kissed John* – clearly – and it’s the passive voice which gets into the aggregate meaning of *John was kissed by Mary* and sets things right.

I think it’s clear that the meaning associated with the passive voice will not be the kind of thing that we typically call an “idea.” This meaning has a grammatical nature, having to do with the order in which John and Mary are mentioned in these sentences. This suggests that, if the idea theory is to work at all, ideas will have to have a language-like nature. That is to say, because the meanings of certain pieces of language are deeply tied into the grammar of that language, ideas themselves will need a grammar. If our ideas have a grammar, they are a language, and we are thinking of them as a *language of thought*.⁴

An important objection to the idea theory arises from the famous “Twin Earth thought experiments.”⁵ The basic idea of the Twin Earth thought experiments is that we can learn a lot about the nature of the mind and language by imagining a world which is exactly like our earth except for some specific differences, and then examining whether those differences seem relevant to how our minds or language work. The following kind of Twin Earth thought experiment is relevant to whether we should accept or reject the idea theory of meaning. One thing we’re absolutely sure about is that when we use a common term like *water*, we are referring to a very definite kind of thing in the natural world. In fact, when we use *water*, we’re referring to H₂O. Now, imagine some people inhabiting a planet very much like ours. This planet is, in fact, so very much like ours that if you went there, you couldn’t tell you weren’t on earth. Everything on Twin Earth is just like on earth. There’s even a copy of you there (and while you’re visiting Twin Earth, he or she is visiting our earth). This Twin Earth only differs from earth in one way, and that is that everywhere we expect to find water there, we find another substance, XYZ, which looks, feels, and acts just like water, but which is actually not water. So, obviously, when the Twin Earth people who speak a language very similar to English say *water*, they are referring to XYZ, and not H₂O. This implies that Twin Earth English and our earth English are not quite the same language. The word *water* differs in meaning between them.

Since Twin Earth is just like earth, each person on earth has a twin there. And this twin is exactly like the true earthling in every respect (except that instead of H_2O , there is XYZ in his or her body). Importantly, the mental life of the earthlings and their twins on Twin Earth are identical. Despite the fact that their mental lives are identical, they don't mean the same thing by the word *water*. This implies that the meaning of *water* is not determined solely by the what goes on inside the head of people who use the language, and this is just to say that meanings are not ideas. Rather, what a word refers to is partially dependent on the environment in which people who use a language live. Given that you live surrounded by H_2O , *water* will refer to water and not XYZ. (If you want to make Twin Earth a bit more consistent, you can say that people who know a lot about chemistry there know that XYZ is common there, but H_2O is not, and conversely on real earth. All that matters for the thought experiment is that some people – not knowing much about chemistry – have no clue about what XYZ or H_2O are. These people and their twins will have the same internal mental concepts associated with the word *water*, but will mean different things by it.)

Another important conclusion about the idea theory can be drawn by considering the meanings of individual words like *dog*. I have formed a concept of dogs through my experience with them. In particular, when I think of dogs I often think of the fluffy gray-and-white, mid-sized ones called Keeshonds, because this is what my dog Shelby is. For me, DOG has many components drawn from Shelby's appearance and behavior. Most people, though, have never heard of or seen a Keeshond, as I can attest from Shelby's celebrity whenever we go for a walk through town. Let us call these people "the unfortunate ones." The concept of DOG held by the unfortunate ones lacks many of the most prominent features of my concept DOG. Since we have different concepts of dogs, and according to the idea theory the concept equals the meaning, this would seem to imply that the word *dog* means something different for me and for the unfortunate ones. So, when they ask me "What kind of dog is that?" their question actually doesn't mean the same thing for me and for them. This conclusion seems somewhat implausible, for even if two people with different experiences with dogs ask that question, they are nevertheless probably looking for the same kind of answer from me.

You might say that my concept of DOG and the unfortunates' concept are similar enough to count as the same because of the fact that they are concepts of the same things out there in the world, namely the members of the species *Canis familiaris*. In other words, all the unfortunates and I aim to use our concept of DOG to classify and think about the members of that species, and on this ground our concepts are different versions of

the same thing. But this concession weakens the idea theory tremendously. It turns the idea theory into a side-trip on the road of understanding meaning. As illustrated in diagram 1, if we're going to explain the meaning of *dog* in terms of the concept DOG, and the concept in terms of the animals which it describes, we might as well explain the meaning directly in terms of the things, avoiding the detour through ideas. This is the view which is adopted within the theory of meaning presented in this book.

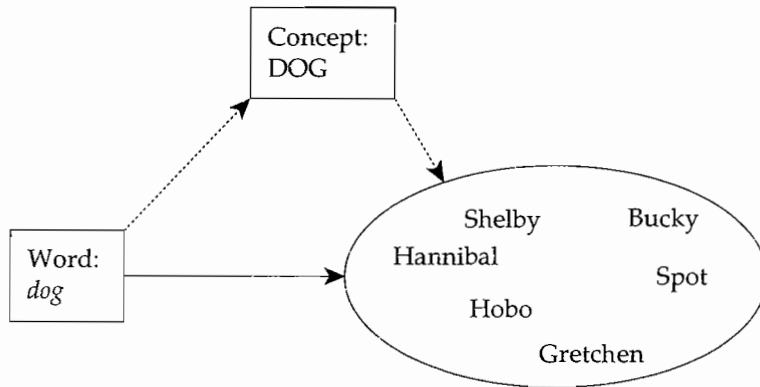


Diagram 1

The reason that the word *dog* means the same thing for you and for me is not that we have the same mental constructs relating to the word. Rather, it's because of our intention to apply the word *dog* to the same things out there in our environment, namely the dogs. (Of course we may not agree about precisely which things these are. You may honestly think that Shelby is a kind of small bear, and refuse to call him a dog. But still our intentions are the same: to call everything by the word *dog* which really is a dog. It's just that you are failing to do so, due to ignorance.)

Notice, by the way, that none of this implies that we don't have concepts of dogs and other things – of course we do! – or that these concepts aren't crucial in how we use meaningful language in the correct way. It doesn't imply that we lack a language of thought (or that we have one). It doesn't mean that semanticists can't contribute to studying the psychology of language. In fact, I am convinced that semantics has an important part to play in the study of cognition. It simply says that meanings aren't ideas. Whatever the relation between meanings and ideas may turn out to be, they aren't the same thing.⁶

If meanings are not in the relations among pieces of language, and are not ideas, what else could they be? Another type of answer which is influential within philosophy says that they are social practices.⁷ The idea is that when

somebody says something, it should be thought of as a kind of move in a giant language game which we all play. This game has rules, and these rules imply things like "If somebody makes the move of saying 'What time is it?' to you, an appropriate move for you to make is to say 'It's 6 o'clock' (if it is 6 o'clock)." Of course this is, in fact, an appropriate response, so any type of semanticist is going to want an understanding of meaning which explains why it is appropriate. But an advocate of the social practice theory of meaning goes beyond this to argue that this is all there is to meaning. So an important task is to explain precisely how each meaningful piece of language comes to have the roles in the language game which it has. Take for example the response "It is 6 o'clock." It is an appropriate response only if you think it is in fact 6 o'clock, and this is somehow based on the fact that it is comprised of the pieces *it*, *be*, present tense, *6* and *o'clock* put together in a particular way. (If these same pieces were put together as "Is it 6 o'clock?" it would not be an appropriate response.) Perhaps we can say that there is a social practice (a rule of the language game) which says something like this:

If the previous move was an utterance of "What time is it",
 then an utterance of the form "It is X o'clock" is a candidate appropriate move, and
 if the clocks in the neighborhood look roughly like this: ⓐ,
 then it is an appropriate move to make the utterance with
 X filled in as "6."

This social practice theory has not had much impact on linguistically oriented semanticists for three reasons. First, there seems to be a fundamental conflict with one of the basic insights of modern linguistics (not just semantics) that our languages are not arbitrary social conventions, but rather reveal deeper universal patterns which spring from the way our minds are built. These patterns call out for scientific explanation, and it seems hard to explain them on the basis of the notion of game-like social practices. In this way this perspective on meaning suffers from the same problem as the first approach we looked at, the one which said that the meaning of an expression comes from its relations with other expressions of language. However, it is not clear to me that there is a fundamental conflict here. It could be that a follower of the social practice theory could say that language is a very special kind of social practice which is instinctual and which has rules that are to a large extent determined by a part of our brains which is responsible for making us follow them.

The second reason that the social practice theory has not been seen as relevant to linguistic semantics is that it does not seem to provide an

important role to the intuition that “It is 6 o’clock” is an appropriate response because it *is* 6 o’clock and you want me to believe this. It is this aspect of language which seems to distinguish it from true games. In baseball, a player will try to catch the ball before it hits the ground because this will help achieve the goal of winning, and that’s it. But you answer my question not just because the rules of the language game tell you that you have to (or else you “lose”), but because you recognize that I want to know a certain piece of information, which I can get if you say “It’s 6 o’clock.” Any other way of giving the same information, or other relevant information, would also be an appropriate move (think about *It’s dinnertime* and *I don’t know*). Gathering these moves together under a collection of rules which tell you what you can do in response to my making the move of saying “What time is it?” would just miss the reason why they are appropriate. So advocates of the social practice theory must provide a notion of “the information provided by a sentence” which is compatible with the fundamental idea that language is a social practice. This is a difficult task, and there is controversy about whether they can succeed. If they can, then it is possible that the social practice theory is compatible with the ideas about meaning presented in this book.⁸

The third reason that linguistically oriented semanticists tend not to pursue the social practice theory is that they feel that we already understand a great deal about how meaning works in language. The main point of this book is to introduce you to the most important of these insights. But these insights have not been achieved in terms of the social practice theory, or the other theories discussed above for that matter. Rather they have been achieved in terms of some ideas about meaning which I’m about to get around to introducing. Until some other way of thinking about meaning shows that it helps us understand something about how language works that we didn’t understand before, semanticists – like other scientists – will see little reason to change.

1.2 Meanings are Out in the World

This discussion so far points to the conclusion that meanings are not internal to language, are not in the mind, and are not merely social practices. Rather, they are based in language- and mind-external reality. The meaning of the word *dog* implies that it describes all of those things that actually are dogs, regardless of our ability to define it with words or to formulate an appropriate mental concept.⁹ The point is even simpler to see in the case of names. The name *Confucius* refers to the ancient Chinese philosopher, and this is the basis of its meaning (indeed this may