

Language and Meaning

language survives everything – corruption, misuse, ignorance, ineptitude. Linking man to man in the dark, it brought man out of the dark. It is the human glory which antecedes all others. It merits not only our homage but our constant and intelligent study. [Anthony Burgess: *Language Made Plain*, revised edition, Fontana/Collins, 1975]

We use language so effortlessly, so automatically, and so unreflectively that it's easy to forget what an extraordinary phenomenon it is. A few sounds or marks can *mean* something; they can be used to enable us to exchange thoughts of a complexity that has no upper bound; they can be used in some of mankind's greatest aesthetic creations; and of course they can be used in the most humdrum ways as well. How is language possible? What breathes meaning into what are otherwise mere sounds or marks? Do dolphins have a language? Could a creature without language think? Is language use a uniquely human accomplishment? Do language users deserve special moral consideration?

I'm not going to be able to address (let alone answer) all these questions. Rather, I'll offer a tour of three main themes from recent philosophical approaches to language. The first theme is rather surprising: I consider two versions of the view that there can be no such thing as meaning, at least as we ordinarily conceive it. Willard van Ormond Quine, a philosopher at Harvard for fifty years, reached this conclusion by reflecting on the evidence we would need in order to be justified in ascribing meaning. Ludwig Wittgenstein, a very influential Austrian-born philosopher who did most of his philosophical work while living in England (first in Cambridge, then in Oxford), challenged our ordinary conception of meaning as involving an indefensible notion of rules which determine in every situation how a word is to be rightly applied: in his striking metaphor, "rails laid to infinity". His challenge has been sympathetically elaborated more recently by Saul Kripke.

The second theme I'll address is the systematic work that has been done in describing meaning in language, work which to some extent presupposes that the kinds of skepticism discussed in the first theme of this chapter are misguided. The work itself involve quite a lot of technicality, but I can give the general flavor of two of approaches without much technical detail. One approach is based on the idea that the meaning of a sentence can be given by the conditions under which it is true, in slogan form: meaning is truth conditions. The other approach is based on the idea that the basic notion of meaning can be defined in terms of speakers' intentions.

The third theme is the relation between the meanings of sentences and what we can reliably use them to communicate. For example, suppose I ask you whether you like music and you reply "I can tell the difference between loud and soft".¹ I would be right to treat you as meaning that you did not like music, although this is not what your words meant, and you did not literally say you did not like music.

We use the word "means" and its cognates for various phenomena with no special connection with language. In "Those spots mean measles", we invoke a notion of meaning centered on the idea of indicating or being a sign of. In "He meant no harm", we invoke a notion of meaning in which it amounts to intending. It may be that both aspects are invoked in linguistic meaning: there's a connection with indicating and being a sign of, and also a connection with the intentions of speakers. However described, most of us believe that we possess a sound and distinctive notion of linguistic meaning. The next section considers challenges to this natural belief.

Skepticism about meaning

Language is a social art. (Quine 1960: ix)

Quine on rabbits

¹ Gilbert Ryle, a famous Oxford philosopher in the years after the Second World War, is said to have given this reply to the question.

Imagine you are a linguist in the field, studying a people who speak a radically alien language (there are no bilinguals or similar languages to help you get started). Let's suppose you notice that, mostly, when the aliens utter one of their strange words, "Gavagai", there's a rabbit in their vicinity; and, mostly, if an alien notices a rabbit, she utters "Gavagai". This behavior makes "Rabbit" a sensible tentative translation of "Gavagai". No doubt there will be times when "Gavagai" is uttered when there is no visible rabbit: perhaps the speaker sees old rabbit-tracks, or wishes she had a rabbit. No doubt there will be times when a rabbit is in plain view but no one says anything: perhaps silence is required for rabbit catching, or the aliens see no reason to state the obvious. But it's natural to assume that these more complex data could be factored in, and in the end we would have a solid case for the hypothesis that "Gavagai" means "Rabbit".

This is the natural view that Quine seeks to undermine. One main line of argument is that we have simply assumed that the aliens share a metaphysical picture of the world that resembles ours. But suppose they think of the world as made up of momentary things, collected together into bundles. What we think of as a rabbit is thought of by them as a collection of temporal parts of rabbits, a "rabbit-at-an-instant-of-time". The data about "Gavagai" equally support the view that it is true not of rabbits, but of temporal stages of rabbits. Or, of undetached parts of rabbits (for a rabbit is present just when an undetached rabbit part is present); and so on. There are endless distinct possibilities, and the data don't sort out which is correct.

Quine admits that if we already shared some parts of language with the aliens, we could reach a more determinate view. For example, if, when a single rabbit is present, we could ask an alien how many gavagai there are, any plural answer (e.g. "Five" or "Many") would count against the hypothesis that "gavagai" means "rabbit" and in favor of some such hypothesis as that it means "undetached rabbit part". But how are we to justify our supposition that we have really asked how many gavagai there are? Perhaps the alien words for "how many" which we use really mean "how many undetached parts of", in which case a plural answer is, after all, consistent with the view that "gavagai" means rabbit.

Quine calls the difficulty of grounding a determinate meaning for a term in the behavior of its users the "inscrutability of reference". His target is to establish a related thesis, the "indeterminacy of translation". He claims that any sentence of a language can be translated with equal correctness into both of two sentences which, intuitively, have different meanings. Perhaps, so far as any observable data about how the aliens behave are concerned, the one-word sentence "Gavagai" can be translated equally well as "Lo, a rabbit!" or as "Lo, an undetached part of a rabbit!". But the translations don't mean the same as each other, so they can't both mean the same as "Gavagai!". Quine invites us to conclude that the notion of *meaning* which underpins this contradiction is unsound. We must replace it by something closer to observable reality, and less governed by what he calls the "museum myth", according to which a meaning is displayed as some one determinate object among others.

The discussions of Quine's indeterminacy thesis are voluminous and detailed, and I will not attempt to indicate even their most salient features. But here's something we can usefully carry forward: Quine thinks that meaning is a property of behavior, of the use of language. This can be interpreted as an anodyne thesis, if no limitation is placed upon how behavior or use is to be regarded. An anodyne thesis connecting meaning and use is that the meaning of a word is a matter of what it is conventionally used to *mean*. Quine, however, thought that meaning needs to be a property of behavior which can be specified independently of semantic terms, like "meaning" or "reference". That plays a crucial role in his discussion, but it is a challenging and controversial thesis. One moral we might draw from the implausibility of the conclusion he reaches is that he is wrong to admit only such an impoverished conception of behavior as the basis for meaning. Perhaps it's wrong to think of meaning as reducible to "observable behavior" thought of as Quine thinks of it: the motion of matter through space. Perhaps we have to view the world, notably the behavior of others, in terms of meanings if we are to understand what it contains.

Wittgenstein and Kripke on invisible rails

Whence comes the idea that the beginning of a series is the visible section of rails invisibly laid to infinity?
(Ludwig Wittgenstein, *Philosophical Investigations*)

We all think there's a definite fact concerning what we mean by "+" (the addition function in arithmetic). Ludwig Wittgenstein, however, broaches skeptical reflections designed to undermine our confidence in this opinion. The essence of this skepticism is that it seems there is nothing we can point to which makes it the case that we mean addition by "+", rather than some related function. Our behavior would have been the same, even if we had been

applying a function which differs from addition with respect to some relatively large numbers, ones we have never used in a sum.

This Wittgensteinian skepticism is made vivid by Saul Kripke (*Wittgenstein on Rules and Private Language*, 1982). Kripke invites us to consider an alternative function, quus (“ \oplus ”), which is defined as follows:

$x \oplus y = x + y$ when x and y are less than 57; otherwise, $x \oplus y = 5$.

The skeptical question can then be posed as follows: what facts about you make it the case that you’ve meant addition by “+” rather than quaddition (the function signified by “ \oplus ”)? Let’s suppose that you have never done a sum involving numbers larger than 57. So even if you’ve got all your sums right, you’ve been behaving just as you would behave if you had been quadding rather than adding. Your behavior doesn’t mark you out as an adder rather than a quadder.

The most natural first answer is that I intended to use “+” for addition, not for quaddition. For this to be so, however, I must have been in a mental state with the content *addition* rather than the content *quaddition*, and the original question arises again: what makes my intention have the one content rather than the other?

The second most natural answer is that I exemplify different dispositions depending on whether I am an adder or a quadder. If I’m an adder, I’m disposed (even if the occasion has never in fact arisen) to answer the question “68 + 57?” by “125”. If I’m a quadder, I’m disposed to answer the question by “5”. Isn’t that difference enough?

Kripke says not, on the grounds that what the skeptic required was an answer that would *justify* the claim that “125” is the right answer (which it would not be if the relevant function were quaddition). Concerning the totality of dispositional facts, he writes “How does any of this indicate that—now *or* in the past—‘125’ was an answer *justified* in terms of instructions I gave myself, rather than a mere jack-in-the-box unjustified and arbitrary response?” (Kripke 1982: 23). At least at this point in the discussion, he thinks the non-skeptic cannot answer the question, and he sums up the predicament thus:

ultimately, if the sceptic is right, the concepts of meaning and of intending one function rather than another will make no sense. For the sceptic holds that no fact about my past history—nothing that was ever in my mind, or in my external behavior—establishes that I meant plus rather than quus. ... But if this is correct, there can of course be no fact about which function I meant, and if there can be no fact about which particular function I meant in the *past*, there can be none in the *present* either. (1982: 13)

Like Quine’s skepticism about meaning, this version (Kripkenstein’s as they say) has been widely discussed, and many different responses have been offered. Kripke himself seems to think that relief from skepticism can be found by reflecting on the social character of language, the fact that we are engaged in the practice as a community. Not everyone has understood why many speakers are somehow better able to resist the skeptic than one.

Another line of response is that two questions have been run together under the heading “What’s the right answer to ‘68 + 57?’”. There’s the question “Which function is being picked out by ‘+’: addition or quaddition? Then there’s the question: given which function is at issue, what’s the correct answer? One view is that dispositions can give a sufficiently determinate answer to the first question. If the person answering the question is generally disposed to answer “5”, and is not disposed to accept corrections, that’s in part constitutive of her being a quadder; *mutatis mutandis* if the disposition is to answer “125”. Justification for the correctness of the answer is a mathematical matter which arises only when it is determinate which function is intended. That is the second question, and is quite different from the first.

This response can be challenged on the grounds that it fails to do justice to the normativity of meaning. Meaning imposes requirements which someone who grasps it may fail to meet. We might mean addition by “+”, but still give the *wrong* answer, say 121, to the question ‘68 + 57?’. Hence we cannot read off what function is being invoked just from the subject’s dispositions to use an expression, and the dispositions can’t be even partially constitutive of what the expression means.

We need to re-raise the thought with which we concluded the discussion of Quine’s skepticism. The skeptic takes for granted that an account of meaning can be given in other terms (for example, dispositions to use). It’s not clear that this reductive approach is justified. No doubt meaning supervenes in some way on use, in that two communities whose use of every expression coincides arguably mean the same by these expressions. But we may not be able to

tell exactly how the supervenience works, so we may not be able to answer the question: what facts about use make these facts about meaning as they are? Our ignorance does not mean that there is nothing to know, nor does our ignorance of the details entail that we don't know that the general fact that meaning supervenes on use.

Meaning and truth conditions

To give truth conditions is a way of giving the meaning of a sentence
(Donald Davidson "Truth and meaning")

Let's suppose that we have somehow managed to resist skepticism about the very notion of meaning. The question arises: how should meaning be described? This question divides into two: (1) What central notions is it appropriate to exploit in describing meaning? (2) What kind of structure should we use?

Let's take a whole sentence, say "Snow is white", and consider what notions it would be legitimate to bring to bear to describe its meaning. In a sense we can't go wrong if we stick with the notion of meaning itself. Perhaps we could say that "Snow is white" means that snow is white. (This might look a bit trivial but it's something a monolingual French speaker doesn't know, so it can't be all that trivial!) We might think that this takes the notion of meaning itself too much for granted. Just as Quine thought we should be able to describe linguistic behavior without using such terms as "meaning", so it's often supposed that we should be able to describe the meanings of sentences without using such terms. Alternative terms may cast light on the notion of meaning itself.

Most people agree that one understands a declarative sentence like "Snow is white" only if one knows what it would be for the sentence to be true. A more controversial, but still widely held opinion, is that the converse obtains also: if one knows what it would be for a declarative sentence to be true, then one understands it. This suggests an approach to meaning which had a large impact in the twentieth century: meaning is truth conditions, because knowledge of meaning is knowledge of truth conditions. Tradition has it that this view is to be traced back to Gottlob Frege in the late nineteenth century. It's most salient advocate, in the second half of the twentieth, is Donald Davidson. The view needs to be understood in a rather nuanced fashion or else it is obviously false. Consider two necessary truths, for example, "Every triangle has three sides" and "Nothing is red and green all over". Under what conditions are these true? Under all conditions; that is, on one understanding of "truth conditions", they have the same truth conditions. Yet they evidently differ in meaning. We'll see shortly how Davidson provides the nuance which enables him to avoid this instant refutation. First, let's see how we might arrive from a different direction at the view that meaning is truth conditions.

We can't hope to describe the meaning of every sentence in a language using a sentence-by-sentence approach. That's because of the "creative" character of language use, which in turn depends upon the "compositionality" of language itself. Any of us can put words together in an entirely new way, producing a sentence which has never been used before. (The sentence you have just read might be one such.) As speakers we have no difficulty constructing, and as hearers no difficulty understanding, such novel sentences. That's creativity. The explanation is that the meaning of a sentence is determined by the meanings of the parts of which it is composed. Once we have understood a relatively small vocabulary and a relatively small number of ways in which words can be meaningfully strung together, we are equipped to understand any of the huge number of sentences that can be made up from those parts. That's compositionality.

In describing meaning in a language, we need to say something about the meaning of its words, and something about ways in which words can be combined; this should deliver information about the meaning of all the sentences which can be formed from these words and these modes of combination, that is, about all the sentences of the language. How to proceed?

What's meant by saying that what we say about the words is to "deliver information" about whole sentences? We need to be able to *derive* that information, by an inferential process that at least in principle can be made fully explicit in a system of logical rules. This suggests that what we need is a "theory", in a somewhat technical use of that word: a system of sentences in which some are axioms, and the remainder follow from the axioms by specified rules of logic. As philosophers put it, a theory, in this sense (and it's how I'll be using the word "theory" in the rest of this chapter), is a set of sentences *closed* under the logical rules: anything which follows by the rules from any sentences in the set is also in the set. (For example, if some proposition *p* is in the set, and so is *q*, then the

proposition “*p* and *q*” is also in the set, since it follows from other things in the set.) This answers a structural question: meaning needs to be described by a theory. But how?

Let’s start with very simple sentences of English, like “Fido barks”. We assume that “Fido” refers to some one dog (and we’ll forget that in fact “Fido” is used for many dogs, and no doubt for many other things as well). We want to be able to say something about “Fido” and “barks” which will enable us somehow to describe the meaning of the whole sentence “Fido barks”. Let’s try this: “Fido” means (or stands for) Fido; “barks” means (or stands for) barking, or being a barker; and (here’s the part that addresses the way words are put together) when a proper name (like “Fido”) precedes an intransitive verb (like “barks”) the resulting sentence means that what the name means has the property the verb means.

A problem is that this is false. “Fido barks” means that Fido barks; it does not mean that what “Fido” means has the property “barks” means. It’s not about words and their meanings, it’s about a dog and his barking. Can we attain a truth by taking a further step? Suppose Fido *is* the meaning of “Fido”, and that barking *is* the meaning of “barks”. Then we should be able to replace “what ‘Fido’ means” by “Fido” and “what ‘barks’ means” by “barking”. Then we’d get: “Fido barks” means that Fido has the property of barking.

This is a whole lot better, but still not quite right, since having the property of barking is one thing and barking is another, seemingly simpler, thing. The very young, for example, could use “Fido barks” with the meaning that Fido barks, but, lacking the concept of a property, arguably could not use it with the meaning that Fido has the property of barking. Here’s one way round the difficulty. Instead of describing the meaning of “barks” by associating it with a property, let’s instead associate it with the things of which the word is true: “barks” is true just of things that bark. The axiom for putting names and verbs together could be modified like this: a sentence that results from prefixing an intransitive verb by a proper name means that the meaning of the name is among the things of which the verb is true. Again the result is not what’s wanted: “Fido barks” does not mean that Fido is among the things of which “barks” is true. We would like to be able to replace “being among the things of which ‘barks’ is true” by “barks”; but this replacement can’t be justified in the context: it’s not a replacement which preserves meaning, as is made apparent by the fact that before the replacement we have the wrong meaning, and after it the right one.

These difficulties can be resolved if, instead of using *meaning* as the central notion by which to describe linguistic meaning, we instead use *truth conditions*. Our target would then be to establish, on the basis of things said about “Fido” and “barks”, that “Fido barks” is true if, and only if, Fido barks. We can do this if we say that “Fido” stands for Fido, that “barks” is true just of things that bark, and that a sentence consisting of a name plus intransitive verb is true if, and only if, the referent of the name is among the things of which the verb is true. So “Fido barks” is true if, and only if, the referent of “Fido” is among the things of which “barks” is true, and now we can go on to infer that this in turn is so if, and only if, Fido barks. In a theory of truth conditions, we don’t have to preserve meaning in making inferences, only truth conditions, and the relevant inference (to *barking*, from being among the things of which “barks” is true) does this.

Somewhat technical considerations about how to devise a theory of meaning coincide with more philosophical considerations to suggest that truth conditions should play a central role in an account of meaning. We need to revisit an earlier worry: don’t any two necessary truths have the same truth conditions? So wouldn’t a truth conditions theory of meaning be committed to the absurdity that any two necessary truths have the same meaning? In response, Davidsonian theorists have to say that they’re not identifying meaning and truth conditions, if truth conditions are thought of in such a coarse-grained way that every necessary truth has the same truth conditions. Rather, a *theory* of truth conditions, taken as a whole, can “serve as” a theory of meaning. A truth condition, from this perspective, is what follows the “if, and only if,” in a specification of truth conditions. A specification for “Every triangle has three sides” will be that it’s true if, and only if, every triangle has three sides. A specification for “Nothing is red and green all over” will be different: it’s true if, and only if, nothing is red and green all over. Just as the sentences are differently composed, their truth conditions will be derived by different routes, and so will count as distinct; this is what will mark the difference in meaning of the sentences for which they give the truth conditions.

Here’s a further qualm: does “‘Fido barks’ is true if, and only if, Fido barks” really tell us anything about meaning? Or, indeed, anything at all that we want to know? Isn’t it just trivial? We didn’t have to derive this from axioms relating to words; we could just have given a rule: take any sentence, write it within quotation marks, append “is true if, and only if,” and then write the sentence again. Truth conditions theorists, like Davidson, are at pains to repudiate

this claim of triviality. They point out that a monolingual Frenchman doesn't know that "Fido barks" is true if, and only if, Fido barks, and that's because he doesn't know English. The specification of truth conditions links the sentence to a state of affairs, Fido's barking, and thus describes the crucial connection between language and world in which meaning consists.

This crucial connection cannot be held in place by magic: it must somehow be effected by what speakers of a language do. We saw that if we conceive this behavior in terms of movement of matter, skepticism looms, whether of the Quinean or the Kripkensteinian variety. Suppose, however, we allow ourselves to describe behavior in terms of the intentions of speakers. Can we then give a satisfying account of the connection between language and the world, based on the way in which intentions relate to the world? That question will be addressed in the next section.

Meaning and intentions

linguistic intentions are very like nonlinguistic intentions.
(Paul Grice, "Meaning")

In apparent opposition to truth conditions theorists of meaning, other philosophers, the most famous being Paul Grice, have urged that we should account for meaning in terms of the intentions of those who speak the language. Grice took as the first notion to analyze that of someone's meaning that such-and-such by an action, where the action may or may not involve language. His ultimate ambitious plan was to move on to describe the meanings of words in terms of conventions governing what speakers should use them to mean. To this ambitious part of the plan, he and others have paid relatively little attention; but the early part, in which Grice offers an analysis of speaker-meaning, has been highly influential.

Herod showed Salome the severed head of St John the Baptist, trying to make manifest to her that he had kept his promise and that John the Baptist was dead. Herod tried, and succeeded, in getting Salome to form various beliefs, but intuitively, in showing her the head, he did not *mean that* John was dead, or that he had kept his promise. He engaged in intentional behavior, and behavior designed to induce beliefs in others, but he did not mean anything. Why not?

Grice puts his finger on a crucial issue in discussing two superficially similar cases. In one of them, I show Mr. X a photograph of his wife displaying undue familiarity with Mr. Y. In the other, I do a sketch which is supposed to represent the very same thing. Grice suggests that in the case of the sketch we are inclined to say that I meant that Mrs. Y was being unfaithful; but we are not inclined to say this in the photograph case. Grice explains the contrast as follows: in the case of the photograph, I can count on Mr. X to form the belief that his wife was unfaithful quite independently of his having any opinion about what I was up to. (Project back to the days before Photoshop, in which it was assumed that the camera could not lie!) In the case of the sketch, however, I would know that Mr. X would only form the belief that his wife was being unfaithful if he recognized my intention to get him to form this belief. I would know that if he thought I was just doodling, or fantasizing with pencil in hand, seeing the sketch would not lead him to form that belief.

The example brings out the central element of Grice's theory: speaker-meaning involves what has come to be called the "Gricean mechanism": the speaker (in the broadest sense, to include one who holds out a photo or a sketch) must intend his audience to recognize his intention that she should form a certain belief, and this recognition should be intended to supply a reason for the audience to form that belief. In the case of the sketch, one could reconstruct how I intend Mr. X to reason along the following lines: "he has shown me a sketch apparently representing my wife displaying undue familiarity with Mr. Y. It would be in terrible taste to do this in my presence merely as a doodle or distinctly off-color fantasy, so he must have had some serious purpose. He must believe my wife is behaving in the way he has sketched. If he believes that then, since he is reliably informed, it is probably true. So, probably, my wife is being unfaithful." By contrast, in the case of the photograph, I can't expect Mr. X to reason in this way. Seeing the photograph is likely to produce the belief quite independently of any beliefs Mr. X may form about my intentions. That's why I mean something by showing the sketch but don't mean something by showing the photo. And that's why Herod didn't mean that John the Baptist was dead by showing Salome the severed head: he would have realized that Salome, seeing the head, would believe he was dead quite independently of any appreciation of Herod's intentions.

Grice's idea has been subject to all kinds of criticism, but there is a wide consensus that he has put his finger on

something interesting and important about how communication works, in particular in those cases in which communication is effected by linguistic means. What is less widely accepted is that the Gricean mechanism can play the role Grice cast it for in an analysis of speaker-meaning; some even doubt the robustness of any such notion.

One battle Griceans have had to fight is against truth conditions theorists. It has sometimes seemed to both parties that their views are mutually opposed: if meaning is truth conditions, then how can it be based on the notion of speaker's intentions and the Gricean mechanism? Maybe this sense of conflict is largely illusory. Truth conditions theorists do specially good justice to the compositional character of language. Griceans do specially good justice to the evident truth that language in some way depends upon use, and use is a species of intentional behavior. But a complete theory must do justice to both aspects: not just to the connection with use, but also to compositionality. Since countless sentences which have never been used are meaningful, Griceans have found it necessary to engage in some detours to incorporate the compositional aspects of language into their story. Some truth conditions theorists forge the link with use when they describe the kind of evidence that would ground a correct truth-theoretic account of a language: the evidence is, of course, how the language is used, and for Davidson the specific aspect of use is a tally of which of the sentences they utter the speakers hold true.

Each party must borrow material from the other: Griceans will somehow need to exploit the theoretical apparatus of truth theorists, or something like it; truth-theorists cannot entirely neglect the use of language and the intentions which animate that use. But one should not exaggerate the degree of rapprochement: truth-theorists can take cognizance of speakers' intentions without necessarily placing any weight on the Gricean mechanism; and although Griceans need to harness a theory, if they are going to provide a description of the meaning of every sentence of a language, it might be that they will regard this as a dispensable ambition: perhaps they are not offering such a description, but are, rather, providing an insight into what lies at the heart of meaning in any possible language.

Meaning and use: semantics and pragmatics

So far in this chapter, linguistic meaning has been at the center of our discussion. Admittedly, the Gricean mechanism relates to speaker-meaning and not to sentence-meaning; but we introduced it in the context of the ambition to describe sentence-meaning. Grice himself was vividly aware that the impact of our use of language often diverges markedly from the meanings of the sentences we use. A simple example is irony: we say the opposite of what we mean. Saying goes with sentence-meaning, and what we mean is, of course, speaker-meaning, and the two diverge in the most obvious way: one is the opposite of the other.

Here is a more interesting example: I ask you whether you want coffee, and you reply "It will keep me awake". Does this amount to yes, or no, or neither? That depends on the context. Suppose I know you need to stay awake in order to finish an essay. Then I'll take you to have meant Yes. Suppose I know you want to have an early night to prepare for the next day. Then I'll take you to have meant No. There is no connection, merely at the level of language, between either "Yes" or "No" and "It will keep me awake". Yet we could hardly give a full account of the role of language in communication unless we observed the way in which such a connection can be forged.

As Grice's work made clear, a great deal of what we gather from the speech of others is sensitive to context in this kind of way. If you say "I'm nearly out of gas" and I reply "There's a gas station around the corner", you will assume that the gas station is open, or at least that I do not know it is closed, though what I strictly said did not mention whether it was open or not. If you ask me for a reference for a student and I write "She has beautiful handwriting" and nothing else, you will take it that I have expressed a negative opinion of her, though I have strictly speaking said nothing negative. If we use "semantic" correlatively with "sentence-meaning" and "pragmatic" correlatively with "speaker-meaning", the phenomenon we have started to uncover is that communication often involves an extra-semantic, pragmatic component; and sometimes this is more important than the semantic component.

As we reflect on these cases, we may begin to wonder whether any linguistic exchange can fail to exploit pragmatic features. In the ancient philosophy seminar, you say "Aristotle's logic is quite sophisticated". How do I know you are talking about Aristotle the philosopher, rather than about Aristotle Onassis, the shipping magnate? The words on offer do not themselves settle the issue. The issue is settled by the context: we are in the ancient philosophy seminar, and the ascription to Mr. Onassis of such a property would be surprising. In this case, features of context help guide the audience to the right Aristotle. One might think that there are two words "Aristotle" spelled and

pronounced the same, and with different semantics (one refers to the philosopher, the other to the tycoon). On this view, the features of context do not fill in semantic properties, but only help determine which semantic properties are present.

In other cases, the contextual information does more. Consider this sentence:

<dsp>The policeman shouted to the robber to stop. He was in full uniform.

Who was said to be in uniform? Most people's first response is that it was the policeman. On reflection, they realize that it could have been the robber, too, but that's a less likely interpretation in most (not all) contexts. (Think of a context in which we have already been told that the robber was an army colonel.) Here again context helps resolve a semantic issue.

In the case of coffee, irony, or beautiful handwriting, the pragmatic contribution goes beyond anything that could be thought of as semantic. For some cases, it's hard to decide whether that's the best description or whether, as with Aristotle and the policeman, contextual information is simply rendering the semantic properties more determinate. Compare the following distinct conversations between me, in Brownsville, and you, in Austin:

- A: You: Is it raining?
Me: Is it raining? You know it never stops here.
B: Me: It is raining?
You: No, clear skies as usual.

In A, my question "Is it raining" asks whether it is raining in Brownsville. In B, my question, asked with the same words, and with the speakers in the same places, relates to whether it is raining in Austin. On one view, "It's raining" has a semantic slot, waiting to be determined contextually. On another view, nothing in the semantics of rain determines a place; that information is a purely pragmatic addition.

These kinds of case raise a more general skepticism about whether sentence-meaning, the purely semantic contribution, ever determines truth conditions. (If it does not, truth conditions theorists are in trouble.) Try to think of a sentence whose meaning is completely independent of context. Charles Travis suggested what seems like a good candidate: "The ball is round", as spoken of a particular ball at a particular moment. At any moment, a ball has just one shape, round or not, so it seems that the sentence must have a single pragmatically independent meaning. Travis, however, thought that even this sentence does not. He illustrated the point by describing a squash game. You know nothing about squash and although you're watching a game your eyesight isn't very good and you ask whether the ball is round, like a cricket ball, or ovoid, like an (American) football. I say "The ball is round", and what I say is (surely!) true. Another spectator is a manufacturer of squash balls who has fitted this particular ball with very small sensors to determine its response to deformation (in the interests of product improvement). When the ball is deformed, the sensors should trigger a reading on the spectator's measuring device. Checking his equipment, he asks the same question, and I reply as before, "The ball is round". But I'm not paying attention: at that moment the ball was flattened against the wall and so not round. The very same sentence is false, on the understanding appropriate to the second question, though true on the understanding appropriate to the first. So perhaps pragmatic phenomena always play a role in settling how we are to understand the sentences we hear or read. Perhaps every sentence can be understood in more than one way, and semantic meaning, understood as rich enough at least to determine truth conditions, is an illusion.

Conclusion

Semantic meaning, the meaning of words and sentences, is determined by what people do, understood not as the motion of matter through space but as intentional action. Our communicative actions have layers of meaning that are much more complex than the meanings of the words we use. Quite how to draw the line between the pragmatic and the semantic aspects of meaning is not clear. But it does seem clear that linguistic meaning cannot be sharply separated from a wider notion of meaning that's applicable to everything we do. Understanding what someone says cannot be sharply separated from understanding the speaker. In John McDowell's words: "We have not properly made sense of forms of words in a language if we have not, thereby, got some way towards making sense of its speakers. If there is a pun here, it is an illuminating one" (1977: 160).

References

Davidson, D. (1967). "Truth and meaning." *Synthese* 7: 304–323.

Grice, H. P. (1957). "Meaning." *Philosophical Review* 66: 377–88.
Kripke, S. (1982). *Wittgenstein on Rules and Private Language*. Harvard University Press.
McDowell, J. (1977). "On the sense and reference of a proper name." *Mind* 86: 159–85.
Quine, W. V. O. (1960). *Word and Object*. New York, Technology Press of MIT and John Wiley and Sons Inc.
Wittgenstein, L. (1953). *Philosophical Investigations*. Oxford, Basil Blackwell.

Suggested reading

For a highly entertaining approach to language, focussing on empirical questions, see:

Pinker, Steven (1994). *The Language Instinct: How the Mind Creates Language*. Penguin.

To explore the questions raised in this chapter, I'd recommend visiting the Stanford Encyclopaedia of Philosophy (plato.stanford.edu). You will find it rewarding to sample at least these search strings: "gavagai", "Kripke's sceptical Wittgenstein", "Davidson meaning and truth", "pragmatics". For a book intruding the main themes, as these are understood in the analytic tradition, I would suggest:

Lycan, William (1999). *Philosophy of Language: A Contemporary Introduction*. Routledge.