

Foundational Issues in the Philosophy of Language

Linguistic expressions are meaningful. Sentences, built from words and phrases, are used to communicate information about objects, properties and events in the world. In philosophy of language, the study of linguistic meaning is central.

1. Meaning and communication: Semantics and pragmatics

Just connecting meaning with communication does not yet tell us which properties of expressions and their uses are to be brought under the umbrella of linguistic meaning. For there is often an intuitive distinction to be drawn between the message that a speaker communicates and the meaning of the sentence that the speaker uses. Offered coffee after dinner, Nigel may utter the sentence, ‘Coffee keeps me awake’, and thereby decline the offer. But declining an offer is no part of the linguistic meaning of the sentence. In a different context – imagine that there is an after-dinner speaker who is known to be boring but that the occasion demands that people not doze off – the very same sentence could be used to accept an offer of coffee. In yet another context, the sentence might be used in a simple factual report about factors that exacerbate insomnia.

It is usual, in philosophy of language, to distinguish between properties that are strictly aspects of an expression’s literal linguistic meaning and other properties of the expression or its use that may contribute to the message that is communicated. The study of literal linguistic meaning is *semantics*. The study of the use of language to communicate messages – very often, messages that go far beyond the literal meanings of the expressions used – is *pragmatics*.

As a rough guide to the semantics-pragmatics distinction, we can say that semantics is concerned with the properties of expressions that help to determine the conditions under which an utterance would be literally true, rather than false. In contrast, pragmatics focuses on the conditions under which an utterance would be helpful, rather than misleading, or more generally appropriate, rather than inappropriate. The semantic properties of the English words ‘coffee’ and so on, and the way that the words are put together, determine that an utterance of the sentence ‘Coffee keeps me awake’ is true under the conditions that coffee keeps the speaker awake, and false otherwise. So if coffee keeps Nigel awake then his utterance is true. But an utterance that is true may be misleading. Late in the evening, with no after-dinner speaker scheduled, Nigel’s utterance may give the impression that he is declining the offer of coffee, so that his next remark, ‘I’d like some’, comes as a surprise. The fact that the utterance has the truth conditions that it does is explained by semantics. The fact that the utterance gives the misleading impression that it does is explained by pragmatics.

However, it is important to recognise that equating semantics with the study of literal linguistic meaning does not quite guarantee that semantics is concerned with truth conditions while pragmatics deals with appropriateness conditions. One kind of example that helps to make this point is provided by pairs of expressions that apply to the same things but differ in literal meaning because one word is a polite or neutral form while the other is an impolite or derogatory form. Thus, for example, the English word ‘cur’ applies to the same things as the word ‘dog’; but, while ‘dog’ is the neutral form, it is part of the literal meaning of ‘cur’ that its use expresses contempt. Suppose that Rover is a dog of mixed, or even indeterminate, breed. Some people might feel contempt towards Rover, but suppose that Fiona’s attitude is one of admiration and affection. Then an utterance by Fiona, pointing at Rover, of ‘This cur slept all night’ would be seriously misleading as to her attitude. But if Rover did sleep all night then Fiona’s utterance

would not be false. It would be true but misleading. In that respect, Fiona's utterance would be similar to Nigel's. But there is also an important difference between the cases. Someone who hears Nigel's utterance and grasps the literal meaning of what he says may very well fail to draw the misleading inference that he is declining the offer of coffee. The hearer might know that Nigel wants to stay awake to write a philosophy essay, for example. But someone who has no tendency to draw the misleading inference that Fiona is contemptuous of Rover has failed to grasp the literal meaning of the word 'cur'.

It appears, then, that two words may differ in their literal meaning even though they make the same contribution to truth conditions. Here we can make use of a distinction drawn by Frege (1892) between *sense*, on the one hand, and illumination, colouring or *tone*, on the other (Dummett, 1973, chapter 1). Many complex issues surround Frege's notion of sense. But for present purposes, the only aspects of meaning that belong to the sense of an expression are those that help to determine the truth conditions of sentences in which the expression occurs. Frege's notion of tone is more heterogeneous, including even the various ideas that may be evoked in individual readers by the eloquence of a poet. But Frege also mentions that conjunctions such as 'but', 'although', and 'yet' *illuminate* the sense of the following clause 'in a peculiar fashion'. This is the point, familiar to every student of logic, that 'but' differs from 'and' in some aspect of meaning but not in the contribution that it makes to truth conditions. The use of 'but' in 'Bruce is Australian but he is cultured' carries the implication that there is some kind of contrast between being Australian and being cultured. If there is, in reality, no such contrast, then this is a misleading thing to say. But that does not make it false. The truth value of the sentence with 'but' is the same as the truth value of the more neutral 'Bruce is Australian and he is cultured'. In Fregean terminology, 'but' differs from 'and', not in sense, but in tone.

We have distinguished between two distinctions. There is the distinction between literal linguistic meaning and communicated messages or implications that go beyond literal meaning. And there is the distinction between aspects of meaning that contribute to the determination of truth conditions and factors that generate implications that are not directly relevant to questions of truth or falsity. We have introduced the terms 'semantics' and 'pragmatics' to mark the first distinction. But we should note that some theorists might prefer to use these terms for the second distinction, restricting semantics to the study of sense rather than tone, and correspondingly expanding the domain of pragmatics. These theorists must, of course, recognise a difference, within the domain of pragmatics, between implications that are generated by aspects of literal meaning (as in Fiona's utterance) and those that result from some kind of interaction between literal meaning and contextual factors (as in Nigel's). Whether or not anything of substance turns on this terminological decision, it is important to be clear about the difference between the two distinctions. In some writings on philosophy of language the difference is apt to be obscured because the phenomenon of tone is set aside, virtually without comment.

2. Meaning, science, and philosophy: Semantics and meta-semantics

We have taken some time to describe the domain of semantics. But it may seem that, whatever the details of the demarcation, there is a puzzle as to how semantics could be of central interest for philosophy of language. Facts about what expressions a language contains, how those expressions are used in utterances, and what those expressions refer to or mean are empirical facts about the natural order. It is an empirical fact about the language spoken in Italy that there is a word 'prosciutto', pronounced with the stress on 'u' and with 'sci' sounding much like English 'sh', that refers to a particular kind of ham.

It is an empirical fact that this Italian word has been incorporated into our language just as the ham has been incorporated into our diet. It is an empirical fact that the Italian word can be used in the sentence ‘Prosciutto è buono con melon’ to mean that this kind of ham is good with melon. These facts are surely of limited philosophical interest. It is difficult to regard them as belonging alongside facts about the nature of free agency or about the conditions for a conception of objects as existing independently of our perception of them, for example. So, how could the study of meaning and reference be a distinctively philosophical project rather than simply a part of natural science?

To this question, put in this way, an answer immediately suggests itself. If there is a science of meaning and reference then, presumably, the philosophy of science includes a division that offers philosophical commentary on the key concepts deployed by that science. So philosophy of language can concern itself with meaning and reference just as philosophy of biology concerns itself with species and selection. Now, in fact, there are reasons to hesitate over this assimilation of the philosophy of language to philosophy of science. There is indeed a science of language, namely linguistics, and philosophy of linguistics is a division of philosophy of science. But we are not obliged to hold that philosophy of linguistics is all there is to philosophy of language. For philosophy of language also includes questions that arise from our ordinary, everyday conceptions of language and meaning. (In a similar way, we should not hold that philosophy of mind is exhausted by philosophy of cognitive science.) However, we do not need to embrace a gratuitous scientism in order to respond to the apparent puzzle about how semantics could be of importance to philosophy.

There is a familiar distinction between questions of first-order ethics (‘How should I behave?’) and questions of meta-ethics (‘What is the nature of ethical judgements?’). Similarly, there is a distinction between aesthetic questions about works of art (‘Is this a beautiful painting?’) and questions in aesthetics (‘What is the nature of beauty?’). So also there is a distinction to be drawn between questions about the meanings of particular linguistic expressions and questions about the nature of linguistic meaning itself. Questions of the first kind belong to *semantics*; questions of the second kind belong to *meta-semantics*.

There are three general points to be made about this kind of distinction between ground-level questions and meta-level questions. First, it is the meta-level questions that are most clearly philosophical in character. Questions about the nature of ethical judgements or about the nature of beauty do belong with questions about free will or about the conditions for a non-solipsistic consciousness. Second, the philosophical character of these meta-level questions does not depend on the scientific credentials of the corresponding ground-level questions. There is no science of how one should behave and there are no scientific laws about beauty as such. Third, while it is the meta-level questions that are most clearly philosophical, ground-level questions may also be philosophical questions. Moral philosophy includes, not only meta-ethics, but also first-order or applied ethics and arguments about the aesthetic properties of particular works of art or music are advanced within philosophical aesthetics.

In the case of semantics and philosophy of language, the three points are these. First, meta-semantic questions about the nature of linguistic meaning are central for philosophy of language. Second, philosophy of language need not start from the assumption that ground-level questions about the semantic properties of expressions are scientific questions. Indeed, whether semantics is a science is a meta-level question that is debated in contemporary philosophy of language. Third, it is not to be ruled out that engagement with ground-level questions about the semantic properties of expressions may also be a

philosophical project. In the remainder of this section, we shall briefly develop the first two of these points. The third point has a section of its own.

2.1 Meta-semantic questions: Metaphysics and epistemology

Meta-semantic questions about the nature of linguistic meaning include metaphysical questions. For example, are meanings entities; and, if so, what kinds of entities are they? Intuitively, utterances of sentences are used to express thoughts, and so the meaning of a complete sentence seems to be something like the content of a complete thought. This idea might lead to the proposal that the meanings of complete sentences are *propositions*, and that proposal would lead, in turn, to questions about the nature of propositions themselves. There may be other proposals about what kinds of entities meanings are, assuming that they are entities at all. But an alternative view would have it that we do not need to include meanings as entities in our ontology. Earlier, we gave an example of an Italian sentence with a particular meaning property. The sentence ‘Prosciutto è buono con melon’ means that prosciutto is good with melon. If meanings are entities then this sentence stands in the meaning relation to such an entity, perhaps the proposition that prosciutto is good with melon. But, on the alternative view, the meaning property of the Italian sentence is not to be conceived as a relation to an abstract entity, though it may involve relations to worldly entities like ham and melons (Devitt, 1996). Consider an analogy. Suppose that Nigel has a sharp hairstyle. On a ‘hairstyles as entities’ view, this is a matter of Nigel standing in a particular relation to a particular abstract entity, a hairstyle. But it might be said, with some plausibility, that we do not need to include hairstyles in our ontology. Having a sharp hairstyle is not really a relational property; it is just a matter of having one’s hair styled in a particular way. Similarly in the case of meaning, it might be said that for a sentence to mean that prosciutto is good with melon is just a matter of the sentence being used in a particular way.

Metaphysical questions always go hand-in-hand with epistemological questions. So, in this case, any metaphysical view as to the nature of facts about meaning must be compatible with an answer to the epistemological question how knowledge about meanings is achieved. Competent speakers of a language know, or can come to know, the meanings of indefinitely many sentences. They *understand* the sentences of their own language. Philosophy of language addresses questions about the nature of this understanding and, indeed, the question how understanding a language – with potentially infinitely many sentences – is even possible for finite creatures like us.

2.2 Is semantics a science?

Questions as to whether semantics is, or can be, a science arise from at least two sources (Chomsky, 2000). One concerns our ordinary notion of a language; the other concerns the role in semantics of worldly objects, properties and events.

Our ordinary notion of language is of a shared language, a public language, a language that is used for communication amongst several people, typically amongst thousands or millions of people. Our ordinary notion of linguistic meaning is thus meaning in a public language and semantics is conceived, at least initially, as the study of the meanings of expressions in a public language such as English or Italian. But there are reasons to doubt whether the meaning properties of a public language as such constitute a proper domain for investigation by the natural sciences. One such reason is that there may be, within what we describe as a single language, considerable dialect variation. But whether dialects are considered as variations of a single language or as languages in their own right may owe nothing to the natural order and everything to social, political or even

military factors. So the notion of a public language is the notion of an artefactual, rather than a natural, kind and the same goes for public language meaning.

A further reason to doubt the scientific credentials of public language meaning arises from the idea that the meaning of an expression is closely linked with the use of that expression in inter-personal communication. For in real-life communication, the hearer draws on information about all kinds of things in order to discern the message that the speaker is trying to communicate. And the speaker has expectations about the information available to the hearer, expectations that are themselves grounded in information about a variety of topics. It is not clear that this essential background to communication, a heterogeneous collection of pieces of information, could fit into a scientific theory. So, if meaning really is closely linked with communication, this casts doubt on the prospect of a scientific theory of public language meaning.

In order to avoid the worry that a public language is a socio-political construction rather than a natural object, we could focus on the language of an individual speaker – an idiolect, the limiting case of a dialect. However, if we do this then it becomes even clearer that inter-personal communication depends on all manner of information, assumptions, hunches, wisdom, and wit (Davidson, 1986). For speaker and hearer must be ready to make adjustments for differences in the meaning properties of their idiolects. So a science of literal idiolectal meaning would need to abstract away from the use of language in inter-personal communication. We would need to conceive of communication as a massive interaction between literal meaning and other factors, perhaps not itself susceptible of scientific description.

Suppose that idiolectal meaning could be separated from communication in this way, so that questions from the first source were set aside. There would remain questions from the second source concerning the role in semantics of worldly objects, properties and events. To see how the worry arises, suppose that we set out to describe the meaning properties of the idiolect of Maria, an Italian speaker. Her idiolect may, of course, be idiosyncratic in various ways. Perhaps she uses the word ‘melon’ to refer only to watermelon and not to other kinds of melon. But suppose that, in her idiolect, the word ‘prosciutto’ refers to that particular kind of ham, also called ‘prosciutto’ in English. Then our account (in English, of course) of the semantics of her idiolect will include the principle:

The word ‘prosciutto’ refers to prosciutto (a particular kind of ham).

If this is a scientific principle then the various notions that figure in it should be scientific notions; that is, notions of natural kinds or categories. Given the existence of the science of linguistics, the notion of a word, and of the particular word ‘prosciutto’, is not especially problematic. We might hesitate rather more over the scientific credentials of the notion of reference. But the most immediate problem is that there is no guarantee that there is or will be a natural science of prosciutto as such (cf. Fodor, 1980). And the problem becomes even more vivid when we consider that, in Maria’s idiolect, the word ‘Roma’ refers to Rome, ‘la giustizia’ refers to justice, and ‘la bellezza’ refers to beauty. For there is no natural science of cities, nor of justice or beauty.

Just indicating this worry about the scientific status of semantics does not, of course, settle the issue. But imagine for a moment that it were to be shown compellingly that semantics is not, and cannot be, a science. What would follow from that? If the domain of science were to be equated with the domain of rational enquiry concerning the natural world then it would seem to follow that doing semantics is not engaging in such rational enquiry. Perhaps, on this view, it is more like reading a novel or undergoing therapy. But we are not bound to accept that consequence because the equation really does constitute a

gratuitous scientism. Whether or not semantics is science, the aim of semantics is to give a systematic description, with some theoretical depth, of the meaning properties of some objects and events in the natural world, including utterances, idiolects, and even public languages conceived as social objects (Wiggins, 1997).

3. Semantics as a philosophical project

Meta-semantic questions about the metaphysics and epistemology of meaning are clearly philosophical in character. In contrast, semantic questions about the meanings of expressions in a public language like Italian or an idiolect like Maria's seem to be straightforwardly empirical and of limited philosophical interest. So it may seem puzzling that philosophy of language should include substantial contributions to semantics as well as to meta-semantics. How could the construction of semantic descriptions of languages or idiolects be a philosophical project?

One reason for philosophical interest in semantics is that spelling out the conditions for a putative semantic description to be *correct* is a way of approaching meta-semantic questions of undoubted philosophical importance. Any elucidation of the concept of literal linguistic meaning furnishes a correctness condition on semantic descriptions. And conversely, correctness conditions on semantic descriptions can be transposed into elucidations of, or constraints on, the concept of meaning. A philosopher of language who adopts this approach to meta-semantic questions obviously needs to know something about the semantic descriptions for which correctness conditions are to be formulated.

However, the possibility of this approach to meta-semantic questions does not really explain why detailed and substantive engagement with the construction of semantic descriptions could be a philosophical project. For the formulation of correctness conditions on semantic descriptions requires only very general knowledge about the form that the descriptions are to take. In order to go further, we must take account of the *compositionality* of linguistic meaning.

3.1 Compositionality and finitely axiomatised semantic theories

The meaning of a sentence is determined by the meanings of its constituent words and the ways in which they are put together. It is part of our everyday conception of knowing a language that, because of this compositionality, it is possible to understand a completely novel sentence on first hearing if it is built out of familiar words put together in familiar ways. So the compositionality of meaning promises an answer to the epistemological question how understanding a language with potentially infinitely many sentences is so much as possible. The general form of the answer must be that the meanings of finitely many words and finitely many ways of putting words together determine the meanings of infinitely many sentences. But it is not a trivial matter to give substance to this answer.

The most convincing way to develop the idea of the compositionality of meaning is to provide, for the language or idiolect under investigation, not just a semantic description, but a certain kind of formal, axiomatised semantic theory. The axioms of this theory should specify the meaning properties of words and of ways of putting words together. The background logic of the theory should then permit the derivation, from these axioms, of theorems that specify the meaning properties of complete sentences. For any particular sentence, the derivation of a meaning-specifying theorem should draw specifically on the axioms for the words and ways of putting words together that are involved in the construction of that sentence. Furthermore, the derivation should follow a natural and direct kind of route that can be specified by a canonical proof procedure. In a formal, axiomatised semantic theory of this kind, the canonical derivation of a meaning-

specifying theorem for a sentence displays how the meaning of the sentence is determined by the meanings of its constituent words and the ways in which they are put together. Such a semantic theory could itself be called *compositional*.

Suppose that we could provide a compositional semantic theory in which meaning-specifying theorems for the infinitely many sentences of some language could be derived from finitely many axioms specifying the meaning properties of words and of ways of putting words together. This would constitute a relatively precise answer to the question how understanding an infinite language is *possible*. For knowledge of the finitely many facts stated by the axioms of the theory *would suffice*, given an ability to carry out logical derivations, for knowledge of the meaning of any sentence of the language. It is for this reason that philosophers of language engage in the construction of formal semantic theories subject to a *finite axiomatisation constraint* (Davidson, 1984).

3.2 *A structural constraint on semantic theories*

It is not clear, however, that finite axiomatisation really goes to the heart of the notion of compositionality. In a finitely axiomatised semantic theory for an infinite language, a single axiom may be involved in the canonical derivations of meaning-specifying theorems for many sentences. Indeed, in general, an axiom will figure in the derivations for infinitely many sentences. This gives clear content to the idea that a word makes a repeatable contribution to the meanings of sentences in which it occurs. It is because the constituents of sentences make repeatable contributions to the meanings of sentences that it is possible to understand a sentence that one has never heard before provided that it is built from familiar words in familiar ways. The problem with the finite axiomatisation constraint is that these features of the compositionality of meaning may be present in a finite language, no less than in an infinite language.

Imagine a toy language with one hundred sentences built from ten names ('a', 'b', 'c', ...) and ten predicate terms ('F', 'G', 'H', ...). A particular name makes a repeatable contribution to the meanings of the ten sentences in which it occurs; and the same goes for a particular predicate term. Someone who is familiar with the name 'a' as it is used in several sentences (to talk about Harry) and with the predicate 'F' as it is used in several sentences (to mean that various people are bald) may understand the sentence 'Fa' on first hearing as meaning that Harry is bald. But the finite axiomatisation constraint does not require that a semantic theory should display the way in which the meaning of 'Fa' is determined by the meanings of 'F' and of 'a' and the way in which they are put together. It does not require that a semantic theory should be compositional. For a semantic description that simply lists the meanings of all the sentences:

'Fa' means that Harry is bald

'Fb' means that Bruce is bald

and so on, counts as a finitely axiomatised semantic theory. The one hundred axioms simply coincide with the one hundred meaning-specifying theorems for complete sentences.

It appears that we need some further constraint on formal, axiomatised semantic theories if they are to explain how it is possible to proceed from understanding of some sentences (of even a finite language) to understanding of other sentences built from the same constituents. In the case of the toy language with just one hundred sentences, the compositionality of meaning would allow someone who understood several sentences containing 'a' and several sentences containing 'F' to proceed, by intuitively rational means of inductive, abductive, and deductive inference, to knowledge of the meaning of 'Fa'. What we want of a semantic theory is that, in such a case, the axioms (and other

deductive resources) that are used in the derivations of meaning-specifying theorems for the initially understood sentences should suffice for the derivation of a meaning-specifying theorem for the consequently understood sentence. A theory meeting this condition would display how a name, such as ‘a’, and a predicate, such as ‘F’, make repeatable contributions to the meanings of sentences in which they occur. Given such a theory, we could see how someone might proceed, by rational inductive or abductive means, from knowledge about the meanings of several sentences containing ‘a’ and several sentences containing ‘F’ to knowledge of the facts stated by the axioms that describe the meaning properties of ‘a’ and ‘F’. And we could see how someone with knowledge of those facts might proceed by deductive means to knowledge of the meaning of the sentence ‘Fa’.

Considerations of this kind motivate a *structural constraint* on semantic theories that goes well beyond the finite axiomatisation constraint. The idea of this constraint is that a formal, axiomatised semantic theory for a language should display the compositional semantic structure that is present in the language, and should display it as structure that could be used by an idealised rational subject. A semantic theory meeting this constraint shows how systematic mastery of a language, a mastery marked by the ability to move from understanding of some sentences to understanding of others – and, in principle, of infinitely many others – could be a rational achievement (Davies, 1981, chapter 3; Wright, 1987).

3.3 Theories of meaning from the armchair

The construction of compositional semantic theories is thus a philosophical project. It makes a contribution to answering the epistemological question how knowledge of language – in particular, knowledge of meaning – is possible.

Devising an axiomatisation that is finite and that satisfies the structural constraint is a fairly technical task for which largely a priori methods, similar to those of logic and mathematics, are appropriate. But it may seem that there must still be a substantial empirical, rather than philosophical, component in the project since it is an empirical question whether the meaning-specifying theorems that are derived from the axioms are correct. It is, for example, an empirical fact about Italian that ‘Harry è calvo’ means that Harry is bald. So the construction of an axiomatised semantic theory even for a toy version of Italian seems to be partly an empirical project.

In practice, there are two ways in which this empirical component can be factored out. On the one hand, a semantic theorist might focus on abstract, formal languages for which the meanings of sentences can be stipulated from the armchair. The theorist could address the epistemological question about natural languages indirectly by devising axiomatised semantic theories for progressively more complex formal languages that contain the same kinds of terms and constructions as natural languages do. On the other hand, a semantic theorist might focus directly on progressively more complex fragments of his or her own natural language. It is an empirical fact about English that ‘Harry is bald’ means that Harry is bald. But the theorist can still avoid the need to rise from the armchair and to conduct substantive empirical investigations. This is because the language for which the semantic theory is to be provided is the same as the language in which the theory is to be cast. There is no more accurate way of specifying, in English, the meaning of the English sentence ‘Harry is bald’ than to use that very sentence, so as to provide a *homophonic* meaning-specification. Thus philosophical work in semantics is often concerned with providing axiomatisations from which homophonic meaning-specifying theorems can be derived.

An axiomatised semantic theory in which the theorems specify meanings surely deserves to be called a *theory of meaning*. But there is another use of this term, by analogy with ‘theory of knowledge’ or ‘theory of justice’. In this second use, a theory of meaning is a philosophical theory about the nature of meaning; it is a meta-semantic theory. In this section, we have seen how the study of theories of meaning in the first sense, semantic theories, can contribute to the theory of meaning in the second sense, the meta-semantic theory about the metaphysics and epistemology of meaning.

4. Approaches to questions in philosophy of language

Foundational questions in philosophy of language concern the nature of meaning, understanding, and communication. But approaches to these questions vary along several dimensions. In this section, we shall briefly consider some factors that may condition choices between these various approaches. First, there are views about the proper order of explanatory priority as between philosophy of language and philosophy of mind. Second, there are conceptions of the relationship between the philosophical study and the scientific study of language and mind.

4.1 Orders of priority

When we use language to communicate information about the world we express our thoughts, and these are also about the world. Just as, in philosophy of language, foundational issues concern the notion of meaning, so also, in philosophy of mind, foundational issues concern the notion of *intentionality*, the way in which our thoughts are about, or represent, worldly states of affairs. It is natural to suppose that issues concerning linguistic meaning and the intentionality of thoughts are closely connected. But it is not obvious which should take priority.

We are concerned here with priority in the order of philosophical explanation, analysis or elucidation. Let us say that the notion of X is *analytically prior* to the notion of Y if Y can be philosophically explained, analysed or elucidated in terms of X, while the explanation, analysis or elucidation of X itself does not have to advert to Y. This kind of priority is to be distinguished from several others including ontological priority, where X is *ontologically prior* to Y if X can exist without Y, although Y cannot exist without X. It is, for example, one thing to ask whether thought can exist without language and quite another thing to ask whether thought can be philosophically explained without adverting to language. On the question of the order of analytical priority as between language (linguistic meaning) and mind (the intentionality of thoughts) four views seem to be possible: mind first, language first, inter-dependence, and independence.

According to the *mind-first* option, it is possible to give a philosophical account of the intentionality of thoughts without essentially adverting to language, and the notion of linguistic meaning can then be analysed or elucidated in terms of the thoughts that language is used to express. The mind-first view finds its boldest and most sophisticated development in the work of Paul Grice (1989; see also Schiffer, 1972). One way of denying the mind-first view is to say that the philosophical explanations of language and mind are *inter-dependent*. There is no way of elucidating the notion of linguistic meaning without bringing in the intentionality of thoughts, nor the other way around. This *no-priority* view is characteristic of the work of Donald Davidson (1984). We shall describe the contrasting programmes of Davidson and Grice in the next section.

According to the *language-first* option, an account of the nature of linguistic meaning can be given without bringing in the intentionality of thoughts, and what a person’s thoughts are about can then be philosophically explained in terms of the use of language.

This option finds expression in the writings of Michael Dummett (1993). Indeed, Dummett holds the view, not only that philosophy of language takes priority over philosophy of mind, but also that it is ‘the foundation of the rest of the subject’ (1978, p. 441). A philosophical account of the nature of linguistic meaning in accordance with the language-first view cannot, of course, be overtly mentalistic. So it would be unsurprising to find that, once this view is adopted, attempts to explain meaning in other terms are apt to sound rather behaviouristic.

The fourth possible view is that philosophical explanations of language and of mind are quite *independent* of each other. Language is one thing and thought is another unrelated thing. This second kind of no-priority view is not plausible at all so long as we regard the objects of study in philosophy of language as being communication systems in use. But the view might be adopted by a semantic theorist whose objects of study are purely formal, abstract languages (Katz, 1984).

4.2 *Philosophy and science*

We have been considering possible views of the relationship between philosophy of language and philosophy of mind. Now we turn to the relationship between philosophy of language and mind, on the one hand, and the sciences of language and mind, on the other. For conceptions of this relationship also condition approaches to foundational issues in philosophy of language.

Towards one end of a spectrum of possible views is a position that says that the only questions about language and mind that are susceptible of rational investigation are questions that belong to the sciences of language and mind. According to this position the business of the philosophy of language and mind is simply to hand all the substantive questions over to cognitive science. We might call this first position *cognitive scientism*. (We mentioned and rejected this kind of view at the end of section 2.2.)

Towards the opposite end of the spectrum is a position that maintains that the philosophy of language and mind offers a distinctive methodology for investigating a class of substantive questions about the notions that figure in our everyday thinking about linguistic and mental matters. The occupants of this position say that cognitive science has little or nothing to contribute to the philosopher’s project of plotting the contours of our conceptual scheme. Let us call this second position *philosophical isolationism*.

These opposed conceptions of the inter-disciplinary relationship go naturally with similarly opposed views about our everyday descriptions of ourselves as talking, thinking, acting, feeling, conscious and self-conscious persons. Cognitive scientism goes naturally with a *reductionist* view of these personal-level descriptions. Philosophical isolationism goes naturally with the view that these descriptions, particularly the ones that are of primary interest to philosophy, are *autonomous* in the sense that their correctness is not answerable to empirical discoveries about cognitive structures and processes.

But there are positions that are intermediate between cognitive scientism and philosophical isolationism, and there are views of our everyday personal-level descriptions that are intermediate between reductionism and autonomy. According to one kind of intermediate position, philosophy, with its distinctive methodology, reveals that some personal-level descriptions of great importance for philosophical theory carry a commitment to the existence of particular kinds of cognitive structures and processes. Yet, we cannot fully reconstruct the personal-level descriptions in cognitive scientific terms. In the case of philosophy of language, an occupant of this intermediate position could maintain that our personal-level descriptions of ourselves as having knowledge of

linguistic meaning are neither reducible to, nor independent from, descriptions of the structures and processes investigated by cognitive science.

5. Two programmes in philosophy of language: Davidson and Grice

Grice approaches the meta-semantic question about the nature of linguistic meaning head-on. He aims to provide an analysis of literal linguistic meaning in terms of conventional practices of communicating messages, and to define the notions of conventional practice and communicating a message in terms of familiar mental notions such as belief and intention. Davidson, in contrast, does not attempt to provide an analysis of the concept of meaning in other terms, but approaches questions about the nature of linguistic meaning indirectly, by way of conditions of adequacy on semantic theories. One key feature of Davidson's programme is that the semantic theories to be considered are theories of *truth conditions*. Another is that the conditions of adequacy on these theories are formulated in terms of the use of a semantic theory in an imagined project of overall *interpretation*.

5.1 Davidson's programme: Truth conditions and interpretation

The basic idea of an axiomatised semantic theory is that the theorems of the theory should specify the meaning properties of complete sentences. A theorem might do this by stating that a sentence stands in the meaning relation to some entity:

The meaning of sentence S is entity m.

Alternatively, without any commitment to meanings as entities, a theorem might adopt the format:

Sentence S means that *p*.

Davidson rejects the 'meanings as entities' view; but he also rejects the 'S means that *p*' format.

The problem that he finds with explicit use of the notion of meaning is that the 'means that *p*' construction presents logical difficulties with the result that formal derivations of meaning-specifying theorems would be problematic. Now, in fact, it is a matter of dispute whether insuperable technical obstacles stand in the way of axiomatised semantic theories that adopt the 'S means that *p*' format. But, however that dispute may turn out, the format that Davidson favours, and the format adopted for most philosophical work in semantics, is:

Sentence S is true iff *p*.

The logical properties of this truth-conditional format are certainly well understood. First, 'is true' is a predicate and 'iff' is the material biconditional of propositional and predicate logic. And second, the way in which theorems specifying truth conditions for complete sentences can be derived from axioms assigning semantic properties to words and ways of putting words together can, to a considerable extent, be carried over from the work of Alfred Tarski (1944, 1956) on certain formal languages.

Davidson imposes both formal and empirical conditions of adequacy on truth-conditional semantic theories. The formal conditions of adequacy include the finite axiomatisation constraint (section 3.1). But our concern in this section is with the empirical constraint of *interpretational adequacy*. This says, roughly, that a semantic theory, θ , for a language in use is adequate to the extent that its theorems contribute to the best overall interpretation of speakers – their utterances, thoughts, and behaviour:

If θ delivers (by a canonical proof) a theorem saying that sentence S is true iff p then it should be the case that interpreting utterances of S as expressions of the proposition that p contributes to the best overall interpretation of speakers.

There are three things to notice about this interpretational adequacy constraint on semantic theories. The first is that, strictly speaking, it is not formulated as a correctness condition on the semantic descriptions that are provided by theorems. The constraint does not treat a specification of truth conditions as a surrogate for a specification of meaning. This is just as well since the logical properties of ‘iff’ guarantee that there is a massive gap between:

Sentence S is true iff p

and

Sentence S means that p .

The closest that we come to a surrogate for ‘Sentence S means that p ’ in the Davidsonian framework is, perhaps, ‘A semantic theory, θ , meeting such-and-such formal and empirical constraints, has a canonical theorem stating that S is true iff p ’. (This idea of a near-surrogate for ‘Sentence S means that p ’ within the framework of truth-conditional semantics could be inspired by what Frege says in a famous passage in *Grundgesetze* (1893, I:32) about the sense and truth conditions of sentences.)

The second thing to notice is that the imagined project of overall interpretation involves simultaneously assigning meanings to utterances and contents to thoughts. The aim of the interpretation is to make the best possible sense of the totality of a person’s linguistic and non-linguistic behaviour given the person’s circumstances. The interpretational adequacy constraint can be transposed into a partial elucidation of the concept of meaning for sentences of a public language:

If sentence S means (in the public language L) that p then interpreting utterances of S as expressions of the proposition that p should contribute to the best overall interpretation of L -speakers.

But the same notion of overall interpretation would figure in an elucidation of the intentionality of thoughts. There is no analytical priority of thought over language; nor the other way around. Linguistic meaning and the intentionality of thoughts have to be elucidated together.

The third thing to notice about the empirical constraint of interpretational adequacy is that it provides only a thin elucidation of the concept of meaning. It does not tell us what makes an interpretation of speakers a good interpretation. An interpretation will involve the description of speakers as engaging in certain speech acts (saying and asserting things, for example) and as having certain propositional attitudes (believing and intending things, for example). The constraint requires that these attributions of speech acts and attributions of propositional attitudes should fit together to make best sense of the speakers. But it does not tell us, for example, which combined attributions of speech acts and propositional attitudes make good sense, and which do not.

What seems to be missing from the Davidsonian framework is an account of which combinations of speech act attributions and propositional attitude attributions are coherent. On the face of it, this would be provided by articulating the analytical connections between the concepts of various speech acts like saying and asserting and the concepts of propositional attitudes like believing and intending. There might, for example, be conceptual connections that require that anyone who asserts that p intends an audience to take him or her (the speaker) to believe that p . There might also be

conceptual connections between concepts of speech acts and the notion of literal meaning or some surrogate notion.

5.2 Grice's programme: Analysing the concept of meaning

The bold proposal of Grice's programme is that there are conceptual connections that actually permit the analysis of the concept of literal meaning in a public language (and the concepts of the various speech acts) in terms of propositional attitudes.

If we are aiming at an analysis of the notion:

Sentence S means that p in the language of group G

then a first suggestion might be this:

Members of G use the sentence S to *communicate the message that p* .

This suggestion is in accordance with the mind-first view. There seems to be a reasonable prospect of explaining the notion of communicating (or getting across) a message in terms of propositional attitudes, without having to re-introduce the concept of literal meaning. Certainly, someone can get across a message even though the sounds used have no literal meaning at all. For example, the perpetrator of a sound that has no literal meaning might get across the message that he or she is angry by relying on some resemblance between the sound produced and the sound made by an angry dog.

However, there is a problem with the suggestion. Suppose that, not just one member, but every member, of the group uses this same kind of sound to get across the message that he or she is angry. According to the putative analysis, this should be sufficient for the sound to have the literal meaning that the perpetrator is angry. But, intuitively, it is not sufficient. Indeed, the intuition that this is not a case of literal meaning is strong if each person who uses a sound like an angry dog to communicate anger takes himself or herself to be making an innovative use of a resemblance between sounds (Schiffer, 1972, chapter 5).

When we have a case of literal meaning, in contrast, it seems that the reason why we use a particular sound is just that it does have the appropriate literal meaning. However, if the aim is to provide an analysis of the concept of literal linguistic meaning in other terms then we cannot appeal directly to this reason why the sound is used. For the resulting analysis would be clearly circular:

Members of G use the sentence S to communicate the message that p and do so *because S means that p* .

We can do better by appealing to the idea of a convention as a *rationaly self-perpetuating* regularity in the practices of a group of people (Lewis, 1969). In an analysis of the concept of convention there is a clause saying that the fact of past conformity to the regularity provides members of the group with a reason to conform to the regularity in the future. The clause does not say that the reason why members of the group conform to the regularity is that there is such a convention, for such a clause would make the analysis circular. But, once the analysis has been provided, we can say, harmlessly enough, that members of the group conform to the regularity 'as a matter of convention' or 'because it is a convention'.

We can combine the concept of convention with the concept of communicating a message to provide an analysis of literal linguistic meaning along the following lines:

In the practices of G there is a regularity of using the sentence S to communicate the message that p and this regularity is a convention.

So, when a member of G communicates the message that he or she is angry by using an expression that literally means just that, the speaker does not rely on any resemblance

between the sound produced and the sound of an angry dog. Rather, the speaker relies on the fact, known to both speaker and hearer, that the expression has regularly been used to communicate that message. There is no evident circularity in such an analysis. (Once the analysis has been provided, we can say, harmlessly enough, that members of the group use the expression to communicate that message ‘because there is such a convention’ or, indeed, ‘because that is what the expression literally means’.)

This, in barest outline, is how the Gricean analytical project is supposed to go. But it will be clear, even from this sketchy account, that such an ambitious project faces many challenges. Indeed, one of the most authoritative exponents of Grice’s programme, Stephen Schiffer (1972, 1987), has reached the conclusion that the project of analysing literal meaning in terms of beliefs and intentions cannot be carried through. In the next section, we present one problem for the project, the problem of *meaning without use*.

6. The problem of meaning without use

The general idea that literal meaning is a matter of there being conventions to use expressions in certain ways confronts a dilemma. Either the link between meaning and convention is to be made only at the level of complete sentences or else the link is to be made at the level of words and ways of putting words together.

If we opt for the first alternative, then literal meaning will be grounded in rationally self-perpetuating regularities in the use of complete sentences. The problem is that there are many sentences that are never used at all. These unused sentences, built from familiar words in familiar ways, are perfectly meaningful. But the analysis of literal meaning for sentences in terms of conventional regularities of communicative use has nothing to say about them.

If we shift to the second alternative, then literal meaning will be grounded in rationally self-perpetuating regularities in the use of words and ways of putting words together. This alternative seems to promise a solution to the problem of meaning without use. Words and ways of putting words together have literal meanings in virtue of their use in complete sentences. And then, because of the compositionality of meaning, these literal meanings for the constituents of sentences determine literal meanings even for sentences that are never used. However, there is a price to be paid for adopting this second alternative.

If the regularities in the use of words and ways of putting words together are to be rationally self-perpetuating then we must credit members of the group (that is, users of the language) with knowledge that they and others participate in these regularities. But it is not really plausible that ordinary speakers know what these regularities of use are. It is more plausible to say that ordinary speakers just *know how* to use words and ways of putting words together in the context of complete sentences.

It seems that the problem of meaning without use cannot be solved within Grice’s programme. But equally, our everyday conception of literal meaning seems to require that the problem must have a solution. Suppose that S is a sentence of my language, built from familiar words in familiar ways, and that I understand it on first hearing. It is part of our conception of literal meaning that S does not come to have a determinate meaning only when I first hear or use it. If S means that *p* in my language then this is what it does and did and would mean whether or not I ever heard or used it.

6.1 Two solutions: Tacit knowledge and compositionality

One kind of solution to the problem of meaning without use involves ideas that go beyond Grice’s programme in the direction of cognitive science. According to the mind-

first view, as we introduced it (section 4.1), the notion of linguistic meaning is to be analysed in terms of mental notions that figure in our everyday personal-level descriptions of ourselves. But, according to this first kind of solution, we need to mention, not only beliefs and intentions, but also cognitive structures and processes.

The envisaged solution would proceed in three steps. First, we would make the Gricean link between meaning and convention only at the level of complete sentences. This would avoid attributing to ordinary speakers knowledge of regularities in the use of words and ways of putting words together. Then, second, we would appeal to the cognitive structures and processes that underpin speakers' assignments of meanings to sentences that are actually used. We would assume that these cognitive structures and processes correspond closely to axioms of a semantic theory and to deductive resources used in (canonical) derivations of meaning-specifying theorems from those axioms. Finally, third, we would appeal to these same axioms and deductive resources to fix the literal meanings of sentences that are unused but are built from the same words, and in the same ways, as sentences that are used.

A natural way of developing this first solution would be to make use of the cognitive scientific notion of tacit knowledge (Chomsky, 1965; Evans, 1981; Davies, 1987). Instead of saying, rather vaguely, that there are cognitive structures and processes that correspond to axioms of a semantic theory and to rules of inference used in derivations, we could say that the axioms and rules are tacitly known.

The possibility of an alternative solution becomes clear if we consider how the problem of meaning without use would arise within Davidson's programme. If a sentence, *S*, is never used then a theorem saying that *S* is true iff *p* cannot make any contribution to an overall interpretation of speakers. So the empirical constraint of interpretational adequacy provides us with no clue as to what *S* means. But a solution to the problem is available within Davidson's programme because semantic theories are subject to formal constraints as well as empirical constraints.

The formal constraints are intended to ensure that a semantic theory is compositional. In a compositional semantic theory, the canonical theorems specifying the truth conditions of used sentences will be derived from axioms that specify the semantic properties of the constituents of those sentences. From those axioms, canonical derivations will lead to theorems specifying the truth conditions of other sentences built from the same constituents, whether those sentences are used or unused.

Although this second solution is suggested by Davidson's programme, it can be applied to the problem of meaning without use that arises for Grice's programme. It avoids the attribution to ordinary speakers of any knowledge, explicit or tacit, about the meaning or use of words and ways of putting words together. Someone who adopts this solution follows the first solution in its first step, by making the Gricean link between meaning and convention only at the level of complete sentences. But the projection of meaning from used sentences to unused sentences is not achieved by following the contours of a semantic theory that is embodied in cognitive structures and processes. Rather, the projection of meaning follows the inductive, abductive, and deductive reasoning of a hypothetical idealised rational subject (section 3.2).

6.2 Comparing the two solutions

The construction of compositional semantic theories contributes to the epistemological project of explaining how knowledge of meaning is possible and how mastery of a language could be a rational achievement. But construction of a compositional semantic theory for a language does not involve any commitment to the idea that ordinary speakers

of the language actually know the facts stated by the axioms of the theory. Indeed, attributing this knowledge to ordinary speakers would seem as implausible as attributing to them knowledge of the regularities in the use of words and ways of putting words together. It would over-intellectualise ordinary language use.

Suppose that the problem of meaning without use is solved in the second way, by appeal to the compositionality of meaning. Then we have an account of how it would be possible for someone – a hypothetical idealised rational subject – to come to know the meaning of a hitherto unused sentence. But this does not bring with it any account of how ordinary speakers actually arrive at such knowledge. If, however, the problem is solved in the first way, by appeal to cognitive structures and processes, then an account of the epistemology of understanding is naturally suggested. When a hitherto unused sentence is heard for the first time, the same cognitive structures and processes whose presence has provided the sentence with its meaning come into play to underpin the speaker's assignment to the sentence of that very meaning.

This epistemological difference between the two solutions might count in favour of the first solution. It might also persuade someone who starts with Davidson's programme to add some cognitive scientific commitments so that ordinary speakers are credited with tacit knowledge of the axioms of a semantic theory. Indeed, there is a programme in philosophy of language that combines Davidsonian truth-conditional semantics with the idea of tacit knowledge of rules and principles that is characteristic of Chomskyan theoretical linguistics (Higginbotham, 1986, 1989).

There is a second difference between the two solutions that is metaphysical rather than epistemological. The two accounts offer different answers to questions about the meanings of unused sentences. Imagine that both Charlene and Bruce have used just fifty – the same fifty – of the one hundred sentences in the toy language considered earlier, and that each name and each predicate term occurs several times in this corpus. Charlene has a structured mastery of these fifty sentences, which is naturally glossed in terms of cognitive structures and processes that embody tacit knowledge of a compositional semantic theory for the language. In contrast, Bruce has an unstructured mastery of these sentences, acquired sentence-by-sentence from a phrasebook.

The first solution, which appeals to cognitive structures and processes, says that each of the one hundred sentences has a determinate meaning in Charlene's language, even though fifty of them have not been used and perhaps never will be used. But the situation with Bruce is quite different. If he has tacit knowledge of a semantic theory, it is a theory with fifty separate axioms, one specifying the meaning of each sentence. So appeal to cognitive structures and processes does not determine a meaning for any sentence beyond the fifty that Bruce actually uses.

While the first solution treats the cases of Charlene and Bruce differently, the second solution groups them together, for it takes no account of actual cognitive structures and processes. Meaning for sentences that are not used is determined by a semantic theory. But it is not a theory that is tacitly known by speakers. Rather, it is a theory that displays the semantic structure that could be used by a hypothetical idealised rational subject; the structure to which Charlene is sensitive, but to which Bruce is blind. So, as in the first solution's treatment of Charlene's case, all fifty unused sentences have determinate meanings.

There are philosophers of language who think that it is definitely wrong to attribute to Bruce, with his phrasebook knowledge of fifty sentences, a language in which a further fifty sentences, of which he knows nothing, have determinate meanings (Schiffer, 1993). If we agree, then we shall prefer the first solution to the problem of meaning without use

over the second solution. If we think that the case for the first solution is compelling, then we shall agree with Brian Loar when he says that ‘the Chomskyan idea of the internalization [tacit knowledge] of the generative procedures of a grammar has got to be invoked to . . . make sense of literal meaning’ (1981, p. 259).

We cannot settle this issue here. But the attractions of the first solution are also attractions of a more general position, encompassing a conception of the relationship between philosophy and science and a view about our everyday personal-level descriptions of ourselves. According to this position, intermediate between cognitive scientism and philosophical isolationism, the philosophically foundational description of ourselves as understanding a language is not reducible to, but is also not independent from, descriptions of the structures and processes that cognitive science investigates.

References

- Chomsky, N. 1965: *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.
- Chomsky, N. 2000: *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.
- Davidson, D. 1984: *Inquiries into Truth and Interpretation*. Oxford: Oxford University Press.
- Davidson, D. 1986: A nice derangement of epitaphs. In E. LePore (ed.), *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*. Oxford: Blackwell, 433–46.
- Davies, M. 1981: *Meaning, Quantification, Necessity: Themes in Philosophical Logic*. London: Routledge and Kegan Paul.
- Davies, M. 1987: Tacit knowledge and semantic theory: Can a five per cent difference matter? *Mind*, 96, 441–62.
- Devitt, M. 1996: *Coming to Our Senses: A Naturalistic Program for Semantic Localism*. Cambridge: Cambridge University Press.
- Dummett, M. 1973: *Frege: Philosophy of Language*. London: Duckworth (Second Edition, 1981)
- Dummett, M. 1978: *Truth and Other Enigmas*. London: Duckworth
- Dummett, M. 1993: *The Seas of Language*. Oxford: Oxford University Press.
- Evans, G. 1981: Semantic theory and tacit knowledge. In S. Holtzman and C. Leich (eds), *Wittgenstein: To Follow a Rule*. London: Routledge and Kegan Paul, 118–37. Reprinted in *Collected Papers*. Oxford: Oxford University Press, 1985, 322–42.
- Fodor, J. 1980: Methodological solipsism considered as a research strategy in cognitive psychology. *Behavioral and Brain Sciences*, 3, 63–109. Reprinted in *Representations: Philosophical Essays on the Foundations of Cognitive Science*. Brighton: Harvester Press, 1981, 225–53.
- Frege, G. 1892: On sense and reference. In P. Geach and M. Black (eds), *Translations from the Philosophical Writings of Gottlob Frege*, Oxford: Blackwell, 1970, 56–78.
- Frege G. 1893: *The Basic Laws of Arithmetic*. Translated and edited, with an Introduction, by Montgomery Furth. Berkeley and Los Angeles: University of California Press, 1967.
- Grice, H.P. 1989: *Studies in the Way of Words*. Cambridge MA: Harvard University Press.
- Higginbotham, J. 1986: Linguistic theory and Davidson's program in semantics. In E. LePore (ed.), *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*, Oxford: Blackwell, 29–48.
- Higginbotham, J. 1989: Knowledge of reference. In A. George (ed.), *Reflections on Chomsky*, Oxford: Blackwell, 153–74.
- Katz, J.J. 1984: An outline of Platonist grammar. In T. Bever, J.M. Carroll and L.A. Miller (eds), *Talking Minds: The Study of Language in Cognitive Science*. Cambridge, MA: MIT Press, 1–33. Reprinted in J.J. Katz (ed.), *The Philosophy of Linguistics*. Oxford: Oxford University Press, 1985, 172–203.
- Loar, B. 1981: *Mind and Meaning*. Cambridge: Cambridge University Press.
- Schiffer, S. 1972: *Meaning*. Oxford: Oxford University Press (Second Edition, 1988).
- Schiffer, S. 1987: *The Remnants of Meaning*. Cambridge, MA: MIT Press.
- Schiffer, S. 1993: Actual-language relations. In J.E. Tomberlin (ed.), *Philosophical Perspectives, 7: Language and Logic*, Atascadero, CA: Ridgeview Publishing Company, 231–58.

- Tarski, A. 1944: The semantic conception of truth and the foundations of semantics. *Philosophy and Phenomenological Research*, 4, 341–75.
- Tarski, A. 1956: The concept of truth in formalized languages. In *Logic, Semantics, Metamathematics*. Oxford: Oxford University Press, 152–278.
- Wiggins, D. 1997: Languages as social objects, *Philosophy*, 72, 499–524.
- Wright, C. 1987: Theories of meaning and speakers' knowledge. In *Realism, Meaning and Truth*. Oxford: Blackwell, 204–238 (Second Edition, 1993).

Further Reading

Chomsky, N. 2000: *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.

This is a collection of Chomsky's recent papers engaging with philosophy of language and particularly with the limits of a scientific investigation of language.

Davidson, D. 1984: *Inquiries into Truth and Interpretation*. Oxford: Oxford University Press (Second Edition, 2001).

For Davidson's programme in philosophy of language, this collection of his essays is essential reading. 'Truth and meaning' is a good place to start, and is also widely reprinted in anthologies including the two edited by Ludlow and by Martinich.

Davis, S. (ed.) 1991: *Pragmatics: A Reader*. Oxford: Oxford University Press.

This anthology covers the key topics in pragmatic theory and includes classic papers by Grice and several papers by Sperber and Wilson.

Dummett, M. 1993: *The Seas of Language*. Oxford: Oxford University Press.

Dummett's papers on philosophy of language, beginning with his two essays addressing the question 'What is a theory of meaning?', are challenging, rewarding, and influential.

Grice, H.P. 1989: *Studies in the Way of Words*. Cambridge MA: Harvard University Press.

This book is a rich resource, both for Grice's work on pragmatics and for his attempts to analyse the concept of literal meaning. On the first theme, 'Logic and conversation' is the place to start; on the second, 'Utterer's meaning and intentions' covers many of the essential points.

Hale, B. and Wright, C. (eds) 1997: *A Companion to the Philosophy of Language*. Oxford: Blackwell.

Pitched at a higher level of difficulty than this Blackwell *Guide*, the *Companion* offers authoritative and sophisticated surveys of many of the topics touched on in this chapter. The first chapter, by David Wiggins, plots a line of thought linking meaning and truth conditions from Frege, through Tarski, to Davidson.

Larson, R. and Segal, G. 1995: *Knowledge of Meaning: An Introduction to Semantic Theory*. Cambridge, MA: MIT Press.

This accessible book offers a very thorough grounding in the approach to semantic theory that combines elements from Davidson's and Chomsky's work, as proposed by James Higginbotham.

Ludlow, P. (ed.) 1997: *Readings in the Philosophy of Language*. Cambridge, MA: MIT Press.

Philosophy of language is an area where most of the reading that is recommended to students is in the form of articles. Ludlow's excellent anthology offers forty-two papers that include good coverage of the foundational issues in this chapter.

Martinich, A.P. (ed.) 2001: *The Philosophy of Language* (Fourth Edition). Oxford: Oxford University Press.

Like Ludlow's, this anthology offers good coverage of the main topics in this chapter and of central issues in the theory of reference.

Schiffer, S. 1987: *The Remnants of Meaning*. Cambridge, MA: MIT Press.

Schiffer's book has been massively influential in persuading most philosophers of language that Grice's analytical programme cannot be carried through. It is not, however, a book for beginners in this area.

Sperber, D. and Wilson, D. 1995: *Relevance: Communication and Cognition*. Oxford: Blackwell (Second Edition).

This is the original and authoritative account of a dominant approach in contemporary pragmatic theory.

Strawson, P.F. 1971: *Logico-Linguistic Papers*. London: Methuen (Revised Edition, Aldershot, UK: Ashgate, 2004).

Strawson's seminal essays on philosophy of language include 'Meaning and truth', in which he offers an important comparison of Davidson's and Grice's programmes.