

Philosophy of Language

Martin Davies

Philosophy of language deals with questions that arise from our ordinary, everyday conception of language. (Philosophy of linguistics, in contrast, follows up questions that arise from the scientific study of language.) But saying this does not yet give a clear idea of the sorts of questions that belong distinctively in philosophy of language. Wittgenstein said (1953, §119), ‘The results of philosophy are the uncovering of one or another piece of plain nonsense and of bumps that the understanding has got by running its head up against the limits of language.’ On this conception, philosophy is about the ways in which we understand and misunderstand language, about how we come to mistake plain nonsense for something that is intelligible, and about what cannot be expressed in language. So, on this view, virtually all of philosophy is concerned with questions about language. It is, indeed, true that language has loomed large in the philosophy of the last hundred years or so. But there is still a specific, recognizable area of the discipline that is philosophy of language. It begins from one absolutely basic fact about language, namely, that expressions of a language have meaning, and can be used to talk about objects and events in the world. For philosophy of language, the central phenomenon to be studied is linguistic meaning. This chapter introduces some of the ways in which that study proceeds. Readers might also like to look at the closely related chapters on PHILOSOPHY OF LOGIC (chapter 4), PHILOSOPHY OF MIND (chapter 5), FREGE AND RUSSELL (chapter 27), and WITTGENSTEIN (chapter 29).

1 Introduction: Questions of Meaning

Questions about meaning are central in the philosophy of language. These questions are of two kinds. On the one hand, there are questions about the meanings of particular linguistic expressions (words, phrases, and whole sentences); on the other hand, there are questions about the nature of linguistic meaning itself. Questions of the first kind belong to *semantics* (section 4); questions of the second kind belong to *meta-semantics* (section 5).

The business of semantics includes questions about the meanings of subject expressions – including PROPER NAMES (chapter 27) (‘Theaetetus’, ‘Fido’) and DESCRIPTIVE PHRASES (pp. 000–00) (‘the man in the gabardine suit’, ‘the present king of France’) – and of PREDICATE (chapter 4) expressions like ‘is sitting’, ‘barks’, ‘is a spy’, and ‘is bald’. It also includes questions about the meanings of complete subject-predicate sentences (‘Theaetetus is sitting’, ‘The present king of France is bald’). There are important philosophical questions about the meanings of other subject terms, including pronouns (‘I’, ‘you’, ‘she’, ‘he’, ‘it’, ‘they’) and demonstratives (‘this’, ‘that’, ‘this knife’, ‘that butter’), and also about the meanings of expressions that go beyond the terms in the basic combination of subject and predicate. Expressions in this latter category include adjectives (‘large’ as it occurs in ‘large flea’, ‘small’ as it occurs in ‘small elephant’, ‘good’ as it occurs in ‘good person’ and in ‘good philosopher’, ‘false’ as it occurs in ‘false sentence’ and in ‘false nose’), modal adverbs (‘possibly’, ‘necessarily’), manner adverbs (‘slowly’, ‘clumsily’), and many more.

Still within the domain of semantics, and closely related to questions about the meanings of words, phrases and sentences, are questions about the ways in which the

meanings of words determine, or at least constrain, the meanings of the phrases and sentences in which they occur. Even if we know what kind of thing the meaning of a sentence is (the meaning of 'Fido barks', or of 'The man in the gabardine suit is a spy'), we also need to understand how component words and phrases make their contributions to the meanings of complete sentences.

On the other side of the semantics versus meta-semantics divide are questions about the nature of linguistic meaning itself. Some of these questions are ontological. Are meanings entities; and, if so, what kinds of entities are they? One putative answer might be that the meanings of complete sentences are *propositions*, and that answer would lead, in turn, to questions about the nature of propositions themselves. An alternative answer might be that, if the meanings of sentences are entities, then they are *states of affairs*. Someone following up this alternative might say that in order for a sentence to be true, the state of affairs that is the sentence's meaning needs to be a state of affairs that *obtains*. *Facts* might then be identified with states of affairs that obtain, and true sentences would be said to be true in virtue of facts. In this case, the answer to the ontological question about meanings would lead to a version of the correspondence theory of truth.

Other meta-semantic questions concern the elucidation or analysis of the concept of meaning. Can we, for example, give any kind of philosophical analysis of the concept of linguistic meaning; and, if we can, what kinds of ideas can legitimately be used in the analysis? How, in general, is the meaning of a linguistic expression related to its *use*? How is the concept of meaning related to the concept of *truth*? In particular, for a complete sentence, what is the relation between the meaning of the sentence and the conditions under which an utterance of the sentence would be true?

The everyday idea of meaning or significance is related to the idea of what is conveyed or communicated in the use of language. In recent philosophy of language, a standard assumption in meta-semantics has been that there is such a thing as the *literal meaning* of linguistic expressions, and that the total communicative significance of a linguistic act is the product jointly of the literal meanings of the expressions used and of contextual factors. According to that meta-semantic view, semantics is the study of the literal meanings of expressions, and of the way that the literal meanings of complex expressions (phrases and sentences) are determined by the literal meanings of their component words. Strictly speaking, questions about the interaction between literal meaning and contextual factors belong, not to semantics, but to *pragmatics* (section 6).

2 Theories of Meaning

Answers to semantic questions and answers to meta-semantic questions can be given by propounding what might be called *theories of meaning*; but that phrase has two very different senses. In the first sense, a theory of meaning answers semantic questions by specifying the meanings of linguistic expressions. In the second sense, a theory of meaning answers meta-semantic questions by providing an elucidatory account of the nature of linguistic meaning.

Consider semantic questions first. These are questions about the meanings of words and phrases, and about the ways in which these contribute to the meanings of whole sentences. For the words, phrases and sentences of a particular language, someone might seek to answer these questions in a very explicit way, by providing a certain kind of formal, axiomatized theory. The idea would be that the axioms of this theory should specify explicitly the meanings of the words of the language under study, and that rules of inference should then permit the derivation, from those axioms, of *theorems* specifying

the meanings of phrases, and ultimately of whole sentences in that language. In the case of any given sentence (say, ‘Fido barks’), the derivation of a meaning-specifying theorem would make use of the axioms of the theory that specify the meanings of the words occurring in that sentence (in this case, the words ‘Fido’ and ‘barks’). This derivation of a theorem from axioms could reasonably be said to display how, in the language under study, the meanings of the component words contribute to the meaning of the complete sentence. A theory that shows how the meanings of sentences depend on the meanings of their parts is sometimes said to be *compositional*.

The construction of theories of the kind envisaged here is not a trivial matter. We might start off by supposing that a typical theorem would say something like:

The meaning of the sentence ‘Fido barks’ is the proposition that the particular dog Fido engages in the activity of barking.

or, avoiding the explicit talk about propositions:

Th1 The sentence ‘Fido barks’ means that the particular dog Fido engages in the activity of barking.

We might also suppose that the axioms from which this theorem is to be derived would say things like:

Ax1 The word ‘Fido’ means a particular dog, namely, Fido.

Ax2 The word ‘barks’ means the activity of barking.

But, even in the context of this extremely simple example, we can see that we would immediately confront at least two important issues. The first issue concerns the difference between subject terms and predicate terms (here, between ‘Fido’ and ‘barks’); the second issue concerns the derivational route from the axioms to the theorem (here, from Ax1 and Ax2 to Th1).

2.1 The difference between subject terms and predicate terms

The issue about the difference between subject terms and predicate terms is this. The two axioms that we proposed have the same form. In each case, the axiom says that an expression stands in the meaning relation to something in the world – in one case, a dog, in the other case, an activity. So, those axioms do not really explain why these two words – ‘Fido’ and ‘barks’ – can go together to make up a sentence, whereas other pairs of words – such as two nouns – cannot together yield a sentence, but merely constitute a list.

We can make this point more vivid if we consider, for a moment, nouns that are closely related to predicate terms. Consider, for example, the noun ‘barking’. This noun can, of course, occur in sentences: ‘The barking of the neighbourhood dogs kept me awake’. Indeed, it can function as the subject term in a sentence: ‘Barking is fun’. But the noun ‘barking’ cannot be juxtaposed with the word ‘Fido’ to make a sentence. Similarly, consider the noun ‘baldness’. If we juxtapose this noun with a subject term, then what we get – say, ‘The present king of France baldness’ – is not a sentence. (In order to see that it is just a list, we can imagine it as an account of the topics discussed at a meeting.)

The point, in short, is that there seems to be an important difference in meaning between predicate terms, such as ‘barks’ and ‘is bald’, and other closely related expressions. The meanings of the predicate terms somehow fit them for combining with subject terms to form sentences, while the meanings of those other expressions (‘barking’, ‘baldness’) do not so fit them. Now, suppose that we were asked to offer a meaning-specifying axiom for the noun ‘barking’ in the style of Ax1 and Ax2. We would be virtually bound to say:

The noun ‘barking’ means the activity of barking.

But, this says just the same about the meaning of the noun ‘barking’ as Ax2 says about the meaning of the verb ‘barks’. Similarly, the model for axioms that we have established thus far would suggest that for the verb phrase ‘is bald’, and for the noun ‘baldness’, the axioms should be:

The phrase ‘is bald’ means the property of baldness.

The noun ‘baldness’ means the property of baldness.

But, as we just saw, while ‘barks’ and ‘is bald’ can function as predicate terms in sentences, ‘barking’ and ‘baldness’ cannot. There must be a difference in meaning between ‘barks’ and ‘barking’, and between ‘is bald’ and ‘baldness’; but the axioms that we have proposed do not reveal what this difference might be.

2.2 The derivational route from axioms to theorems

The issue about the route from axioms to a theorem is this. The derivation of a theorem from axioms is supposed to have the status of a logical PROOF (pp. 000–00), and it should make use of well understood, and well-behaved, logical resources. So, suppose that we apply these requirements to the theorem that we have highlighted (Th1) and the two axioms (Ax1 and Ax2).

Certainly the theorem does not follow from the axioms just by way of the logical resources of the propositional calculus or the PREDICATE CALCULUS (pp. 000–00), for example. The inference from Ax1 and Ax2 to Th1 is not a logically valid one. In order to obtain a valid inference, we would need to add an extra premise, in the form of another axiom.

Clearly, if we were to add the hypothetical statement:

If the word ‘Fido’ means a particular dog, namely, Fido, and the word ‘barks’ means the activity of barking, then the sentence ‘Fido barks’ means that the particular dog Fido engages in the activity of barking.

then we could proceed from this statement plus Ax1 and Ax2 to the conclusion Th1 by way of familiar logical rules of inference. The form of the inference would be:

A, B, if A and B then C; therefore, C.

which is clearly valid. But there would be good reasons to aspire after something more general than this hypothetical statement about the specific words ‘Fido’ and ‘barks’. We should seek an axiom that speaks in general terms of the effect of putting together two terms – subject and predicate – to make a simple sentence.

Provided that we can find some account of what a subject term is and what a predicate term is, something along the following lines might suggest itself:

Ax3 If a subject term M means something, say, X, and a predicate term N means something, say, Y, then the sentence made up of M followed by N means that X engages in, or exemplifies, Y.

This is pleasingly general in the way that it talks about M and N and their meanings X and Y, and it is a generalization that we can, it seems, instantiate in order to yield the more specific hypothetical statement about the subject term ‘Fido’ and the predicate term ‘barks’. If that is right, then the conclusion Th1 can be validly derived from Ax1, Ax2 and Ax3. So, by making explicit the generalization in Ax3, we have made some progress with the requirement that the derivation of a meaning-specifying theorem from axioms is supposed to have the status of a logical proof.

Furthermore, we could instantiate the same generalization, Ax3, to give a hypothetical statement about the subject term ‘Theaetetus’, the predicate term ‘is sitting’, and the sentence that they go together to make up. Taken in conjunction with specific

axioms about the meanings of ‘Theaetetus’ and ‘is sitting’, this hypothetical statement would enable us to derive a theorem specifying the meaning of the sentence ‘Theaetetus is sitting’.

But, although this looks promising (apart from the issue about subject terms and predicate terms), there are still legitimate causes for concern as to whether statements like Ax1, Ax2 and Ax3 interact with standard logical resources in a well-behaved way. In fact, most work in semantics proceeds on the assumption that, in order to be sure of the logical status of our derivations, we should not make use of axioms and theorems that talk explicitly about meaning. We shall return to this point later (section 4.1).

The Notion of Meaning and Standard Logic

In standard logical systems we have a principle – in essence, Leibniz’s Law – saying, amongst other things, that if Fido = Rover then from any statement about Fido we can infer the corresponding statement about Rover, and vice versa. If we now imagine that one and the same dog has two names, ‘Fido’ and ‘Rover’, we can ask whether from the premise:

Fido = Rover

plus:

Ax1 The word ‘Fido’ means a particular dog, namely, Fido.

it follows that:

Ax1’ The word ‘Fido’ means a particular dog, namely, Rover.

Either answer – ‘yes’ or ‘no’ – is potentially problematic.

If Ax1’ does follow, then we can see – by a few more steps – that everything that can be said about the meaning of the word ‘Fido’ can be equally truly said about the meaning of the word ‘Rover’. Furthermore, the same will go for any pair of words that pick out the same object in the world, such as the pair ‘Hesperus’ and ‘Phosphorus’, which both pick out the planet Venus. But, we should want to leave it open, at this early stage of our enquiry, whether such pairs of words can be distinguished in point of their meaning. Indeed, ahead of detailed theoretical considerations, we might reasonably expect that such pairs of words would sometimes differ in meaning, despite picking out the same object.

So, perhaps we should prefer the alternative that says that Ax1’ does not follow from Ax1 plus the premise that Fido = Rover. But, in that case, we shall have to admit frankly that we are departing from what is familiar, and moving to a kind of theory whose logical behaviour is not so well understood.

We have been considering theories of meaning, in the first of two possible senses of that phrase. These are theories that seek to answer semantic questions. A semantic theory may be a theory in a quite formal sense (with axioms and rules of inference by means of which theorems can be derived) and we have noted that there is an issue about the role of the notion of meaning in formal derivations or proofs. But, apart from the worry about logical good behaviour, a semantic theory might take the concept of meaning as an unanalysed, and unexplained, primitive notion. After all, it is the job of a semantic theory to tell us what linguistic expressions *mean*. A theory of meaning in the second sense – a meta-semantic theory – will, in contrast, set out to explain the concept of meaning in other terms. A meta-semantic theory will provide an analysis, or some other kind of

philosophical elucidation, of the notion of linguistic meaning, perhaps by plotting connections between that notion and the notion of use, or the notion of truth.

3 Language, Mind and Metaphysics: Questions of Priority

A semantic theory relates pieces of language to pieces of the world. We use language to talk about the world, and to express our thoughts, which are also about the world. (The aboutness of thoughts is often called *intentionality*.) Talk, thought, and world form a triangle, and in philosophy of language, PHILOSOPHY OF MIND (chapter 5), and METAPHYSICS (chapter 2) we move around this triangle (figure 3.1).

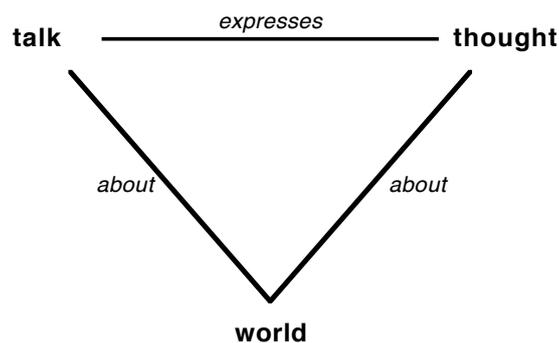


Figure 3.1 The relationship between talk, thought and the world, which is explored in the philosophy of language, the philosophy of mind and metaphysics.

Thus, for example, we might try to give a philosophical account of some distinctions in reality – say, between OBJECTS AND PROPERTIES (pp. 000–00), or between particulars and universals – in terms of differences amongst words, or in terms of differences in the realm of thought, provided that we already had some understanding of those linguistic or mental differences. Or, going the other way about, we might assume some account of the metaphysical differences, and use it in our philosophical work in the domains of talk or thought. We shall shortly consider just such a question of relative priority between philosophy of language and metaphysics.

There are also important questions of priority between philosophy of language and philosophy of mind. Indeed, any strategy for elucidating the concept of linguistic meaning will inevitably depend on our general view of the order of priority as between talk and thought. We need to be clear, first, just what notion of priority is at issue here. Then we shall consider three possible views about language and mind.

3.1 Philosophical priorities: Language and mind

The kind of priority that concerns us here is priority in the order of philosophical analysis or elucidation. To say that the notion of X is *analytically prior* to the notion of Y is to say that Y can be analysed or elucidated in terms of X, while the analysis or elucidation of X itself does not have to advert to Y. Thus to say that the notion of BELIEF (chapter 1) is analytically prior to the notion of KNOWLEDGE (chapter 1), for example, is to say that knowledge can be analysed in terms of belief, while a good analysis of belief does not need to reintroduce the notion of knowledge. (This is just to say what the claim would amount to, not whether it would be correct.)

Analytical priority should be distinguished from ontological priority and from epistemological priority. To say that *X* is *ontologically prior* to *Y* is to say that *X* can exist without *Y*, although *Y* cannot exist without *X*. For example, it might plausibly be said that individuals are ontologically prior to nations. To say that *X* is *epistemologically prior* to *Y* is to say that it is possible to find out about *X* without having to proceed via knowledge about *Y*, whereas finding out about *Y* has to go by way of finding out about *X*. So, it is plausible, for example, that the positions and trajectories of medium-sized material bodies are epistemologically prior to the positions and trajectories of subatomic particles. We can find out about material bodies without investigating subatomic particles but, it might be said, our route to knowledge about subatomic particles has to go via observations of material bodies.

Having distinguished these three kinds of priority, we can make the working assumption that they are logically independent of each other. According to that assumption we can, for example, suppose that the notion of *X* is analytically prior to the notion of *Y*, without being obliged to hold that *X* is either ontologically or epistemologically prior to *Y*. Our question about the order of analytical priority as between language and mind relates, particularly, to the notion of linguistic meaning and the notion of intentionality (aboutness) for mental states. The three possible views that we need to be aware of are these.

Mind first: This is the view that it is possible to give a philosophical account of the intentionality of *thoughts* without essentially adverting to language, and that the notion of linguistic meaning can then be analysed in terms of the thoughts that language is used to express.

Language first: This is the view that an account of linguistic meaning can be given without bringing in the intentionality of thoughts, and that what a person's thoughts are about can then be analysed in terms of the use of language.

No priority (both together): This is the view that there is no way of elucidating the notion of what a person's thoughts are about without bringing in the notion of linguistic meaning, nor the other way around. The two notions have to be explained together.

(There is in fact a fourth possible view, which involves a different kind of no priority claim, namely that the two notions are quite unrelated. But this view, while possible, is quite implausible.)

All three views have adherents. The no analytical priority view is characteristic of the work of Donald Davidson (b. 1917), who couples it with ontological and epistemological no priority claims. Thus, Davidson argues that there can be no thought without language (1975) nor language without thought, and that there is no finding out in detail what a person believes without interpreting the person's speech (1974). The language first view finds expression in Michael Dummett's (b. 1925) writings (see 1973, 1991, 1993), while the mind first approach is taken by Paul Grice (1913–88) (1989; see also Schiffer, 1972). Given what we have assumed about the logical independence of the three kinds of priority, we can see that one option that is available to us is to agree with Davidson in denying the ontological or epistemological priority of mind over language, and yet to follow Grice in trying to analyse the notion of linguistic meaning in terms of the thoughts that language is used to express.

3.2 *Philosophical priorities: Language and the world*

Our example of a question of relative priority as between philosophy of language and metaphysics comes from the work of Peter Strawson (b. 1919) (1959, Part II; 1970a) on subject terms and predicate terms. The distinction in language between subject and predicate terms is intuitively closely related to the metaphysical distinction between objects or particulars, on the one hand (corresponding to subject terms), and properties or

universals, on the other (corresponding to predicate terms). Can we use the metaphysical distinction to help us understand the linguistic one? (As we have already seen, we certainly need some account of the subject-predicate distinction.) Or should we, alternatively, seek a more purely logico-linguistic account of the difference between subject terms and predicate terms, and then try to understand the metaphysical distinction in terms of the linguistic one?

One aspect of the subject-predicate distinction to which Strawson gives considerable attention is what he calls the *asymmetry of subjects and predicates regarding negation*. We can explain the basic idea like this. When a subject-predicate sentence ('Theaetetus is sitting') is negated ('It is not the case that Theaetetus is sitting', or 'Theaetetus is not sitting'), the negation can be taken together with the predicate term to yield a new expression ('not sitting') that is of the same kind – a new predicate term. But the negation cannot be taken together with the subject to yield a new expression ('not Theaetetus') that is of the same kind – a new subject term. In short, predicate terms have negations while subject terms do not.

In order to provide (at least the beginnings of) an explanation of this fact about language, Strawson invites us to consider propositions in which, as he puts it, a general character or kind (or property or universal) is assigned to, or predicated of, a particular, or spatio-temporal individual. The important point is that these general characters come in incompatibility groups vis-à-vis such empirical particulars, while the particulars do not conversely come in incompatibility groups vis-à-vis the general characters. Thus, for example, consider the various colours (general characters) and the sundry items of furniture in the room (spatio-temporal individuals). An item of furniture cannot be more than one colour (all over); whereas more than one item of furniture can be the same colour. Or, consider the various postures that people can adopt (general characters) and the philosophers in a courtyard. Perhaps Theaetetus is sitting, and so not standing or lying, while Theodorus is also sitting and Socrates is standing. As both examples illustrate, general characters come in groups that 'compete' for spatio-temporal individuals; whereas the individuals do not come in groups that 'compete' for the general characters. There is at least the prospect, here, of an explanation of the linguistic distinction between subject terms and predicate terms that appeals to the metaphysical distinction between spatio-temporal individuals and the general characters that they may exemplify.

This approach exhibits an order of explanation that is the opposite of the one that is usually associated with the German mathematician and philosopher Gottlob FREGE (1848–1925) (chapter 27). On the Fregean approach, the metaphysical or ontological category of objects is to be read off from an antecedently fixed linguistic category of logical subject terms (roughly, names). So someone adopting this approach would need to be able to give a purely logico-linguistic criterion for an expression to be a name. It is not an easy matter to discharge this obligation without smuggling back in something like the distinction between objects and properties. Strawson himself, after reviewing some of the logico-linguistic marks of the subject-predicate distinction that might be drawn on in the Fregean order of explanation, remarks (1974a, pp. 13–14):

The general relative natures of the terms combined in the fundamental combination of predication – the differences or alleged differences between subjects and predicates – have so far been discussed in almost exclusively formal terms, in terms belonging to formal logic itself or to grammar. To understand the matter fully, we must be prepared to use a richer vocabulary, a range of notions which fall outside these formal limits. We assume that the subject-predicate duality, and hence the differences so far remarked on, reflect some fundamental features of our thought about the world.

Now, it would not, in fact, be quite accurate to say that Strawson awards absolute priority to metaphysics; but we can, at least, say that he does not award a pre-eminent role to philosophy of language. In this, Strawson's position contrasts with that of Dummett. As we have already noted, when it comes to the question of priority as between language and mind, Dummett adopts the language first option. As we can now add, he also takes a Fregean view on the question about language and metaphysics. So, Dummett does regard the philosophy of language as pre-eminent. Indeed, he once wrote (1978, p. 441):

[W]e may characterise analytical philosophy as that which follows Frege in accepting that the philosophy of language is the foundation of the rest of the subject.

This may not be an entirely happy characterization of analytical philosophy as a whole – since it would leave many analytical philosophers on the wrong side of the classification – but it does indicate one influential line of thought within analytical philosophy.

We should not, however, think that the only possible reason for interest in philosophy of language is that it might hold the key to questions in philosophy of mind and metaphysics. Questions about meaning are no less deep, and no less important, for the fact that, in order to investigate them fully, we may well need to draw on resources coming from other areas of philosophy.

4 Semantic Theories: Davidson's Programme

We have already met the idea of a formal, axiomatized theory that permits the systematic derivation of theorems specifying the meanings of complete sentences drawn from some language (section 2). The sentences might be simple or complex, and the language might be a formal language or a natural language. If we were interested in actually constructing theories of this type, then we might do well to proceed by stages.

The first and simplest stage would involve just a finite stock of sentences from a formal language, and would treat each of the sentences as an unstructured unit. Perhaps there are just three sentences, 'A', 'B' and 'C', where 'A' means that Theaetetus is sitting, 'B' means that Fido barks, and 'C' means that the man in the gabardine suit is a spy. In this case, the idea of an axiomatized semantic theory is simple to the point of triviality. For we can just take three axioms, one stating the meaning of each of the three sentences, and then the project of deriving a meaning-specifying theorem for each sentence of the language quite literally involves taking no steps at all.

The second stage would involve subject-predicate sentences from a formal language. We might have three predicate terms, 'F', 'G' and 'H' and three subject terms, 'a', 'b' and 'c', allowing us to construct nine subject-predicate sentences, amongst which, let us suppose:

'Fa' means that Theaetetus is sitting

'Gb' means that Fido barks

and

'Hc' means that the man in the gabardine suit is a spy.

If we want to be able to derive a meaning-specifying theorem for each sentence in a systematic way, then we might hope to formulate a theory in which there is an axiom talking about the meaning of each subject term and an axiom talking about the meaning of each predicate term (rather along the lines of Ax1 and Ax2 in section 2 above).

At these first two stages, the total number of sentences under consideration is finite. The third stage would involve an infinite collection of sentences built from some finite stock by applying constructions over and over. We might take the three sentences, 'A', 'B' and 'C', from the first stage, and allow the application of a negation operator '¬'.

This would give three more sentences, ‘ $\neg A$ ’, ‘ $\neg B$ ’, ‘ $\neg C$ ’ (meaning that Theaetetus is not sitting, that Fido does not bark, and that the man in the gabardine suit is not a spy, respectively), and then three more, ‘ $\neg\neg A$ ’, ‘ $\neg\neg B$ ’, ‘ $\neg\neg C$ ’, and then three more, indefinitely. Alternatively, we could allow the application of a conjunction operator to two sentences at a time to give, ‘ $A\&B$ ’, ‘ $A\&C$ ’, ‘ $B\&C$ ’, and (if we mark the order of the conjuncts) also ‘ $B\&A$ ’, ‘ $C\&A$ ’, ‘ $C\&B$ ’, and (if we allow repetition of conjuncts) ‘ $A\&A$ ’, ‘ $B\&B$ ’, ‘ $C\&C$ ’ (thus, nine new sentences to add to the original three). Then, applying the conjunction operator again, we get ‘ $(A\&B)\&C$ ’, ‘ $(A\&C)\&C$ ’, ‘ $(B\&C)\&C$ ’, and ‘ $(A\&B)\&(A\&C)$ ’, and many more – in fact, 144 new sentences to add to the twelve.

We could, of course, allow the application of both the negation operator and the conjunction operator, to give us sentences like ‘ $A\&\neg B$ ’. Suppose that we do that, and then try to provide an axiomatized semantic theory for this simple, but infinite, collection of sentences. Along with axioms stating the meanings of our three building blocks, ‘ A ’, ‘ B ’ and ‘ C ’, we shall need an axiom saying something about the meaning of ‘ \neg ’ and one saying something about the meaning of ‘ $\&$ ’.

A first thought in this direction might be to suggest these two axioms:

‘ \neg ’ means negation

‘ $\&$ ’ means conjunction.

But it is quite unclear how, from this axiom about ‘ \neg ’, say, and the axiom stating the meaning of ‘ A ’, we could derive a theorem specifying the meaning of ‘ $\neg A$ ’. What looks more promising, for the negation operator, is something like:

If S is any sentence at all, then the sentence made up by putting ‘ \neg ’ together with S means the negation of whatever S means.

This builds into the specification of the meaning of the operator an indication that it is indeed an *operator* that can be applied to any sentence to yield a new sentence. Something similar can be done for the conjunction operator; but we shall not pause any longer over these details. (In fact, this also suggests how we might make some progress with the issue about the difference between subject terms and predicate terms, discussed in section 2.1. We could build into the axiom for a predicate term an indication that it is indeed a predicate – an expression that can go together with a subject term to make a sentence.)

A formal language with subject terms and predicate terms, plus a couple of sentence-building operators, is clearly very far indeed from the rich complexity of natural language. But, apart from the difference of scale, the step from formal languages to natural languages might seem simple in principle. If we know how to provide an axiomatized semantic theory for a formal language with subject terms and predicate terms, for example, then we also know how to provide a theory for a similarly structured fragment of natural language that uses real words and phrases instead of letters (‘Theaetetus’ instead of ‘ a ’, ‘is sitting’ instead of ‘ F ’).

However, in many cases, the relationship between sentences of a formal language and sentences of a natural language is much less clear. There is a tiny indication of this already in the case of negation, where the formal language operator attaches at the front of a complete sentence whereas natural language negation usually occurs within the predicate term (‘is *not* sitting’, ‘does *not* bark’). And, quite generally, the complexity of the relationship between formal and natural languages is shown by the difficulty of the task (familiar to most philosophy students) of regimenting into logical notation arguments that are expressed in natural language. So, one question that has to be faced is just how to regard the relationship between the superficial forms of natural language sentences, on the one hand, and the regimented forms (sometimes called *logical forms*;

see Sainsbury, 1991) to which an axiomatized semantic theory could be applied, on the other hand.

Even if we ignore these differences between formal languages and natural languages, it is still the case that we have so far only envisaged a semantic theory for a simple subject-predicate language with sentence operators. Subsequent stages of the project of constructing semantic theories would involve bringing further types of expression within the scope of the enterprise.

We should, however, pause here to notice that, while it is clear that axiomatized semantic theories could be of interest to logicians, and – to the extent that they display the meanings of particular kinds of expressions such as pronouns, demonstratives, adjectives, and so on – also to linguists, still it may not be evident why they should merit a philosopher’s attention. In fact, the project of constructing compositional semantic theories has been a central concern in recent philosophy of language, and Davidson’s work in particular reveals several reasons for focusing philosophical attention on semantic theories (Davidson, 1967a; and many other papers collected in Davidson, 1984). We shall review five of these reasons.

4.1 *The format of semantic theories*

The first reason for paying attention to axiomatized semantic theories concerns the proper format for the theorems that such a theory is supposed to yield. Since these theorems are supposed to be meaning-specifying, the initially obvious format would be one that relates a sentence to its meaning. If meanings are themselves regarded as entities of some kind, then we might expect a format like:

The meaning of sentence S is *m*.

If meanings are not regarded as entities, then we might expect instead:

Sentence S means *that p*.

Thus, on the first alternative, if meaning-specifying theorems are explicitly to relate sentences to propositions, say, then an example of such a theorem might be:

The meaning of the sentence ‘Theaetetus is sitting’ is the proposition that Theaetetus is sitting.

On the second alternative, similarly, an example might be:

The sentence ‘Theaetetus is sitting’ means that Theaetetus is sitting.

If, for some general reason, those formats have to be rejected – if the concept of meaning cannot, after all, figure in the theorems of a compositional semantic theory – then that would be a significant meta-semantic result.

Davidson (1967a) does indeed reject both those formats, and argues that the target theorems of a semantic theory should, instead, exhibit the format:

Sentence S is true if and only if *p*

(where ‘if and only if’ (‘iff’) expresses the *material biconditional*). Since theorems like this specify conditions under which a sentence is true, Davidson is said to favour *truth-conditional semantics*. Davidson’s argument for this conclusion comes in two steps.

The first step is intended to rule out the idea that, to each word, each phrase, and each sentence, there should be assigned some entity as its meaning. In this step, the so-called *Frege argument* (Frege, 1892, pp. 62–5; cf. Quine, 1960, pp. 148–9) is used to show that, under certain assumptions, all true sentences would be assigned the same entity. Clearly, no such indiscriminating assignment of entities could be an assignment of meanings, since it is certainly not correct that every true sentence has the same meaning. (Essentially the same line of argument allows Davidson (1969a) to conclude that there is no point in saying that true sentences correspond to facts. The Frege argument shows

that, under certain assumptions, there is only one fact. For further discussion, see Neale (1995).)

The Frege Argument

The Frege argument shows that if an equivalence relation, E , on sentences meets two conditions then it classifies all true sentences as equivalent, and likewise all false sentences. The two conditions are these.

If two sentences are logically equivalent then they are equivalent according to E .

If two sentences are the same except for the fact that, where one contains a subject term i , the other contains another subject term, j , referring to the same thing, then they are equivalent according to E .

Suppose that we take any two true sentences. Let us choose, say 'Penguins waddle' and 'The earth moves'. Then, firstly, the sentence

(1) Penguins waddle

is logically equivalent to

(2) The number that is 1 if penguins waddle and is 0 if penguins do not waddle is 1.

So, by the first condition, (1) and (2) are equivalent according to E . Similarly,

(3) The earth moves

and

(4) The number that is 1 if the earth moves and is 0 if the earth does not move is 1.

are equivalent according to E .

But secondly, the subject term in (2), namely 'The number that is 1 if penguins waddle and is 0 if penguins do not waddle' refers to the number 1, since penguins do indeed waddle. And the subject term in (4), namely 'The number that is 1 if the earth moves and is 0 if the earth does not move', also refers to the number 1. So, sentences (2) and (4) are the same except for the fact that, where one contains a (rather lengthy) subject term, the other contains another (equally lengthy) subject term referring to the same thing. By the second condition, then, sentences (2) and (4) are equivalent according to E .

Putting all this together, (1) is equivalent to (2), and (2) to (4), and (4) to (3). So, (1) is equivalent to (3) according to E ; and that is what we needed to show.

The second step in Davidson's argument for the truth-conditional format points out that the 'means that p ' construction presents logical difficulties, so that the formal derivations of meaning-specifying theorems will be highly problematic (see section 2.2). In contrast, the truth-conditional format is logically well understood. The way in which theorems specifying truth conditions for complete sentences are derived from axioms assigning semantic properties to words and phrases can, to a considerable extent, be carried over from the work of Alfred Tarski (1902–83) on certain formal languages (Tarski, 1944, 1956).

This second step in the argument is technical in character, and it is a matter of dispute whether a compositional semantic theory making direct use of the 'S means that p ' format is feasible. (See Schiffer, 1972, p. 162; Taylor, 1982; Davies, 1984.)

As for the status of the first step in the argument, it is now widely agreed that there are reasons for rejecting Davidson's use of the Frege argument. The Frege argument can only establish that all true sentences would be assigned the same entity as their meaning if it is legitimate to suppose that the equivalence relation of having the same meaning would meet the two conditions that figure in the argument. But this is very far from being obvious. So, pending further argument to the contrary, we can admit that it would be possible for a compositional semantic theory to work by assigning an entity to each sentence as its meaning or semantic value. In particular, one possibility would be to assign to each sentence a certain kind of structured entity – a state of affairs with objects and properties as constituents. The situation semantics programme of Barwise and Perry (1983) constitutes one development of this possibility (see also Taylor, 1976; Forbes, 1989).

4.2 *The correctness of semantic theories*

The second reason for philosophical interest in semantic theories relates to the conditions of adequacy or correctness on any such theory. Whatever is the right format for a semantic theory, we need some account of the conditions under which a theory in that format is the *correct* theory for the language of a particular speaker or group of speakers. We need to know, for example, what makes it correct to say that, in our language, the particular string of letters 'p'- 'e'- 'n'- 'g'- 'u'- 'i'- 'n'- 's' . . . and so on means that penguins waddle.

In order to be able to consider a possible reformulation of this question, we need to make use of the notion of an abstract or possible language. In the abstract, a language can be considered to be a collection of sentences together with a stipulated assignment of meanings to those sentences. So, for example, there are possible languages in which there is a sentence made up of those same letters – 'Penguins waddle' – but with the meaning that the earth moves. Given the notion of a possible language, the question whether a semantic theory is correct for the language of a given group of speakers can be reformulated as the question whether the possible language for which the semantic theory is stipulated to be correct is the *actual language* of a given group (Lewis, 1975; Peacocke, 1976; Schiffer, 1993). What is sometimes called the *actual language relation* is thus a relation between languages (in the abstract) and groups of language users. Under the reformulation that we are envisaging, conditions of adequacy on semantic theories become constraints on the actual language relation.

Any philosophical elucidation of the key semantic concept used in semantic theories, such as meaning or truth, can be transposed into a condition of adequacy on those theories (or, equivalently, into a constraint on the actual language relation). Thus, suppose for example that an elucidation of the concept of meaning says that any sentence *S* has meaning *m* in the language of a group *G* if and only if some condition $C(S, m, G)$ holds. This can be transposed into a condition of adequacy as follows:

If a semantic theory for the language of a group *G* delivers a theorem saying that the meaning of sentence *S* is *m* then it should be the case that $C(S, m, G)$.

Similarly, it can be transposed into a constraint on the actual language relation:

A possible language in which *S* has the meaning *m* is the actual language of a group *G* only if $C(S, m, G)$.

This kind of transposition can be carried out in the opposite direction too. Any condition of adequacy on semantic theories (or any constraint on the actual language relation) can help us elucidate the key semantic concept used in those theories. Thus, for example, consider semantic theories that adopt the truth-conditional format:

Sentence S is true if and only if p .

Tarski's *Convention T* (1956) imposes a condition of adequacy on such theories, namely, the condition that the sentence that fills the ' p ' place should *translate* (or else be the very same sentence as) the sentence S.

This condition of adequacy on truth-conditional semantic theories constitutes a partial elucidation of the semantic concept of truth in terms of the concept of translation. Intuitively, the concept of translation is very closely related to the concept of meaning; and what Convention T requires, in effect, is that the sentence that fills the ' p ' place should have the same meaning as the sentence S. If a truth-conditional semantic theory meets that condition, then the truth condition specifications that it yields are guaranteed to be correct. Thus, Convention T's elucidation of the semantic concept of truth involves a connection between that concept and the concept of meaning:

If a sentence S means that p then S is true iff p .

But Convention T provides no further help with the concept of meaning itself.

One way to shed further light on the concept of meaning – and so, via Convention T, on the concept of truth – would be to spell out other conditions of adequacy on specifications of meanings (or of truth conditions). The concept for which we seek elucidation here is the concept of meaning (or truth conditions) in the language of any group. The concept of meaning is the same whether we consider a group of English speakers or a group of Chinese speakers; so we expect that a condition of adequacy should relate to groups in a very general way. One quite general thing that we can say about specifications of meaning is that they help us to describe members of a group as engaging in linguistic acts. The theorems of a semantic theory for the language of a group, G, can license the re-description of utterances of sentences by a members of G as acts of saying or assertion. For example, if the semantic theory says that a sentence S means that Theaetetus is sitting, then we might reasonably construe an utterer of S as saying or asserting that Theaetetus is sitting. Construing a person's utterances as particular linguistic acts is one aspect of *interpretation*, and what we have just seen is that we can make a link between the theorems of a semantic theory and the project of interpreting the members of a language community.

A semantic theory can play a role in facilitating an overall interpretation of the behaviour of members of a group, by offering a way of understanding their specifically linguistic behaviour. If there are legitimate constraints on ways of interpreting people, then these may yield constraints on the particular aspects of interpretation that are licensed by a semantic theory. In short, the needs of interpretation may lead us to conditions of adequacy on meaning specifications, or truth condition specifications, and so to elucidations of the concepts of meaning and truth.

In Davidson's programme, this idea of a link between semantic theories and interpretation is implemented by describing the constraints on the project of *radical interpretation*. This is the imaginary project of constructing an overall scheme of interpretation for language users about whom we know nothing at the outset. One putative constraint on radical interpretation is that speakers should be so interpreted that what they say about the world – and, presumably, what they believe about the world – turns out to be by and large correct. This is the *Principle of Charity* (see Davidson, 1967a, 1973). Thus, for example (Davidson, 1984, p. 169):

A theory of interpretation cannot be correct that makes a man assent to very many false sentences: it must be generally the case that a sentence is true when a speaker holds it to be.

The Principle of Charity has significance outside philosophy of language, since it appears to offer some prospect for ANTISKEPTICAL ARGUMENTS (pp. 000–00), in particular,

for arguments to the conclusion that most of what we ourselves say and think is correct (Davidson, 1983).

But, in later work by Davidson and others (Grandy, 1973; McDowell, 1976; Wiggins, 1980), the Principle of Charity has given way to a principle that focuses on intelligibility, rather than on correctness: speakers should be so interpreted that what they say and believe about the world turns out to be by and large reasonable or intelligible. As Davidson himself says (1984, p. xvii): ‘The aim of interpretation is not agreement but understanding.’ Sometimes a speaker’s being wrong is quite understandable, while being right would be almost miraculous. An eloquent statement of this principle – sometimes called the *Principle of Humanity* – is provided by Wiggins (1980, p. 199):

Let us then constrain the theory . . . that provides sentence by sentence interpretations of the language L by the requirement that [it] should combine with a plausible anthropology . . . in such a way that *in concert* the two theories make the best sense possible . . . of the total life and conduct of L-speakers.

If the semantic theory that provides sentence by sentence interpretations were to make use of the ‘S means that *p*’ format, then the Principle of Humanity could be conceived as a condition of adequacy on that semantic theory, roughly along the following lines:

If the semantic theory delivers a theorem saying that S means that *p* then it should be the case that interpreting utterances of S as expressions of the proposition that *p* contributes to making the best sense possible of the total life and conduct of speakers.

This condition of adequacy would help us to elucidate the concept of meaning and thence, via the meaning-truth connection, the notion of truth. But it is also possible to regard the Principle of Humanity as providing direct elucidation of the concept of truth, without explicitly going via the concept of meaning.

One way to do this is to introduce a purely formal predicate ‘T’ that applies to sentences and then to consider theories that deliver theorems of the form:

Sentence S is T if and only if *p*.

We can impose a condition of interpretational adequacy on such a theory, roughly along the lines of:

If the semantic theory (canonically) delivers a theorem saying that S is T iff *p* then it should be the case that interpreting utterances of S as expressions of the proposition that *p* contributes to making the best sense possible of the total life and conduct of speakers.

(We need the restriction to *canonical* proofs of T-condition specifying theorems because, for any given sentence S, a theory will deliver many different theorems of the required form (see Davies, 1981, chapter 2).) And then we can see that if a theory meets this condition of adequacy its predicate ‘T’ will apply to precisely the true sentences of the language under study. Wiggins (1980) regards this as showing the way to a ‘substantial theory of truth’. Truth simply *is* the property that plays the ‘T’-role specified by the condition of adequacy that the Principle of Humanity furnishes (see also McDowell, 1976, section 1; Wiggins, 1997).

4.3 Semantic theories and mental states

Some aspects of natural languages pose particular problems for a semantic theorist working within the Davidsonian framework. One of these aspects is the use of natural language to report on people’s mental states. So, a third reason for being interested in semantic theories is that reflection on these theories may yield insight into the meanings of sentences about mental states, and so, also, insight into the nature of those mental states themselves.

To see how the problems arise, consider a belief report sentence like:

(BRep) Theaetetus believes that Fido barks.

Intuitively, the meaning of this sentence depends on the meaning of the name 'Theaetetus', the meaning of the verb 'believes', and the meaning of the contained sentence 'Fido barks'. So, we might have some expectations about the way in which a meaning-specifying theorem for (BRep) would be derived in a compositional semantic theory. We might expect that the derivation would involve, firstly, proving a meaning-specifying theorem for the contained sentence 'Fido barks', and then using that, plus axioms about the meanings of 'Theaetetus' and 'believes', in order to prove the target theorem for the belief report sentence.

In the framework of truth-conditional semantic theories, similarly, we might expect that the proof of a truth-condition-specifying theorem for the belief report sentence would make use of a theorem for 'Fido barks' plus axioms about 'Theaetetus' and 'believes'. But we can see – at least, in an impressionistic kind of way – why this is liable to be problematic.

Statements about people's mental states display some of the same logical properties as statements about meaning (see the box in section 2). With statements about meaning, it is not usually thought to be the case that from:

Fido = Rover

plus:

Sentence S means that Fido barks

it follows that:

Sentence S means that Rover barks.

This logical behaviour is a departure from standard logical systems, and it raises a question about the logical resources that would be needed in a semantic theory using the 'Sentence S means that p ' format. It also helps to motivate use of the alternative 'Sentence S is true iff p ' format, since that is logically well behaved. Thus, for example, from:

Fido = Rover

plus:

Sentence S is true iff Fido barks

it certainly does follow that:

Sentence S is true iff Rover barks.

But, while this well-understood logical behaviour is one of the attractions of truth-conditional semantic theories, it also makes it hard for truth-conditional theories to give an adequate treatment of sentences that report on people's mental states.

With statements about belief, it is not usually thought to be the case that from:

Fido = Rover

plus:

Theaetetus believes that Fido barks

it follows that:

Theaetetus believes that Rover barks.

The first belief report might be true, while the second was false, for example. This is one aspect of what is called the *intensionality* of belief reports.

The problem posed for truth-conditional semantics is that it is hard to see how a truth-condition-specifying theorem for the sentence:

(BRep) Theaetetus believes that Fido barks

can be derived in the way that we envisaged, without also involving the denial of the aspect of the intensionality of belief reports that we just highlighted. The derivational route that we envisaged goes via a truth-condition-specifying theorem for the contained sentence ‘Fido barks’, to the effect that the contained sentence is true iff Fido does indeed bark. Given the assumption that:

Fido = Rover

we can move from this theorem:

(FF) ‘Fido barks’ is true iff Fido barks

to:

(FR) ‘Fido barks’ is true iff Rover barks.

But now, consider whatever route leads from the original theorem (FF) to a truth-condition-specifying theorem for the belief report sentence (BRep), to the effect that it is true iff Theaetetus does indeed believe that Fido barks. Given the assumption that Fido = Rover, an exactly similar route will also lead via (FR) to the conclusion that the same sentence (BRep) is true iff Theaetetus believes that Rover barks. So, given the assumption that:

Fido = Rover

we are driven to the overall conclusion that:

Theaetetus believes that Fido barks iff Theaetetus believes that Rover barks.

But, this goes flatly against the presumed intensionality of belief reports.

The issues here are controversial, and there have certainly been proposals for bringing within the compass of truth-conditional semantics sentences that are usually reckoned to exhibit intensionality (see Davidson, 1969b). For now, however, it is enough to notice that attention to compositional semantic theories could bring questions about mental states into sharper focus.

4.4 Semantic theories and epistemology

The fourth reason for directing philosophical attention onto the construction of semantic theories is epistemological. We have indicated the way in which a finite set of axioms can permit the derivation of a meaning-specifying theorem for each of infinitely many sentences. So, in the case of a given language with infinitely many sentences, someone who knew the facts stated by the finitely many axioms of a compositional semantic theory would thereby be in a position to come to know what each sentence of the language means. The existence of finitely axiomatized semantic theories thus sheds light on an epistemological problem: How is it possible for a finite being to master an infinite language?

So far as the purely formal notion of an axiomatized theory goes, the set of axioms of a theory does not have to be finite. But, Davidson (1965) uses the requirement that a language should be learnable (and learnable by a finite being) to motivate a *finite axiomatization constraint* on semantic theories. Since Davidson favours truth-conditional semantic theories, this yields the idea that the aim of a semantic theorist is to construct finitely axiomatized theories of truth conditions for (regimented fragments of) natural languages – building, so far as possible, on the work of Tarski. The sorts of theories envisaged are sometimes spoken of as Tarski-Davidson truth theories, and there is a very considerable body of work – much of it quite technical in character – in pursuit of this aim.

We should note, however, that there is some dispute over the finite axiomatization constraint. At least one prominent philosopher of language argues (Schiffer, 1987,

chapter 7) that it may be possible for a speaker to master a language, even though no finitely axiomatized truth-conditional semantic theory can be provided for it.

Even if there is a finitely axiomatized semantic theory for a language, there is still a pressing question about the explanatory relevance of the existence of such a theory. Knowledge of the facts stated by the axioms of a compositional semantic theory would suffice, in principle, for knowledge as to what each sentence of the language means. But ordinary speakers of a natural language usually lack conscious knowledge of any compositional semantic theory for their own language. So, the question that presses is how the mere existence of a finitely axiomatized semantic theory, unknown to ordinary language users, can solve the epistemological problem that they apparently face.

One approach to answering this question begins by considering a wider range of knowledge that ordinary language users have. Along with knowing what the sentences of their language mean, ordinary speakers of English also know, for example, that in the sentence:

Nigel shaved him

the word ‘him’ cannot refer back to Nigel, while in the sentence:

Nigel wanted Bruce to shave him

the word ‘him’ can refer back to Nigel. They know that the sentences:

Less than two books were on the table

More than one book was on the table

sound fine, while:

Less than two books was on the table

More than one book were on the table

sound somehow wrong – despite the fact that ‘Less than two books’ is equivalent to ‘At most one book’, and ‘More than one book’ is equivalent to ‘At least two books’. And we could multiply examples of speakers’ linguistic knowledge indefinitely.

Ordinary language users have conscious access to these pieces of knowledge. But they do not have conscious knowledge of any set of axioms or rules from which these particular pieces of knowledge follow. Theoretical linguists articulate sets of linguistic (especially, syntactic) rules or principles called *grammars*. But, the explanatory claims of theoretical linguistics do not end with the mere existence of grammars from which particular pieces of linguistic knowledge could be derived. In a famous passage, Noam Chomsky (b. 1928) says (1965, p. 8):

Obviously, every speaker of a language has mastered and internalized a generative grammar that expresses his knowledge of his language. This is not to say that he is aware of the rules of the grammar or even that he can become aware of them Any interesting generative grammar will be dealing, for the most part, with mental processes that are far beyond the level of actual or even potential consciousness.

According to Chomsky (1965, 1986, 1995), then, ordinary language users possess a body of linguistic knowledge which is, for the most part, inaccessible to consciousness. For this reason, it is often spoken of as *tacit knowledge*. It is because they have this body of tacit knowledge of linguistic rules that they are able to know the vast host of particular things about their language.

While Chomsky is concerned primarily with knowledge of syntax, it seems that the idea that ordinary language users possess a body of tacit knowledge could also be applied in the area of semantics. Thus, one way of answering the question about the explanatory relevance of the existence of compositional semantic theories would be to credit ordinary language users with tacit knowledge of the axioms of such a theory (see below, section

5.3). Some philosophers of language are critical of any appeal to the notion of tacit knowledge, claiming that it embodies conceptual confusions (Baker and Hacker, 1984). But accounts of the notion of tacit knowledge, particularly as it applies to semantic theories, have been offered (see Evans, 1981a; Davies, 1986, 1987, 1989; Peacocke, 1986, 1989), and the proposal that ordinary language users have tacit knowledge of the axioms of truth-conditional semantic theories has been developed in some detail (see Higginbotham, 1985, 1986, 1988, 1989a, 1989b).

We should note, however, that there are other proposals for answering the epistemological question about compositional semantic theories and ordinary language users' knowledge about the meanings of complete sentences. For example, Dummett (1976, 1991) makes use of a notion of *implicit* knowledge that is importantly different from Chomsky's idea of tacit knowledge, and Wright (1986) sees the construction of compositional semantic theories, not as an articulation of language users' actual knowledge, but as a matter of RATIONAL RECONSTRUCTION (pp. 000–00).

4.5 *Semantic theories and metaphysics*

The fifth reason for philosophical interest in semantic theories is that the task of bringing specific linguistic constructions within the scope of a compositional semantic theory sometimes sheds light on issues in other areas of philosophy, and particularly in metaphysics. We have already discussed the relationships between philosophy of language, philosophy of mind, and metaphysics, so we shall be rather brief here.

Semantic theories for natural languages have to deal with pronouns ('I', 'you', 'she', 'he', 'it', 'they'), demonstratives ('this', 'that', 'this knife', 'that butter'), and other indexical expressions ('now', 'then', 'today', 'tomorrow', 'here', 'there'). Questions about the semantic contributions of these expressions (see Kaplan, 1989) inevitably highlight issues in ontology, as well as in philosophy of mind. For, when we consider the mental states that are typically expressed using these expressions, we find that they are mental states that involve thinking about persons, things, times and places in particular ways. So, an adequate account of the thoughts expressed by the use of these expressions seems to involve a commitment to talking about *ways of thinking* or, in Fregean terminology, about SENSES (pp. 000–00) (see Frege, 1918; Perry, 1977, 1979; Evans, 1981b, 1982; Peacocke, 1981, 1983, 1986, 1992, chapter 4; Forbes, 1987; Davies, 1982).

Another example of a link between recent work in philosophy of language and issues in metaphysics – this time not mediated by philosophy of mind – is provided by work on the semantics of names ('Theaetetus'), definite descriptions ('the man in the gabardine suit'), and modal adverbs ('necessarily'), on the one hand, and metaphysical issues about necessity and essentialism, on the other (Kripke, 1980; Wiggins, 1976). A key claim in this area has been that names are, while definite descriptions typically are not, *rigid designators*. A subject term is said to be a rigid designator if it designates the same object with respect to all different possible situations, so long as we hold the meaning of the term constant.

To see the importance of the notion of rigid designation to questions about essentialism, consider the sentence:

Necessarily, the tallest animal in the courtyard is human.

To set the scene, suppose that Theaetetus is the tallest of a group of people talking together in the courtyard, and that there are no non-human animals in the courtyard other than a dog, Fido, and some fleas. Now, whether the sentence containing the modal adverb 'Necessarily' is true depends on whether the truth of the subject-predicate sentence:

The tallest animal in the courtyard is human.

is a matter of *necessity* or *contingency*. This subject-predicate sentence is true as evaluated with respect to the described situation; but could it have been false, without any change in its meaning?

Clearly, the answer is that the sentence could easily have been false. If Theaetetus and his friends had not been in the courtyard, then Fido would have been the tallest animal, and Fido is not human. So, the question over the modal status of the subject-predicate sentence (whether it is necessarily or contingently true) does not turn on any interesting metaphysical question about whether human beings are essentially, or merely contingently, human, for example.

Now consider the sentence:

Necessarily, Theaetetus is human.

Whether this sentence is true depends on whether the truth of the subject-predicate sentence:

Theaetetus is human

is a matter of necessity or contingency; and this time the question of modal status does seem to turn on a question about essentialism. Might this particular human being, Theaetetus, have been other than human?

The definite description, 'The tallest animal in the courtyard', is a non-rigid designator. With respect to the situation as described, it designates Theaetetus, while with respect to an alternative possible situation, it designates Fido – and without any difference in the meanings of the words ('tallest', 'animal', and so on) that make up the definite description. In contrast, the name 'Theaetetus' is a rigid designator. With respect to the situation described, or with respect to the imagined alternative, that name designates Theaetetus – so long as it retains its meaning. What we have just seen is that, where the subject term is a rigid designator, a question about the truth of a sentence containing a modal adverb leads directly to a question about the essential properties of a particular object.

Davidson's own work provides a third example of a link between semantic theories and issues in metaphysics – particularly, of the way in which the construction of a semantic theory may reveal the metaphysical commitments of the conceptual scheme that is expressed in a natural language. Sentences that contain action verbs plus adverbs, such as:

John buttered the toast slowly, in the bathroom

pose a challenge for semantic theories because, after we deal with the subject ('John') and the predicate ('buttered the toast'), it is unclear how we should cope with the adverbs. We shall not pause over the details of the challenge that adverbs pose, but simply note that Davidson (1967b, 1969c, 1985; see also Wiggins, 1986) argues that, in order to bring these sentences within the scope of a truth-conditional semantic theory, we need to regiment them as involving quantification over *events*. Roughly, the idea is that an adverbially modified action sentence should be regimented in logical notation as beginning with ' $(\exists e)$ ___' (in English: 'there is an event which ___') so that, for example, the sentence about John and the toast is regimented as equivalent to:

There is an event e which is a buttering of the toast by John, is slow, and is (that is, occurred) in the bathroom.

If this is correct, then our ordinary talk about people doing things, like buttering toast, carries metaphysical commitments, not only to material things like people and toast, but also to an ontology of events.

Adverbs and Events

Davidson's proposal leads to questions about the notion of logical form, about alternative metaphysical theories of events, and about the role of material bodies as the *basic particulars* in our conceptual scheme (Strawson, 1959, chapter 1).

Firstly, the proposal can be motivated in either of two ways, corresponding to two rather different ideas of logical form (Evans, 1976). On the one hand, there is a pattern of inference to be explained. From 'John buttered the toast slowly, in the bathroom' we can validly infer each of:

John buttered the toast slowly

John buttered the toast in the bathroom

John buttered the toast.

If the adverbially modified sentence is regimented as involving a conjunction ('*e* is a buttering of the toast by John & *e* is slow & *e* occurred in the bathroom') then these inferences can be seen as valid in virtue of the logical properties of '&' (Davidson, 1967b). On the other hand, the regimentation might be seen as required simply in order to bring the sentences within the scope of a compositional truth-conditional semantic theory (Evans, 1976). This motivation might not seem very convincing, since an alternative scheme of regimentation (involving operators that modify the predicate terms corresponding to natural language verbs) appears to be available (Strawson, 1974b). But the apparent alternative is not, in the end, formally viable (Taylor, 1985; Davies, 1991).

Secondly, the proposal leads naturally to the question whether the detailed semantic behaviour of adverbs might provide us with grounds for choosing between two conceptions of events. On the one hand, Davidson (1967b, 1969c) regards events as particulars that can be described in many different ways. One and the same event might be both a *pulling* of a trigger, a *firing* of a rifle, and a *killing* of a president, for example. On the other hand, a number of theorists (Goldman, 1971; Kim, 1976; Taylor, 1985) propose that events should be much more finely discriminated – with firings, pullings, and killings distinguished from each other. There is a substantial metaphysical difference between these accounts, but it turns out that purely semantic considerations do not provide any compelling reason for favouring one over the other (Davies, 1991).

Thirdly, the proposal raises questions about an argument for Strawson's thesis that material bodies are, while events and processes are not, basic particulars in our conceptual scheme. Part of what this thesis comes to is that our ability to pick out individual events in our thought and talk depends on our ability to pick out individual material objects, but not vice versa. At one point, Strawson's argument turns crucially on the claim that our most basic understanding of the conceptually true sentence 'Every animal was born' involves possession of the concept of being born, but not possession of the concept of a birth as an individual particular. Clearly, this argument needs to be assessed in the light of Davidson's proposal for regimenting such sentences as 'This animal was born quickly, in the barn'.

5 Analyzing the Concept of Meaning: Grice's Programme

In the last section, we considered five ways in which semantic theories might reward philosophical attention. The philosophical significance of meta-semantic theories is much more obvious; for they are attempts at an immediately recognizable philosophical project – a project of conceptual elucidation, or even conceptual analysis. Thus, for example, the programme begun by Grice (1957, 1968, 1969) aims at an analysis of the concept of literal meaning in a public language in terms of psychological concepts such as intention and belief. In this section, we shall give an outline description of Grice's programme for analysing the notion of literal meaning. (Grice's papers on the analysis of the concept of meaning are collected in Grice, 1989.)

The analytical project can be regarded as having two stages. The first stage aims to characterize a concept of *speaker's meaning* that corresponds, roughly, to the idea of communicating, or attempting to communicate, a proposition (trying to get across a message). The second stage then aims to use that concept of speaker's meaning, along with the notion of a *conventional regularity*, to construct an analysis of the concept of literal linguistic meaning. The basic idea is that a sentence has as its literal meaning the proposition that it is conventionally used to communicate.

Since the project is an analytical one, it is naturally regarded as being subject to a requirement of non-circularity. The accounts of speaker's meaning and of conventional regularity – notions that are drawn on in the analysis of literal linguistic meaning – should not themselves make essential appeal to linguistic notions. Thus, Grice's programme presumes that mind is analytically prior to language. But the success of Grice's programme is not similarly conditional on any ontological or epistemological priority of mind over language (Avramides, 1989; cf. Evans and McDowell, 1976, pp. xv–xxiii). As we mentioned in section 3.1, it is logically consistent to pursue Grice's analytical programme while accepting Davidson's claims about no ontological priority, and no epistemological priority, of mind over language (though it is true that it would be surprising for someone to combine Grice's analytical views with the no ontological priority claim – no thought without talk).

The claim of no epistemological priority of mind over language says that it is not possible to find out about a person's mental states first, and only then to go on to investigate the meanings of the person's utterances. Rather, the assignment of both meanings and mental states, such as beliefs and intentions, is a single integrated epistemological project governed by the Principle of Charity or the Principle of Humanity (section 3.2). But Davidson goes beyond this epistemological claim. His view is that, by spelling out the constraints on interpretation, we provide all the philosophical elucidation that can be provided of the concept of linguistic meaning. The constraints on interpretation are of two kinds. On the formal side, there is, for example, the finite axiomatization constraint on the semantic theory that is used (in concert with a plausible anthropology) for interpretation. On the empirical side, there are the constraints of agreement (Charity) or intelligibility (Humanity). But, the problem with these empirical constraints is that they seem to leave a lacuna.

An interpretation deploys a semantic theory, which generates an assignment of meanings to sentences, in order to provide a description of language users 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 Principle of Humanity constrains the specifications of meaning by imposing the requirement that the resulting overall description of the language users in terms of meanings, speech acts, and propositional attitudes should make them out to be reasonable or intelligible. But the

Principle of Humanity does not itself tell us which combinations of meanings, speech acts and propositional attitudes can be intelligibly attributed.

On the face of it, an account of which combinations are coherent would be provided by articulating the analytical connections between the concept of meaning, 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 does so by using a sentence that literally means that p , and that anyone who asserts that p intends an audience to take him or her (the speaker) to believe that p . Whether there are connections like this, and if so, what exactly they are, is not a trivial question; it is something that requires detailed investigation. The bold proposal of Grice's analytical programme is that there are connections of this kind that actually permit the analysis of the concept of linguistic meaning (and the concepts of the various speech acts; Schiffer, 1972, chapter 4) in terms of propositional attitudes.

Grice's analytical project begins from a distinction between *natural* and *non-natural* meaning (Grice, 1989, p. 214). We are talking about natural meaning (or *indicator* meaning) when we say, for example:

Those spots mean (indicate) measles.

Those clouds mean (indicate) rain.

This notion of meaning – which is just a matter of causal co-variation between two kinds of states of affairs – is not the concern of Grice's programme, though it takes on considerable importance in recent work in philosophy of mind, aimed at providing an account of the contents of psychological states themselves (see Dretske 1981, 1986).

Spots indicate measles and clouds indicate rain whether or not anyone takes them in that way, and whether or not anyone intends them to be taken in that way. The various notions of meaning that Grice proposes to analyse are distinguished from natural meaning by the fact that they involve, in one way or another, the intentions of a perpetrator or utterer.

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 *say* (to each other) that p .

There are several reasons why this would not be adequate, but the one that concerns us now is that the concept of *saying* seems already to depend on the concept of literal meaning. In order to say something (for example, to say that penguins waddle), a speaker has to use an expression that literally means that thing (here, a sentence that literally means that penguins waddle). So, any analysis along these lines would be circular.

It might be replied to this that there is something stipulative about this conception of saying – that is not obligatory to equate saying with *strictly and literally saying*. It might also be replied that perhaps, in the end, we shall have to be satisfied with an analysis that is circular – a circular analysis might be elucidatory provided that the circle of inter-defined terms is large enough. There is some justice in both these replies. But still, a second suggestion for an analysis of the concept of literal meaning might be this:

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

The advantage of this second suggestion is that the concept of communicating (or getting across) a message does not seem to re-introduce the concept of literal meaning. A person can get across a message even though the sounds used have no literal meaning at all.

However, there is still a problem with this second suggestion. A sound that has no literal meaning can be used by one member of a group to communicate a message to another. The perpetrator of the sound, realizing that the sound has no literal meaning, might rely on some resemblance between the sound produced and some other sound, for example, the sound of an angry dog (see Schiffer, 1972, p. 7, for this example). But then, the same sound might be used by each member of the group when he or she wants to get across the message that he or she is angry. So, the conditions in the suggested analysis would be met. But, this still might not be a case of literal meaning; and, intuitively, the reason for this is that each person who uses a sound like an angry dog to communicate anger might take himself or herself to be making an innovative use of a resemblance between sounds.

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. A little more accurately, we can say that when we make communicative use of an expression with a literal meaning, we rely upon a shared recognition on the part of speaker and hearer that the expression can be used – has been used – in just that way. Communicative use of an expression that has a literal meaning is a rationally self-perpetuating practice; and this fits well with the idea that literal meaning is conventional meaning, since a convention is arguably a rationally self-perpetuating regularity. (The notion of convention is most famously explored by Lewis, 1969.)

5.1 *Speaker's meaning*

We have reached the basic idea that a sentence has as its literal meaning the proposition that it is conventionally used to communicate. Grice's programme aims to develop a detailed philosophical analysis along the lines of this basic idea. The first stage of that programme aims to characterize a concept of *speaker's meaning* roughly corresponding, we said, to the notion of communicating, or attempting to communicate, a proposition. However, we should notice that, given the overall aim of the programme, it is not essential that the concept of speaker's meaning should coincide with any antecedently given notion. It would be consistent with the overall aim that speaker's meaning should come to be regarded as a theory-internal construct.

In Grice's own exposition, the notion of speaker's meaning (utterer's occasion meaning) to be analysed is:

Utterer *U* meant something by his or her utterance *x* directed at an audience *A*.

The putative analysis initially offered by Grice (1957) is as follows (Grice, 1989, p. 92):

U uttered *x* intending

- (1) *A* to produce a particular response *r*;
- (2) *A* to think (recognize) that *U* intends (1);
- (3) *A* to fulfil condition (1) on the basis of his fulfilment of condition (2).

In the case where the speaker is attempting to communicate the message that *p* to the audience, the primary intended response in clause (1) is that *A* should believe that *p*. If we feed this into the analysis, and also unpack clause (3) a little, then we have the proposal that:

Utterer *U* meant that *p* by his or her utterance *x* directed at an audience *A*

should be analysed as:

U uttered *x* intending

- (1) *A* to believe that *p*;
- (2) *A* to think (recognize) that *U* intends (1);

- (3) *A*'s recognition of *U*'s primary intention (1) to be at least part of *A*'s reason for believing that *p*.

The utterance type that *U* uses might or might not have a literal linguistic meaning; and if it does then the communicated message might or might not coincide with that literal meaning. But, crucially, this analysis of speaker's meaning does not itself make use of the notion of literal meaning, and so is available for deployment in a non-circular analysis of that notion.

A host of revisions and extensions have been visited on Grice's initial three-clause analysis (Grice, 1969; Strawson, 1964; Schiffer, 1972). Various counterexamples seem to show that Grice's three conditions are neither sufficient nor necessary for speaker's meaning.

There are two main kinds of reason why the three conditions might not be *sufficient* for the intuitive notion of communicating a message or telling somebody something. Firstly, there is nothing in the three conditions to require that there should be some property of the utterance that the utterer should intend or expect to guide the audience towards recognition of the primary communicative intention (Grice, 1989, p. 94). Secondly, the three conditions do not rule out certain kinds of convoluted plans for influencing a person's beliefs in ways that are not, intuitively, straightforwardly communicative (Grice, 1969; Strawson, 1964; Schiffer, 1972, pp. 17–18). We shall not take time over the details here, but the upshot is that – at least so far as sufficiency is concerned – a better analysis of:

Utterer *U* meant that *p* by his or her utterance *x* directed at an audience *A*

would be:

- (1) *U* intended *A* to believe that *p*;
- (2) there is some feature *F* of the utterance *x* such that *U* intended *A* to recognize *U*'s primary intention (the intention in (1)) in part by recognising *x* to have *F*;
- (3) *U* intended *A*'s recognition of *U*'s primary intention to be at least part of *A*'s reason for believing that *p*;
- (4) *U* did not intend that *A* should be deceived about *U*'s intentions (such as those in (1), (2) and (3)).

The point of clause (4) is, specifically, to rule out cases where *U* intends *A* to be deceived about *U*'s intentions as to how *A* should arrive at the belief that *p* (Grice, 1989, p. 99). It is not supposed to rule out the possibility that *U* might intend *A* to be deceived about other intentions that *U* has. For example, clause (4) is certainly supposed to allow that an act of speaker's meaning might also be a case of lying.

The main reason why the three or four conditions might be regarded as not *necessary* for speaker's meaning is that there are so many things that ordinary speakers do with sentences that would not be included under the three- or four-clause analysis. Indeed, there are many quite literal uses of sentences where the speaker does not even intend to produce in the hearer a belief corresponding to the literal meaning of the sentence (so not even clause (1) would be satisfied). Grice (1989, pp. 105–109) reviews several of these cases, including examination answers, confessions, reminding, giving a review of already known facts, spelling out the conclusion of an argument, and talking to a counter-suggestible person. Davidson (1979) also mentions 'stories, rote repetitions, illustrations, suppositions, parodies, charades, chants, and conspicuously unmeant compliments' (1984, p. 111).

There are various suggested remedies for the problems that these examples are supposed to pose. But, for present purposes, it is enough to note that the massive variety of literal uses of language is not necessarily problematic for the putative analysis of

speaker's meaning, if the fundamental model for speaker's meaning is the case of telling someone something with the intention of providing new knowledge. Arguably, that notion of speaker's meaning – with straightforward telling as the central case – can be analysed roughly along the lines of Grice's four-clause definition. But, if we retain an analysis of speaker's meaning on which the majority of uses of language turn out not to be cases of speaker's meaning, then we must take that fact into account in the second stage of the analytical programme.

5.2 Conventional meaning

Given the basic idea that a sentence has as its literal meaning the proposition that it is conventionally used to communicate, the second stage of the programme should characterize the concept of convention, and most recent work in the Gricean tradition makes use of developments of Lewis's (1969) notion of a convention as a rationally self-perpetuating regularity in behaviour. According to Lewis (1975, pp. 164–6) a convention is a regularity R in action, or in action and belief, that meets the following conditions:

- (1) everyone conforms to R;
- (2) everyone believes that everyone else conforms to R;
- (3) the belief in (2) furnishes each person with a reason to conform to R;
- (4) there is a general preference for general conformity to R, rather than slightly less than general conformity;
- (5) there is at least one alternative regularity which would serve reasonably well; and
- (6) the facts listed in (1)–(5) are matters of *common knowledge*.

The most straightforward way to employ this notion and that of speaker's meaning in tandem in an analysis of literal linguistic meaning is to say:

Sentence S means that *p* in the language of group G
iff

there is a convention in G to use utterances of S in order to communicate that *p* (that is, in order to mean that *p*, in the sense of speaker's meaning).

In a slogan:

Speaker's meaning + Convention = Literal meaning.

This would need to be refined to allow for the ambiguity and context-dependence characteristic of natural language. But, much more importantly, the suggested analysis imposes a requirement that is far too strict to be a necessary condition for literal meaning.

We have already seen that the four-clause definition of speaker's meaning makes it a relatively rare phenomenon. So, there can certainly be literal meaning even where there is no regularity of speaker's meaning (as presently defined). If we want to give an analysis of the concept of literal meaning along the lines of the slogan, then we need to weaken the notion of speaker's meaning, or weaken the notion of convention (or weaken both). (The analyses of literal meaning suggested by recent theorists all depart in one way or another from that most straightforward way of pursuing the second stage of Grice's programme. See Bennett, 1976; Blackburn, 1984, chapter 4; Davies, 1981, chapter 1; Lewis, 1975; Peacocke, 1976; Schiffer, 1972, chapter 5.)

Let us first consider weakening the notion of convention. The trouble with the notion of convention does not seem to lie in the idea of rational self-perpetuation, but rather in the commitment to there being a regularity that members of the group actually conform to. The sentence 'Penguins waddle' literally means, in our language, that penguins waddle. But there is not a regularity of using that sentence to communicate the message that penguins waddle. It is not the case, for example, that whenever anyone wants to get

across the message that penguins waddle he or she uses that sentence. There are surely other ways in which that same message might be communicated. But, more important for present purposes is that, although people do sometimes use the sentence ‘Penguins waddle’, it is not regularly used with the intentions specified in the four-clause definition of speaker’s meaning.

Grice himself (1989, p. 126) suggests moving from speaker’s meaning to literal meaning via the notion of ‘having a certain procedure in one’s repertoire’. It would not be adequate to say:

Sentence S means that p in the language of group G
iff
each member of G has it in his or her repertoire to use utterances of S in order to communicate that p (that is, in order to mean that p , in the sense of speaker’s meaning).

That proposed definition fails to build in the idea of rational self-perpetuation (cf. Schiffer, 1972, pp. 132–6). But a more plausible analysis could be built from the concept of speaker’s meaning (as presently defined) plus a concept analogous to the concept of convention but with the notion of regularity replaced by the notion of having a certain procedure in one’s repertoire. Thus Grice suggests that more is required than each member of the group G having the communicative procedure in his or her repertoire (1989, p. 127):

At least some (many) members of group G have in their repertoires the procedure [to use utterances of S in order to communicate that p] the retention of this procedure being for them conditional on the assumption that at least some (other) members of G have, or have had, this procedure in their repertoires.

If we do not weaken the notion of convention, then we have to introduce a concept weaker than speaker’s meaning. Some literal uses of language that fall outside the scope of speaker’s meaning as presently defined might well fall within the scope of a revised notion in which clause (1) is replaced by:

(1)′ U intended A actively to believe that p .

Arguably, examples of reminding are covered by this modification. Some other literal uses are covered by a different modification:

(1)′′ U intended A to believe that U believed that p .

Arguably, this deals with the case of examination answers. But, in order to make it plausible that we have covered all, or at least most, literal uses of language, we need to abstract away from these details about just what attitude U intends A to take to the proposition that p .

The completely general form would be:

(1)′′′ There is *some* propositional attitude ψ , such that U intends that A should ψ that p .

We could take clause (1)′′′, together with clauses (2)–(4) of the four-clause definition of speaker’s meaning to define a new notion of *weak speaker’s meaning*. This might then be combined with the notion of convention in a putative analysis of literal meaning (Peacocke, 1976; Davies, 1981).

However the details are worked out, a Gricean analysis of the concept of literal meaning for sentences presents the use of a public language as a psychologically highly complex matter. On a Gricean account, everyday literal use of language for communication involves beliefs about beliefs, intentions about beliefs, beliefs about intentions about beliefs, and much more. It may seem implausible that such complicated propositional attitudes are involved in the ordinary use of language and this may motivate an objection to Grice’s programme for analysing the notion of literal linguistic meaning.

But, on the other hand, the psychological plausibility of Gricean accounts has been defended (Loar, 1981). There are important questions here. But, even if they are answered in Grice's favour, there are further difficulties to be faced. For we have not yet taken any account of the internal structure of sentences.

5.3 *Expression meaning and the structure of language*

The way in which component words and phrases make their contributions to the meanings of complete sentences is the starting point for Davidson's programme in philosophy of language (see section 3). But the analysis of the concept of literal meaning as it applies to expressions smaller than complete sentences is an extremely difficult problem within Grice's programme (Grice, 1968; Schiffer, 1972, chapter 6; Loar, 1976, 1981, pp. 253–60). Whether or not it is an absolutely insoluble problem is, at best, not clear. The issues here are very complex, and we shall only give a brief indication of the way in which the difficulty arises.

In essence, the problem is that the structure of language presents a dilemma for the general idea that literal meaning is a matter of there being conventions to use expressions with certain intentions. The primary link between meaning and convention has to be made either at the level of words and phrases, and ways of putting them together (such as subject-predicate combination), or else at the level of complete sentences. Either alternative presents problems.

Suppose that we opt for the first alternative. Then, we shall credit language users with having procedures for using words, and phrases, and ways of putting them together, and also with knowledge that they and others have these procedures. We need to credit speakers with this knowledge, in order to maintain the idea of rational self-perpetuation. But it is here that the problem with this first alternative is presented. For the semantic rules that govern the use of sub-sentential components, and ways of combining components, are far from obvious. (That is why the construction of compositional semantic theories is a challenging project.) So, on the first alternative, it is practically inevitable that the analysis of meaning will commit us to attributing to ordinary language users beliefs and knowledge that they do not have: detailed beliefs and knowledge about the components and construction of their sentences (Loar, 1981, p. 256).

Suppose, instead, that we try the second alternative. In that case, literal meaning will be grounded in rationally self-perpetuating regularities of use of complete sentences. But the problem this time is that there are many perfectly meaningful sentences that are never used at all (*meaning without use*); and there are many others that are used, but would never be used to communicate the proposition that is their literal meaning (*meaning despite use*). An example of the latter kind is provided by the sentence: 'No head injury is too trivial to be ignored' (Bennett, 1976, p. 17). This sentence is sometimes used; but it is typically used with the intention that the hearer should believe that no head injury is *sufficiently* trivial to be ignored.

A possible response to this dilemma would be to avoid the first alternative and adopt the second alternative, but only for those sentences for which there is a practice of literal use. According to this response, the Gricean notion of conventional meaning would not apply directly to unused sentences. So there would remain the problem of extending the notion of literal meaning from the core of used sentences to the rest of the sentences in the language. One plausible suggestion for dealing with this problem – perhaps the only plausible suggestion – is that an account of literal meaning applicable to all sentences, whether used or not, should explicitly advert to the psychological mechanisms that underpin the use of sentences in the core set (Loar, 1976, p. 160; see also Schiffer, 1987,

pp. 249–55, and 1993). According to this suggestion, the analysis of literal meaning would be extended from used sentences to unused sentences in two steps.

Firstly, consider a speaker who participates in a conventional practice of using a sentence with a certain literal meaning. We assume that the speaker's assignment of a meaning to the sentence is underpinned by cognitive mechanisms that correspond closely to some of the axioms of a compositional semantic theory for the language. In particular, we assume that the cognitive mechanisms subserving the assignment of a meaning to a sentence include mechanisms corresponding to the semantic theory's axioms for the words (and ways of combining words) in that sentence.

Secondly, consider an unused sentence that is built from words that occur in used sentences (and is built in ways that are also found in used sentences). Then we allow the literal meaning for the unused sentence to be determined by the semantic axioms that correspond to certain cognitive mechanisms; namely, the mechanisms that are implicated in the speaker's assignment of meanings to regularly used sentences built from those same resources. (Strictly speaking, this deals only with the problem of meaning without use. But a similar strategy could be adopted to deal with the problem of meaning despite use.)

A natural way of implementing this suggestion would be to make use of Chomsky's notion of tacit knowledge. Instead of saying, rather vaguely, that there are cognitive mechanisms corresponding in the appropriate way to axioms of a semantic theory, we could say that those axioms are tacitly known (see above, section 4.4). Indeed, Loar (1981, p. 259) suggests that 'the Chomskyan idea of the internalization of the generative procedures of a grammar has got to be invoked to . . . make sense of literal meaning' (see also Davies, 2000).

Appealing to the notion of tacit knowledge of a compositional semantic theory may be the best way for a Gricean to solve the problem of meaning without use. But we do not have any detailed prospectus for a marriage between Chomsky's and Grice's rather different projects in the study of language. It is certainly not obvious just how Gricean the resulting union could be (Chomsky, 1976, pp. 55–77; 1980, pp. 81–7; 1986, pp. 19–24; see also Laurence, 1996, 1998, for a Chomskyan account of literal meaning presented as a competitor to a Gricean, or 'convention-based' account).

In summary, then, we can say that there is something of a consensus that Grice's programme faces serious problems. Lycan (1991, p. 84) lists some of them:

that most sentences of a language are never tokened at all; that since hearers instantly understand 'novel' sentences, this cannot be in virtue of pre-established conventions or expectations directed on those sentences individually; that sentences are ambiguous and have more than one standard linguistic meaning; that sentences are often (not just abnormally) used with other than their standard meanings; and that indeed some sentences are normally used with other than their standard meanings.

In the last few paragraphs, we have sketched a possible strategy for responding to some of these problems. But we should also note that one of the most authoritative exponents of Grice's programme has, after considering just such a strategy, reached the conclusion that the project of analysing literal meaning in terms of intentions and beliefs (IBS: *intention-based semantics*) cannot be carried through (Schiffer, 1987, p. 261): 'I have exhausted my wits. . . . I cannot see how to devise an IBS account of expression-meaning that does not require us to have knowledge that we seem not to have'.

5.4 Minimalism about meaning and truth

Schiffer (1987, chapter 10) goes beyond the claim that Grice's analytical programme cannot be completed to the more general conclusion that there is no correct and

substantive theory of meaning, in the sense of a philosophical theory about the nature of linguistic meaning, to be given. According to Schiffer, there are no such entities as meanings to be assigned to sentences and, given that Grice's analytical project cannot be completed, there is no prospect of a philosophical analysis of the concept of meaning. There is still room for a theory about how language understanding takes place, about the information processing that does on in our heads when we hear an utterance. But, as Schiffer says, 'such a theory would not be a *philosophical* theory' (1987, p. 269).

This 'no-theory' theory of meaning is a species of *minimalism* about meaning. According to the minimalist view (Johnston, 1988), meaning has no substantial nature and so there is nothing substantive for a theory about meaning to say. Such general elucidation of the notion of linguistic meaning as may be required is provided, not by a substantive theory or conceptual analysis, but by truisms or platitudes, such as:

If a sentence S means that *p* then utterances of S can be used to say that *p*.

Someone working within Grice's programme might hope to provide an analysis of the notion of a speaker using an utterance of S to say that *p* and, in that way, to provide a substantive, if partial, account of the notion of meaning. But the minimalist says that there is no such analysis to be provided.

The minimalist may go on to say that the platitudes about meaning highlight a connection between meaning and *use*: understanding S, or knowing that S means that *p*, is a matter of being able to use utterances of S to say that *p*. If that is right, then knowledge of meaning may appear to be so closely connected to language use that it cannot play any role in the causal explanation of language use. This claim about the lack of an explanatory role for knowledge of meaning is, indeed, one aspect of minimalism about meaning. But the proposal that there is a close connection between meaning and use seems to go against the strictest construal of the 'no-theory' theory of meaning. For the proposal that meaning is closely connected with use, or even that meaning *is* use (Wittgenstein, 1953; Horwich, 1998), does seem to be a substantive piece of philosophical theory concerning the notion of meaning.

It may be useful, here, to distinguish two versions of the 'no-theory' theory of meaning. The strict version says that there is nothing substantive at all to be said about meaning – not even that meaning is use. This might be accompanied by the suggestion that there is nothing more to know about meaning than is constituted by acceptance of such trivial-seeming statements about meaning as:

The sentence 'Fido barks' means that Fido barks.

The sentence 'The man in the gabardine suit is a spy' means that the man in the gabardine suit is a spy.

The sentence 'Theaetetus is sitting' means that Theaetetus is sitting.

and so on. The more moderate version of the 'no-theory' theory of meaning allows that there may be something of philosophical interest to be said about linguistic meaning. It would be consistent with the moderate 'no-theory' theory to offer a philosophical theory according to which facts about meaning are determined by facts about use. But the moderate 'no-theory' theory still insists that there is no substantive philosophical theory about meaning that can be packaged as an analysis of the concept of meaning along the lines of:

Sentence S means that *p* iff ___ S ___ *p* ___ .

Certainly there is nothing of the kind that Grice was aiming to provide.

The idea that meaning depends on use, or that meaning is constituted by use, is not obviously inconsistent with the aim of providing an analysis of the notion of linguistic

meaning. After all, Grice's analytical programme might be regarded as one way of developing the idea that meaning depends on use:

Sentence S means that p iff S is conventionally used to communicate the message that p .

But it is certainly possible to offer a 'use theory' of meaning while explicitly disavowing any analytical ambitions. The term 'use theory of meaning' is usually applied to meta-semantic theories that appeal to the notion of use and also meet the requirement for being a moderate 'no-theory' theory of meaning.

A use theory of meaning, in this sense, is sometimes coupled with minimalism (or a *deflationary* theory) about truth (Horwich, 1990, 1998). The main claim made by a minimalist about the notion of truth is that there is nothing more to grasping the notion of truth than accepting instances of the schema (Horwich, 1990, p. 7):

(E) It is true that p if and only if p

such as:

It is true that Fido barks iff Fido barks.

It is true that the man in the gabardine suit is a spy iff the man in the gabardine suit is a spy.

It is true that Theaetetus is sitting iff Theaetetus is sitting.

Because (E) uses the expression 'It is true that ___', it does not speak of the truth of sentences. There is a clear difference between:

The sentence 'Fido barks' is true iff Fido barks.

which is about a linguistic expression, 'Fido barks', and the instance of (E):

It is true that Fido barks iff Fido barks.

which might be said to concern the *proposition* that Fido barks. So what we have described so far is minimalism about truth for propositions, or minimalism about 'It is true that'.

But a minimalist about truth for propositions is likely to maintain, in addition, that there is nothing more to grasping the notion of truth as it applies to sentences (or utterances of sentences) than acceptance of statements of the following kind:

(An utterance of) the sentence 'Fido barks' is true iff Fido barks.

(An utterance of) the sentence 'The man in the gabardine suit is a spy' is true iff the man in the gabardine suit is a spy.

(An utterance of) the sentence 'Theaetetus is sitting' is true iff Theaetetus is sitting.

In these statements, the same sentence occurs first in quotation marks and then without quotation marks, and statements of this form are said to encapsulate the *disquotational* conception of truth (Field, 1994, p. 250; see also Field, 1986).

Even if minimalism about truth for propositions is plausible, minimalism about truth as it applies to sentences may seem to face a challenge. As we saw in section 4.2, Tarski's Convention T can be conceived as a partial elucidation of the concept of truth for sentences in terms of the concept of translation or the concept of meaning. We also observed that conditions of adequacy on semantic theories (or equivalently, constraints on the actual language relation) can help us to elucidate the concept of meaning and thence, via the meaning-truth connection, the concept of truth for sentences. And we described the way in which Wiggins (1980) offered a 'substantial theory of truth' by imposing a condition of adequacy (based on the Principle of Humanity) directly on Tarski-Davidson theories of truth conditions. By offering the prospect of substantial philosophical elucidation of the concept of truth, these considerations seem to threaten minimalism about truth for sentences.

But the minimalist about truth can respond to this threat of unwanted substantiveness. He might, for example, focus on the application of Tarski's Convention T to the case where a theory of truth conditions for a language is given in an extension of that same language. For, in that case, what Convention T requires of:

Sentence S is true if and only if p

is simply that it should meet the 'disquotational' condition that the sentence that fills the ' p ' place should be the sentence S itself, as in:

The sentence 'Fido barks' is true iff Fido barks.

And he might combine this with insistence that there is nothing more substantive to be said about the conditions under which a theory of truth conditions is correct and, in particular, nothing substantive to be said about the actual language relation (Field, 1994, section 6).

Alternatively, a minimalist about truth who also holds a use theory of meaning (Horwich, 1998) may offer his own account of the meaning-truth connection:

If a sentence S means that p then S is true iff p

drawing on no resources that go beyond the use theory of meaning plus disquotational equivalences such as:

(An utterance of) the sentence 'Theaetetus is sitting' is true iff Theaetetus is sitting

It is plausible, in this case, that the threat of unwanted substantiveness flowing across the meaning-truth connection can be met provided that the use theory of meaning is indeed a moderate 'no-theory' theory. Minimalism about truth avoids a potential problem and, to that extent, gains a measure of support, by being coupled with a moderate 'no-theory' theory of meaning, such as a use theory.

There is also a relation of support in the opposite direction. As Horwich says (1998, p. 113):

By itself, the deflationary [minimalist] view of truth does not motivate the use theory of meaning. But it provides a vital part of the argument by showing that a common reason for rejecting the theory was based on a mistake.

If there were a substantive philosophical theory of truth then, because of the meaning-truth connection, there would be at least a substantive necessary condition on: Sentence S means that p . This would not amount to an analysis of the notion of linguistic meaning, but it would go against the spirit of the 'no-theory' theory of meaning. So a moderate 'no-theory' theory of meaning, such as a use theory, is more plausible when it is accompanied by minimalism about truth.

A use theory of meaning and minimalism about truth for sentences are mutually supporting. Each allows the other to avoid a problem. The combination might be called 'semantic deflationism' (Horwich, 1998, p. 11) – or, perhaps better, 'meta-semantic deflationism'. This deflationary position on meaning and truth is opposed, not only to Grice's analytical programme, but also to Davidson's programme with its focus on the construction of compositional, truth-conditional semantic theories. The evaluation of meta-semantic deflationism is a pressing task for the philosophy of language.

6 Pragmatics: Conversational Implicature and Relevance Theory

A sentence that literally means that p can be used to communicate a message other than the proposition that p . Thus, suppose that someone asks you, 'Would you like more coffee?' and that you answer: 'Coffee keeps me awake.' In a suitable context, this answer can convey the message that you do not want any more coffee (Sperber and Wilson, 1995, p. 34); in a different context (as when you are obliged to stay alert through a boring

lecture), it might convey the message that you do want some more coffee. But neither of these possible messages is any part of the literal meaning of the sentence ‘Coffee keeps me awake’. In each case, the communicated message is, rather, something implied (in some sense) in the context, by what is literally and explicitly said.

The distinction between semantics and pragmatics is, roughly, the distinction between (i) the significance conventionally or literally attached to words, and thence to whole sentences, and (ii) the further significance that can be worked out, by more general principles, using contextual information. Grice’s theory of *conversational implicature* is the seminal contribution to recent pragmatic theory. It is a theory – or the beginning of a theory – of how this further significance is generated.

6.1 Conversational implicature

The theory of conversational implicature was first developed in the context of examples such as these.

and: Strawson (1952) discusses differences between the connectives of propositional logic and the corresponding words of English, and he notes that ‘a statement of the form “p and q” may carry an implication of temporal order’ (1952, p. 81).

or: Strawson also notes that ‘the alternative statement [statement of the form “p or q”] carries the implication of the speaker’s uncertainty as to which of the two it was’ (1952, p. 91).

but: The statement ‘She was poor but she was honest’ carries the implication that there is some kind of contrast between being poor and being honest.

looks: The statement ‘That looks red to me’ carries the implication that it is in some way doubtful whether the object really is red.

We can ask whether these implications are all of the same kind. Are they the products of special features of key words in the sentences – ‘and’, ‘or’, ‘but’, ‘looks’ – or are they generated by some kind of interaction between the literal meanings of the sentences that are uttered and some very general conversational principles?

Grice’s (1961) first exposition of the idea of conversational implicature is particularly concerned with the example of ‘looks’. He wants to maintain that the statement ‘That looks red to me’ is literally true, even when the object in question quite obviously *is* red, so that there is no question of any doubt about its colour. In that paper, Grice develops two notions – the *detachability* and the *cancellability* of an implication – that we can use to help us to identify conversational implicatures and to distinguish them from *conventional implicatures*.

In the example with ‘looks’, the implication is not detachable. This means that there is no form of words that can be used to state just what ‘That looks red to me’ can be used to state (that has the same truth conditions as ‘That looks red to me’), but without carrying the implication. The obvious candidates – ‘That appears red to me’, ‘That seems red to me’ – carry the same implication that there is some doubt about the object’s colour. On the other hand, the implication is cancellable. It is quite consistent to override the implication of a doubt by saying: ‘That looks red to me . . . and it really is red – there’s no doubt about it’. In the examples with ‘and’ and ‘or’, the implication is likewise cancellable but not detachable; and, indeed, that is the typical profile for conversational implicature.

In the case of ‘but’, however, the profile is reversed. The implication is detachable, but not cancellable. The implication is detachable (as Grice uses that term) because there is another word, ‘and’, that makes just the same contribution to truth conditions as ‘but’, but does not generate the same implication: ‘She was poor and she was honest’ does not carry the implication of a contrast between poverty and honesty. The implication is not

cancellable because it would be extremely odd – indeed, inconsistent, in some sense – to say: ‘She was poor but she was honest, though of course I do not mean to imply that there is any contrast between poverty and honesty’.

In the case of the word ‘but’, then, the implication that is generated appears to be a feature of the literal or conventional meaning of that specific word – even though it is not a feature that helps to determine the truth conditions of sentences containing the word. In the other cases – ‘and’, ‘or’, ‘looks’ – the implication is generated by some more general principles governing conversation. Grice (1961) offers a conversational principle that at least applies to the cases of ‘or’ and of ‘looks’. He suggests that the implications are generated, in those cases, by the assumption that there is in operation a general principle something like (1961, p. 94):

One should not make a weaker statement rather than a stronger one unless there is a good reason for so doing.

If someone makes a disjunctive statement, then the hearer can suppose – given this conversational principle – that the speaker is not in a position to make a stronger statement by using just one of the disjuncts. Similarly, ‘That looks red to me’ is intuitively a weaker statement than ‘That is red’, and a speaker’s use of the weaker statement about how it looks to her (or appears to her, or seems to her) generates the implication that she is not in a position to make the stronger statement about how it is in reality. The example with ‘and’, however, is not dealt with by this conversational principle.

In a later paper, Grice (1975) offers a more fully worked out theory that covers the example of ‘and’ as well. (Grice’s papers on pragmatics are collected in Grice, 1989.) The fundamental idea is that participants in a conversation expect each other to observe a *Cooperative Principle*:

Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged.

This principle is elaborated by Grice into a series of maxims. He groups these maxims under four headings: Quantity, Quality, Relation and Manner. Examples of the maxims include (one example from each category):

Make your contribution as informative as is required (for the current purposes of the exchange).

Do not say that for which you lack adequate evidence.

Be relevant.

Be orderly.

The first of these sample maxims is one of the maxims of Quantity, and does the work of the earlier conversational principle. The last is one of the maxims of Manner, and is used to explain the implicature generated by ‘and’.

The way that conversational implicatures are discerned involves the hearer in inference. Suppose that a participant in a conversation says – quite literally and explicitly – that p , and gives no indication that he is not observing the Cooperative Principle and the maxims. Then, on Grice’s theory, a conversational implicature that q is generated if the supposition that the speaker believes that q is required in order to make his literally saying that p consistent with the presumption that he is operating in accordance with the maxims ‘or at least the Cooperative Principle’ (Grice, 1989, p. 31).

According to Grice, conversational implicatures can be worked out or calculated using the Cooperative Principle and the maxims together with information about the literal meaning of the sentence used, information about the context, and other background assumptions. There is no way of telling, ahead of time, what background assumptions might be helpful in working out an implicature. So, this aspect of communication

involves a great deal that is not specifically to do with language. Furthermore, on Grice's view, the Cooperative Principle and the maxims themselves are not specific to language, but are simply the instantiations to the case of conversation of principles and maxims that it would be reasonable to observe in any co-operative activity – such as cooking a meal, or helping a friend to change the wheel on a car.

The mechanism of conversational implicature can be used to communicate more than is literally said – more than the literal meaning that is encoded in the words that are used. If an implicature turns out to be false, then the hearer may well be misled; but the speaker has not, strictly speaking, spoken falsely. This distinction between what is false, and what is only misleading, is of considerable general philosophical importance. If Bruce is standing squarely in front of a standard letter box in the United Kingdom in good daylight, and says, 'That looks red to me', then what he says may be misleading, to the extent that it suggests that there is some doubt about the colour of the letter box. But still, what he says is true: the letter box does look red to Bruce. (Grice (1961) uses this point to defend a version of the sense data theory of perception.)

6.2 Three problems for Grice's pragmatic theory

Although Grice's pragmatic theory is important and has been influential, there are some quite serious problems with it. Here we mention three.

The first problem is that the theory provides no account of how the various maxims are to be weighted against each other. Thus consider the following example:

- A: Where does C live?
B: Somewhere in the South of France.

Grice's own gloss on the example is as follows (1989, pp. 32–3):

There is no reason to suppose that B is opting out [from the Cooperative Principle and the maxims]; his answer is, as he well knows, less informative than is required to meet A's needs. This infringement of the first maxim of Quantity [Make your contribution as informative as is required (for the current purposes of the exchange)] can be explained only by the supposition that B is aware that to be more informative would be to say something that infringed the maxim of Quality, 'Don't say what you lack adequate evidence for', so B implicates that he does not know in which town C lives.

If we look at this gloss in the light of Grice's account of the generation of implicatures, then we see that it is implicit that measuring up to the maxim of Quality excuses violating the maxim of Quantity, when there is a clash between them. The maxim of Quality is thus implicitly ranked as more important than the maxim of Quantity. But there is nothing explicit in Grice's account about the relative ranking of the maxims.

In Grice's pragmatic theory, the Cooperative Principle is developed in a large number of maxims, but no general account is provided of the relationships between the maxims. It would thus be tempting to look for some one overarching aim in conversation, which would justify a single overarching maxim.

The second problem arises from the fact that what is conversationally implicated is supposed to be part of the message that is communicated: it is supposed to fall within the scope of speaker's meaning. It is unclear how, given this fact, Grice's pragmatic theory can allow for the open-ended character of many implicatures – seen clearly in the open-ended interpretation of metaphor. What Grice actually says about open-endedness is this (1989, pp. 39–40):

Since, to calculate a conversational implicature is to calculate what has to be supposed in order to preserve the supposition that the Cooperative Principle is being observed, and since there may be various possible specific explanations, a list of which may be open, the conversational implicature in such cases will be disjunction of such specific explanations; and if the list of these is open, the

implicatum will have just the kind of indeterminacy that many actual implicata do in fact seem to possess.

But, what is conversationally implicated is supposed to fall within the scope of speaker's meaning and it is far from clear that it makes sense to say that the proposition that a speaker overtly intended to communicate was an indefinitely extended disjunction. The problem here is that it is unclear exactly what belief the speaker would intend the speaker to end up with (clause (1) of the definition of speaker's meaning in section 5.1), and equally unclear exactly what intention the speaker would intend the hearer to recognize her as having (clause (2)).

We need to allow that, although sometimes a speaker may overtly intend the hearer to recognize a specific conversational implicature as the communicated message, this case lies at one end of a spectrum. Further along the spectrum, there are cases where the speaker intends that the hearer will recognize one or more implicatures in a more or less closely demarcated range. Further along again, there are cases where the speaker is doing little more than to invite the hearer to explore a space of possible implicatures.

What this suggests is that we might do best to focus on the ways – some specifically intended, others not – in which a hearer can *exploit* what a speaker literally and explicitly says. Conversational implicatures, meeting the standards of speaker's meaning, would then be regarded as the limiting case of a much more general phenomenon.

The third problem is that the account of the generation of conversational implicatures is underspecified. To see this, consider what Grice says about metaphor and irony (1989, p. 34). In each case (metaphor: 'You are the cream in my coffee'; irony: 'X is a fine friend') what is literally and explicitly said is obviously something that the speaker believes to be false. So, on the face of it, the speaker is flouting one of the maxims of Quality: Do not say what you believe to be false. So the hearer has to find a proposition q such that, by supposing the speaker to believe that q , the hearer can see the speaker as nevertheless observing the Cooperative Principle.

In the case of metaphor, Grice's gloss is (1989, p. 34):

The most likely supposition is that the speaker is attributing to his audience some feature or features in respect of which the audience resembles (more or less fancifully) the mentioned substance.

In the case of irony, it is (ibid.):

[U]nless A's utterance is entirely pointless, A must be trying to get across some other proposition than the one he purports to be putting forward. This must be some obviously related proposition; the most obviously related proposition is the contradictory of the one he purports to be putting forward.

But Grice provides no explanation of the way in which the particular relationship between literal meaning and communicated message is to be calculated in each case (but see Grice, 1978, for some further discussion).

6.3 Relevance theory

Dan Sperber and Deirdre Wilson (1981, 1986a, 1986b, 1995; Wilson and Sperber, 1981) have developed a pragmatic theory that is based on a quite general account of the way in which a hearer makes use of available information. They propose that the key to a theoretical understanding of communication – and, indeed, of cognition in general – is provided by the notion of *relevance*. Thus (Sperber and Wilson, 1986b/1991, p. 586):

Our suggestion is that humans tend to pay attention to the most relevant phenomena available; that they tend to construct the most relevant possible representations of these phenomena, and to process them in a context that maximises their relevance. Relevance, and the maximisation of relevance, is the key to human cognition.

In order to explain the notion of relevance, we need to make clear the idea of a *context* that Sperber and Wilson use.

In relevance theory, a context is (1995, p. 15) ‘a subset of the hearer’s assumptions about the world’ that provides potential premises to be used in inferences. The idea is that, at any given point in a conversational exchange, the hearer has some assumptions ‘in the forefront of his attention’. When a new proposition is introduced – say, as the result of a communicative act on the part of the speaker – the newly introduced proposition and the pre-existing context interact inferentially, so as to produce *cognitive effects*. An implication that depends on the new proposition and also on some of the assumptions in the context is said to be a *contextual implication* of the new proposition; and contextual implication is the simplest example of a cognitive effect.

Relevance is a matter of having cognitive effects (1986b/1991, pp. 586):

We claim that information is relevant if it interacts in a certain way with your existing assumptions about the world.

Roughly, the more cognitive effects a proposition has, the more relevant it is. However, the cost of deriving cognitive effects has to be taken into account, as well. So we add a second condition; namely, that the less effort that is required in order to derive cognitive effects from a proposition, the more relevant it is. The notion of relevance then provides an account of utterance interpretation via a *Communicative Principle of Relevance* (1995, p. 158):

Every act of ostensive communication communicates the presumption of its own optimal relevance.

This says, roughly, that an utterance carries a guarantee that it can yield an adequate range of cognitive effects without too much processing effort. (For further discussion of the presumption of optimal relevance, see Sperber and Wilson, 1995, pp. 266–71.)

Relevance theory avoids the first of the three problems for Grice’s pragmatic theory – that it provides no account of the relative weight attached to the various maxims – by having a single principle take over the work that is done by the Cooperative Principle and the various maxims.

Relevance theory avoids the second problem for Grice’s account – that it does not easily allow for the open-endedness of interpretation – because a speaker can communicate to a hearer that worthwhile cognitive effects can be obtained without too much effort, yet not make plain to the hearer just what those cognitive effects will be. Relevance theory thus allows for varying degrees of determinacy of cognitive effects. At one end of the spectrum are the cases that Grice’s pragmatic theory handles well. The speaker makes it completely clear what the intended cognitive effects are, overtly intends the hearer to recognise a specific implicature as the communicated message, and can be held as responsible for the truth of the implicature as if she had asserted it outright. At the other end of the spectrum are cases in which the speaker provides a guarantee of relevance, but leaves it totally up to the hearer to explore a range of cognitive effects. Such exploration is likely to involve considerable processing effort, but the speaker communicates that the effort will be repaid (see Sperber and Wilson, 1995, pp. 193–202).

Sperber and Wilson avoid the third problem for Grice’s account by giving detailed and substantive accounts of metaphor and irony. The key ideas are that metaphor is a kind of *loose talk* while irony is an *echoic* use of language (see Sperber and Wilson, 1995, chapter 4).

Loose talk – as when someone who lives just outside the city limits of Paris says, ‘I live in Paris’ – is a pervasive feature of ordinary language use. Relevance theory accounts for loose talk by saying that the utterance of a sentence that is strictly speaking false may

be a highly effective way of communicating a whole set of propositions that can be easily derived as logical or contextual implications from the initial piece of (mis-)information. A hearer can use a presumption of relevance to select some implications and ignore others. According to relevance theory, metaphor is a kind of loose talk. The perpetrator of a metaphor intends a hearer to derive a set of contextual implications from a proposition that is literally false (Sperber and Wilson, 1986a/1991, p. 548):

The most creative metaphors require of the hearer a greater effort in building an appropriate context, and deriving a wide range of implications. In general, the wider the range of potential implicatures and the greater the hearer's responsibility for constructing them, the more creative the metaphor.

The echoic use of language is also widespread. Thus consider the following exchange (Sperber and Wilson, 1981/1991, p. 556):

- a. I've got a toothache.
- b. Oh, you've got a toothache. Open your mouth, and let's have a look.

Here, the second speaker echoes the first speaker's utterance, thereby indicating that it has been heard and understood and also expressing a reaction to it. Often, an echoic utterance is used to convey a speaker's attitude towards the thought of the person whose utterance is echoed, as in the following example (Sperber and Wilson, 1995, p. 239):

He: It's a lovely day for a picnic.
[They go for a picnic and the sun shines.]
She (happily): It's a lovely day for a picnic, indeed.

It is a very short step from here to a clear example of irony (*ibid.*):

He: It's a lovely day for a picnic.
[They go for a picnic and it rains.]
She (sarcastically): It's a lovely day for a picnic, indeed.

According to relevance theory, irony is a kind of echoic use of language (*ibid.*): 'The speaker dissociates herself from the opinion echoed and indicates that she does not hold it herself.' In general, the interpretation of an utterance as ironical involves the recognition of the utterance as echoic, the identification of the person or kind of person whose utterance (real or imagined) or thought is being echoed, and the recognition that the speaker's attitude towards this thought is one of rejection or dissociation (1995, p. 240).

There is an important similarity between relevance theory's account of irony and the account of metaphor. Both metaphor and irony are said to involve the use of 'very general mechanisms of verbal communication' (*ibid.*, p. 242). There is no need for a distinctive theory of *figurative* language; the relevance-theory accounts of metaphor and irony are continuous with accounts of utterances that are neither metaphorical nor ironical.

Implication and Implicature

We can convey some of the basic ideas of relevance theory by looking at an example involving a question, along with four possible answers (based on Wilson and Sperber, 1981). We are to assume that the set of background assumptions shared between speaker and hearer already includes the assumption that Bruce is a philosopher, but does not include the assumption that Bruce is boring.

Q: Are you inviting Bruce to your party?

A: a. No, I am not inviting Bruce.

b. Bruce is a philosopher, and I am not inviting any philosophers.

c. I am not inviting any philosophers.

d. I am not inviting anyone who is boring.

Each of these four replies could serve to communicate the same message – the negative answer to the question. That message is literally and explicitly expressed in (a), but is less directly conveyed in each of (b)–(d). In the case of (b), the message that the speaker is not inviting Bruce can be inferred from what is literally and explicitly said. In the case of (c), that message can be inferred from what is literally and explicitly said together with a shared background assumption. In the case of (d), the message can be inferred from what is explicitly said together with an additional premise, namely, the proposition that Bruce is boring. The idea is that, in order to maintain the presumption that the speaker's utterance is relevant in case (d), the hearer of that answer will supply the additional premise.

To give a description of even this simple example, we need at least four notions:

the proposition that is literally and *explicitly expressed* in an utterance;

a proposition that is *logically implied* by what is explicitly expressed;

a proposition that is contextually or *pragmatically implied* – where the inference relies on both the proposition explicitly expressed and other premises drawn from the set of background assumptions (the context); and

a proposition that is not already part of the context, but is supplied as an *additional premise*.

These four notions can be used to describe cases of a hearer's exploiting information provided by the speaker, as well as cases – such as the example here – where the speaker overtly intends to communicate a message that is not explicitly expressed.

It may not be immediately obvious how Grice's notion of conversational implicature maps onto this relevance theory framework. In the example given, answer (b) would certainly not be a case of conversational implicature on Grice's account, since (1989, p. 39) 'the truth of a conversational implicatum is not required by the truth of what is said'. Answer (c) does pass that test, but in fact the Gricean notion of conversational implicature corresponds more closely to that of an additional premise, as in answer (d).

6.4 Pragmatic contributions to explicit content

Grice's pragmatic theory offers an account of the way in which an implicature that *q* can be worked out, given that the speaker literally and explicitly advances the proposition that

p. This may suggest that pragmatic principles are not involved in discerning that a speaker has explicitly advanced a certain proposition. But that idea would not be correct. Thus, Sperber and Wilson (1995, pp. 175–6):

We regard verbal communication, then, as involving two types of communication process: one based on coding and decoding, the other on ostension and inference. The coded communication process is not autonomous: it is subservient to the inferential process. The inferential process is autonomous: it functions in essentially the same way whether or not combined with coded communication (though in the absence of coded communication, performances are generally poorer). The coded communication is of course linguistic: acoustic (or graphic) signals are used to communicate semantic representations. The semantic representations recovered by decoding are useful only as a source of hypotheses and evidence for the second communication process, the inferential one. Inferential communication involves the application, not of special-purpose decoding rules, but of general-purpose inference rules, which apply to any conceptually represented information.

On this view, pragmatic principles are involved in the recognition of the proposition that is explicitly expressed. They are also involved when a hearer discerns what attitude the speaker takes towards the proposition expressed, or what kind of speech act – an assertion, a command, a question – the speaker is performing. Let us consider the recognition of speech acts first and then the recognition of the proposition explicitly expressed.

If we are to appreciate what is involved in discerning what kind of speech act a speaker is performing then it is important to distinguish the idea of a kind of speech act from the grammatical notion of *mood*. The mood of an uttered sentence – indicative, imperative, interrogative – is part of its encoded literal meaning. But, uttering a sentence in a particular mood does not guarantee the performance of a particular kind of speech act. Uttering a sentence in the indicative mood, for example, is not the same thing as making an assertion. An indicative sentence can be used to express a conjecture or to make a joke; it can be used ironically or questioningly. Nor is there any further conventional sign, over and above the mood of the uttered sentence, that guarantees that an act of assertion has been performed. Uttering a sentence in a serious tone of voice, for example, does not inevitably make the utterance into an assertion.

It is true that there are some conceptual connections between moods and kinds of speech acts, and an adequate theory of language and its use must spell those connections out. Davidson (1979) makes the proposal that uttering a sentence in a non-indicative mood (for example, the imperative mood or the interrogative mood) is rather like making a pair of utterances. An utterance of ‘Put on your hat’ is rather like an utterance of ‘My next utterance is imperatival in force’ (or ‘My next utterance constitutes an act of commanding’) followed by an utterance of the indicative sentence ‘You will put on your hat’. Whether or not Davidson’s proposed account is ultimately satisfactory, it certainly establishes a connection ‘between the mood indicators and the *idea* of a certain illocutionary act’ (1984, p. 275). But the crucial point for present purposes is that even though the presence of a particular mood may constitute a *prima facie* reason for taking an utterance to be a speech act of a particular kind, still that *prima facie* reason may be overridden by all manner of other considerations. There are no ‘special-purpose decoding rules’ for calculating what kind of speech act a speaker is performing.

We turn now to the recognition of the proposition explicitly expressed. It is beyond dispute that the literal meanings of words and the ways in which words are put together to make sentences may leave us far short of determinate truth conditions. A familiar example is provided by the sentence ‘I am tired’. This sentence does not, by itself, have

determinate truth conditions; the conditions for the truth of a particular utterance of ‘I am tired’ depend on who is speaking.

But this kind of case of the under-determination of truth conditions by literal meaning is apt to be misleading. For it is very plausible that there is a rule of language that specifies the way in which the truth conditions of an utterance of ‘I am tired’ depend on a particular feature of the context in which it is made, namely, the speaker in that context. (Here, we use the ordinary notion of context, rather than the relevance-theory notion of a set of assumptions.) In general, however, the step from the literal meaning of a sentence to the truth conditions of an utterance of that sentence in context is not governed by any such neat rules of context-dependence.

For a pronoun, such as ‘he’, ‘she’ or ‘it’, there is no rule that assigns a reference on the basis of some pre-determined feature of the context in which the pronoun is uttered. It is not true, for example, that the reference of ‘he’ is always the nearest male in the context or the last-mentioned male in the context. It may be said that the reference of ‘he’ is always the most *salient* male in the context; but there is no limit to the information that might, in principle, be involved in an assessment of salience in a context. What does for the assignment of reference to pronouns goes also for the removal of lexical ambiguity, the resolution of vagueness or indeterminacy, and the restoration of ellipsis. These are all sensitive to global features of the discourse situation.

Thus, many aspects of the recognition of the proposition explicitly expressed by an utterance involve pragmatic interpretation; they are matters of inference rather than of decoding (Sperber and Wilson, 1986b/1991, p. 585):

Pragmatic interpretation seems to us to resemble scientific theorizing in essential respects. The speaker’s intentions are not decoded but non-demonstratively inferred, by a process of hypothesis formation and confirmation which, like scientific theorizing and unlike grammatical analysis, has free access to contextual information.

6.5 Utterance interpretation and public language meaning

The fact that even explicit content – even an initial proposition from which further inferences might be drawn – depends on pragmatic, as well as semantic, factors is of the greatest importance for a theoretical understanding of the communicative use of language. But we must also note that this fact has been used to cast doubt on the theoretical significance of the very notion of literal meaning in a public language.

A hearer not only uses inference and contextual information in order to bridge the gap between literal meaning and explicit propositional content and then to draw out contextual implications. A hearer may also use those same resources in order to revise his assessment of a speaker’s literal meanings themselves. Literal meaning provides an outline that is then filled in under the guidance of a presumption of optimal relevance. But, if no way of filling in the outline measures up to this presumption, then the outline itself may be redrawn. This is what usually happens when a speaker makes a slip of the tongue or commits a malapropism.

It may happen, then, that a speaker and a hearer bring to a discourse different assignments of meanings to words. A hearer may begin with one set of expectations and, finding them unfulfilled, may set about devising a new assignment of meanings to words to serve for the discourse in which he is engaged. This process of assigning new meanings is like pragmatic interpretation. It involves non-demonstrative inference rather than decoding, and there is no limit to the information on which it might draw.

Davidson says that, in devising a new assignment of meanings, a hearer will operate in essentially the same way as a radical interpreter: he will try to make the best sense

possible of the speaker's total life and conduct (section 4.2). Thus (Davidson, 1984, p. 278):

agreement on what a speaker means by what he says can surely be achieved even though a speaker and hearer have different advance theories as to how to interpret the speaker. The reason this can be is that the speaker may well provide adequate clues, in what he says, and how and where he says it, to allow a hearer to arrive at a correct interpretation.

The notion of convention looms large in a Gricean analysis of literal meaning. But the moral that Davidson draws from consideration of a hearer's ability to revise his assignment of meanings to a speaker's words is that conventions are, in the end, of merely practical significance (ibid., p. 279):

Knowledge of the conventions of language is thus a practical crutch to interpretation, a crutch we cannot in practice afford to do without – but a crutch which, under optimum conditions for communication, we can in the end throw away, and could in theory have done without from the start.

A convention is a rationally self-perpetuating regularity in the behaviour of members of a group. Literal linguistic meaning is, on a Gricean account, meaning that is assigned as a matter of convention. The assignment of meanings is shared by members of the group and is perpetuated through time. But, Davidson says, it is not essential for successful linguistic communication that there should be an assignment of meanings that is shared between speaker and hearer over an extended period of time. All that is needed is that the speaker and the hearer should be able to converge on an assignment of meaning at the time of their communicative exchange.

Utterance interpretation always involves the inferential processes of pragmatic interpretation; decoding of shared literal meanings is never sufficient by itself. Where literal meanings are not shared, pragmatic interpretation (or similar processes of non-demonstrative inference) may in principle still allow the hearer to work out what proposition the speaker was trying to communicate. Sperber and Wilson say that pragmatic interpretation is autonomous while decoding is not autonomous (1995, pp. 176). The conclusion that Davidson draws is that all utterance interpretation is to be seen, for theoretical purposes, on the model of radical interpretation. Since shared literal meanings are not crucial for successful communication, a philosophical theory of linguistic communication will not assign a crucial role to the notion of public language meaning.

It is not clear that the facts about utterance interpretation really license the conclusion about the theoretical dispensability of the notion of literal linguistic meaning. Nor is it obvious that the facts about communication undermine the Gricean analytical programme. After all, Grice himself began with the notion of speaker's meaning, a notion of communication that is supposed to be analytically prior to the notion of literal meaning. But we shall not attempt to adjudicate these issues here.

How much importance attaches to the notions of a shared public language and of literal linguistic meaning? If the notion of literal linguistic meaning is to be retained, can anything more be said about meaning than meta-semantic deflationism suggests? If there is more to be said, then what concepts might be drawn on in a substantive theory about the nature of meaning? Does the notion of convention have a role to play? Should a philosophical account of meaning appeal to the concept of tacit knowledge that figures in theoretical linguistics? These are amongst the important questions for future philosophy of language.

Further Reading

Philosophy of language is an area where most of the reading that is recommended to students is in the form of articles. The collection edited by A.P. Martinich, *The Philosophy of Language* (Fourth Edition, 2001) contains forty one papers. These offer good coverage of some of the main topics in this chapter and also include thirteen articles on referring expressions such as names, definite descriptions, and demonstratives. Peter Ludlow's excellent anthology, *Readings in the Philosophy of Language* (1997) offers forty two papers, including seventeen on referring expressions. Adrian Moore's more compact *Meaning and Reference* (1993) has fourteen papers, some of which take up questions that we have noted, but have not discussed, in this chapter. The collection, *Truth and Meaning: Essays in Semantics*, edited by Gareth Evans and John McDowell (1976) contains high quality papers at an advanced level. The *Blackwell Companion to the Philosophy of Language* (1997), edited by Bob Hale and Crispin Wright, offers authoritative and sophisticated surveys of many of the topics touched on in this chapter.

Useful recent textbooks, at a fairly introductory level, include Michael Devitt and Kim Sterelny, *Language and Reality: An Introduction to the Philosophy of Language* (Second Edition, 1999), William Lycan, *Philosophy of Language: A Contemporary Introduction* (2000), and Kenneth Taylor, *Truth and Meaning: An Introduction to the Philosophy of Language* (1998). Amongst older books, Simon Blackburn's *Spreading the Word* (1984) is engaging and challenging, and offers good coverage of both Davidson's and Grice's programmes. Mark Platts's *Ways of Meaning* (1979; Second Edition, 1997) is useful for an accessible – though far from neutral – account of Davidson's programme.

For Davidson's programme, his book of essays, *Inquiries into Truth and Interpretation* (1984) is essential reading. 'Truth and meaning' (1967a) is a good place to start, and is widely reprinted in anthologies. For an introduction to Chomsky's approach to the study of language, the first chapter of the seminal *Aspects of the Theory of Syntax* (1965) can be combined with a more recent paper, such as 'Language and problems of knowledge' (1996). Chomsky's *New Horizons in the Study of Language and Mind* (2000) is a collection of recent papers engaging with philosophy of language. Richard Larson and Gabriel Segal's *Knowledge of Meaning: An Introduction to Semantic Theory* (1995) 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.

For Grice's programme, his *Studies in the Way of Words* (1989) is a rich resource. 'Utterer's meaning and intentions' (1969) covers many of the essential points in his attempt to analyse the concept of literal meaning. The first thirty pages or so of Stephen Schiffer's *Meaning* (1972) provide an authoritative introduction to Grice's programme, and the Introduction to the Second Edition (1988) sets the programme against a background of more recent concerns. Peter Strawson's Inaugural Lecture, 'Meaning and truth' (1970b) offers an important comparison of Davidson's and Grice's programmes. Schiffer's *Remnants of Meaning* (1987) 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. Reading on minimalism about meaning and truth should begin with Paul Horwich's two books *Truth* (Second Edition, 1998) and *Meaning* (1998), and Hartry Field's 'Deflationist views of meaning and content' (1994).

Grice's 'Logic and conversation' (1975) should be the starting point for reading on conversational implicature. Sperber and Wilson's 'Pragmatics and modularity' (1986b) provides an excellent introduction to their book *Relevance: Communication and Cognition* (Second Edition, 1995). For readers who are particularly interested in

metaphor, the few comments in Grice (1975) could lead on to John Searle, 'Metaphor' (1979) and Davidson's 'What metaphors mean' (1978), and then to Sperber and Wilson's excellent 'Loose talk' (1986a). Steven Davis's reader, *Pragmatics* (1991), covers the key topics in pragmatic theory thoroughly (thirty five papers in all) with classic papers by Grice, several papers by Sperber and Wilson, and a fine section on metaphor and irony.

References

- Avramides, A. 1989: *Meaning and Mind: An Examination of a Gricean Account of Language*. Cambridge MA: MIT Press.
- Baker, G.P. and Hacker, P.M.S. 1984: *Language, Sense and Nonsense: A Critical Investigation into Modern Theories of Language*. Oxford: Blackwell.
- Barwise, J. and Perry, J. 1983: *Situations and Attitudes*. Cambridge MA: MIT Press.
- Bennett, J. 1976: *Linguistic Behaviour*. Cambridge: Cambridge University Press.
- Blackburn, S. 1984: *Spreading the Word*. Oxford: Oxford University Press.
- Chomsky, N. 1965: *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.
- Chomsky, N. 1976: *Reflections on Language*. London: Fontana/Collins.
- Chomsky, N. 1980: *Rules and Representations*. Oxford: Blackwell.
- Chomsky, N. 1986: *Knowledge of Language: Its Nature, Origin, and Use*. New York: Praeger.
- Chomsky, N. 1990: Language and problems of knowledge. In Martinich (ed.), 2001, 581–99.
- Chomsky, N. 1995: *The Minimalist Program*. Cambridge, MA: MIT Press.
- Chomsky, N. 2000: *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.
- Davidson, D. 1965: Theories of meaning and learnable languages. In Y. Bar-Hillel (ed.), *Logic, Methodology and Philosophy of Science*, 2, Amsterdam: North Holland. Reprinted in Davidson, 1984, 3–15.
- Davidson, D. 1967a: Truth and meaning. *Synthèse*, 17, 304–23. Reprinted in Davidson, 1984, 17–36. Also reprinted in Ludlow (ed.), 1997; in Martinich (ed.), 2001; and in Moore (ed.), 1993.
- Davidson, D. 1967b: The logical form of action sentences. In N. Rescher (ed.), *The Logic of Decision and Action*, Pittsburgh: University of Pittsburgh Press. Reprinted in Davidson, 1980, 105–22. Also reprinted in Ludlow (ed.), 1997.
- Davidson, D. 1969a: True to the facts. *Journal of Philosophy*, 66, 748–64. Reprinted in Davidson, 1984, 37–54.
- Davidson, D. 1969b: On saying that. In D. Davidson and J. Hintikka (eds), *Words and Objections: Essays on the Work of W.V. Quine*, Dordrecht: Reidel. Reprinted in Davidson, 1984, 93–108. Also reprinted in Ludlow (ed.), 1997; and in Martinich (ed.), 2001.
- Davidson, D. 1969c: The individuation of events. In N. Rescher (ed.), *Essays in Honor of Carl G. Hempel*, Dordrecht: Reidel. Reprinted in Davidson, 1980, 163–80.
- Davidson, D. 1973: Radical interpretation. *Dialectica*, 27, 313–28. Reprinted in Davidson, 1984, 125–39.
- Davidson, D. 1974: Belief and the basis of meaning. *Synthèse*, 27, 309–23. Reprinted in Davidson, 1984, 141–54. Also reprinted in Martinich (ed.), 2001.
- Davidson, D. 1975: Thought and talk. In S. Guttenplan (ed.), *Mind and Language*, Oxford: Oxford University Press. Reprinted in Davidson, 1984, 155–70.
- Davidson, D. 1978: What metaphors mean. *Critical Inquiry*, 5, 31–47. Reprinted in Davidson, 1984, 245–64. Also reprinted in Davis (ed.), 1991; and in Martinich (ed.), 2001.
- Davidson, D. 1979: Moods and performances. In A. Margalit (ed.), *Meaning and Use*, Dordrecht: Reidel. Reprinted in Davidson, 1984, 109–21.
- Davidson, D. 1980: *Essays on Actions and Events*. Oxford: Oxford University Press.

- Davidson, D. 1983: A coherence theory of truth and knowledge. In E. LePore (ed.), *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*, Oxford: Blackwell, 307–19.
- Davidson, D. 1984: *Inquiries into Truth and Interpretation*. Oxford: Oxford University Press.
- Davidson, D. 1985: Adverbs of action. In B. Vermazen and M.B. Hintikka (eds), *Essays on Davidson: Actions and Events*, Oxford: Oxford University Press, 230–41.
- Davies, M. 1981: *Meaning, Quantification, Necessity: Themes in Philosophical Logic*. London: Routledge and Kegan Paul
- Davies, M. 1982: Individuation and the semantics of demonstratives. *Journal of Philosophical Logic*, 11, 287–310. Reprinted in Ludlow (ed.), 1997, 745–67.
- Davies, M. 1984: Taylor on meaning-theories and theories of meaning. *Mind*, 93, 85–90.
- Davies M. 1986: Tacit knowledge, and the structure of thought and language. In C. Travis (ed.), *Meaning and Interpretation*, Oxford: Blackwell, 127–58.
- Davies, M. 1987: Tacit knowledge and semantic theory: Can a five per cent difference matter? *Mind*, 96, 441–62.
- Davies, M. 1989: Tacit knowledge and subdoxastic states. In A. George (ed.), *Reflections on Chomsky*, Oxford: Blackwell, 131–52. Reprinted in C. Macdonald and G. Macdonald (eds), *Philosophy of Psychology: Debates on Psychological Explanation*, Oxford: Blackwell Publishers, 1995, 309–30.
- Davies, M. 1991: Acts and scenes. In N. Cooper and P. Engel (eds), *New Inquiries into Meaning and Truth*, Hemel Hempstead: Simon and Schuster, 41–82.
- Davies, M. 2000: Persons and their underpinnings. *Philosophical Explorations*, 3, 43–62.
- Davis, S. (ed.) 1991: *Pragmatics: A Reader*. Oxford: Oxford University Press.
- Devitt, M. and Sterelny, K. 1987: *Language and Reality: An Introduction to the Philosophy of Language*. Oxford: Blackwell (Second Edition, 1999).
- Dretske, F. 1981: *Knowledge and the Flow of Information*. Oxford: Blackwell.
- Dretske, F. 1986: Misrepresentation. In R. Bogdan (ed), *Belief: Form, Content and Function*, Oxford: Oxford University Press, 17–36. Reprinted in S.P. Stich and T.A. Warfield (eds), *Mental Representation: A Reader*, Oxford: Blackwell, 157–73.
- Dummett, M. 1973: *Frege: Philosophy of Language*. London: Duckworth.
- Dummett, M. 1976: What is a theory of meaning? (II) In Evans and McDowell, 1976, 67–137. Reprinted in Dummett, 1993.
- Dummett, M. 1978: *Truth and Other Enigmas*. London: Duckworth.
- Dummett, M. 1991: *The Logical Basis of Metaphysics*. Cambridge, MA: Harvard University Press.
- Dummett, M. 1993: *The Seas of Language*. Oxford: Oxford University Press.
- Evans, G. 1976: Semantic structure and logical form. In Evans and McDowell (eds), 1976, 199–222. Reprinted in Evans, 1985, 49–75. Also reprinted in Ludlow (ed.), 1997.
- Evans, G. 1981a: 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 Evans, 1985, 322–42.
- Evans, G. 1981b: Understanding demonstratives. In H. Parret and J. Bouveresse (eds), *Meaning and Understanding*, Berlin: W. de Gruyter. Reprinted in Evans, 1985, 291–321. Also reprinted in Ludlow (ed.), 1997.
- Evans, G. 1982: *The Varieties of Reference*. Oxford: Oxford University Press.
- Evans, G. 1985: *Collected Papers*. Oxford: Oxford University Press.

- Evans, G. and McDowell, J. (eds) 1976: *Truth and Meaning: Essays in Semantics*. Oxford: Oxford University Press.
- Field, H. 1986: The deflationary conception of truth. In G. Macdonald and C. Wright (eds), *Fact, Science and Morality: Essays on A.J. Ayer's Language, Truth and Logic*. Oxford: Blackwell, 55–117.
- Field, H. 1994: Deflationist views of meaning and content. *Mind*, 103, 249–85.
- Forbes, G. 1987: Indexicals and intensionality: A Fregean perspective. *Philosophical Review*, 96, 3–31.
- Forbes, G. 1989: *Languages of Possibility: An Essay in Philosophical Logic*. Oxford: Basil Blackwell.
- 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. Also reprinted in Ludlow (ed.), 1997; in Martinich (ed.), 2001; and in Moore (ed.), 1993.
- Frege, G. 1918: The thought: A logical inquiry. In P.F. Strawson (ed.), *Philosophical Logic*, Oxford: Oxford University Press, 1967, 17–38. Reprinted in Ludlow (ed.), 1997, 9–30.
- Goldman, A.I. 1971: The individuation of action. *Journal of Philosophy*, 68, 761–74.
- Grandy, R. 1973: Reference, meaning and belief. *Journal of Philosophy*, 70, 439–52.
- Grice, H.P. 1957: Meaning. *Philosophical Review*, 66, 377–88. Reprinted in Grice, 1989, 213–23. Also reprinted in P.F. Strawson (ed.), *Philosophical Logic*, Oxford: Oxford University Press, 1967, 39–48; and in Martinich (ed.), 2001.
- Grice, H.P. 1961: The causal theory of perception. *Proceedings of the Aristotelian Society, Supplementary Volume*, 35, 121–52. Reprinted in Grice, 1989, 224–47. Also reprinted in G. Warnock (ed.), *The Philosophy of Perception*, Oxford: Oxford University Press, 1967.
- Grice, H.P. 1968: Utterer's meaning, sentence-meaning, and word-meaning. *Foundations of Language*, 4, 1–18. Reprinted in Grice, 1989, 117–37. Also reprinted in J. Searle (ed.), *Philosophy of Language*, Oxford: Oxford University Press, 1971.
- Grice, H.P. 1969: Utterer's meaning and intentions. *Philosophical Review*, 78, 147–77. Reprinted in Grice, 1989, 86–116. Also reprinted in Ludlow (ed.), 1997.
- Grice, H.P. 1975: Logic and conversation. In P. Cole and J. Morgan (eds), *Syntax and Semantics, Volume 3: Speech Acts*, New York: Academic Press, 41–58. Reprinted in Grice, 1989, 22–40. Also reprinted in Davis (ed.), 1991; and in Martinich (ed.), 2001.
- Grice, H.P. 1978: Further notes on logic and conversation. In P. Cole (ed.), *Syntax and Semantics 9: Pragmatics*, New York: Academic Press, 113–28. Reprinted in Grice, 1989, 41–57.
- Grice, H.P. 1989: *Studies in the Way of Words*. Cambridge MA: Harvard University Press.
- Hale, B. and Wright, C. (eds) 1997: *A Companion to the Philosophy of Language*. Oxford: Blackwell.
- Higginbotham, J. 1985: On semantics. *Linguistic Inquiry*, 16, 547–93.
- 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. 1988: Is semantics necessary? *Proceedings of the Aristotelian Society*, 88, 219–41.
- Higginbotham, J. 1989a: Elucidations of meaning. *Linguistics and Philosophy*, 12, 465–518. Abridged version reprinted in Ludlow (ed.), 1997, 157–78.

- Higginbotham, J. 1989b: Knowledge of reference. In A. George (ed.), *Reflections on Chomsky*, Oxford: Blackwell, 153–74.
- Horwich, P. 1990: *Truth*. Oxford: Blackwell (Second Edition, Oxford: Oxford University Press, 1998).
- Horwich, P. 1998: *Meaning*. Oxford: Oxford University Press.
- Johnston, M. 1988: The end of the theory of meaning. *Mind and Language*, 3, 28–42.
- Kaplan, D. 1989: Demonstratives. In J. Almog, J. Perry and H. Wettstein (eds), *Themes from Kaplan*, Oxford: Oxford University Press, 481–566.
- Kim, J. 1976: Events as property-exemplifications. In M. Brand and D. Walton (eds), *Action Theory*, Dordrecht: Reidel, 159–78.
- Kripke, S. 1980: *Naming and Necessity*. Oxford: Blackwell. Excerpts reprinted in Ludlow (ed.), 1997; and in Martinich (ed.), 2001.
- Larson, R. and Segal, G. 1995: *Knowledge of Meaning: An Introduction to Semantic Theory*. Cambridge, MA: MIT Press. Excerpt reprinted in Ludlow (ed.), 1997, 179–99.
- Laurence, S. 1996: A Chomskian alternative to convention-based semantics. *Mind*, 105, 269–301.
- Laurence, S. 1998: Convention-based semantics and the development of language. In P. Carruthers and J. Boucher (eds), *Language and Thought: Interdisciplinary Themes*, Cambridge: Cambridge University Press, 201–17.
- Lewis, D. 1969: *Convention*. Cambridge, MA: Harvard University Press.
- Lewis, D. 1975: Languages and Language. In K. Gunderson (ed.), *Language, Mind and Knowledge*, Minneapolis: University of Minnesota Press. Reprinted in Lewis, 1983, 163–88. Also reprinted in Martinich (ed.), 2001.
- Lewis, D. 1983: *Philosophical Papers: Volume 1*. Oxford: Oxford University Press.
- Loar, B. 1976: Two theories of meaning. In Evans and McDowell (eds), 1976, 138–61
- Loar, B. 1981: *Mind and Meaning*. Cambridge: Cambridge University Press.
- Ludlow, P. (ed.) 1997: *Readings in the Philosophy of Language*. Cambridge, MA: MIT Press.
- Lycan, W.G. 1991: Review of Avramides, *Meaning and Mind: An Examination of a Gricean Account of Language*. *Mind and Language*, 6, 83–6.
- Lycan, W.G. 2000: *Philosophy of Language: A Contemporary Introduction*. London: Routledge.
- McDowell, J. 1976: Truth conditions, bivalence, and verificationism. In Evans and McDowell, 1976, 42–66.
- Martinich, A.P. (ed.) 2001: *The Philosophy of Language* (Fourth Edition). Oxford: Oxford University Press.
- Moore, A.W. (ed.), 1993: *Meaning and Reference*. Oxford: Oxford University Press.
- Neale, S. 1995: The philosophical significance of Gödel's slingshot. *Mind*, 104, 761–825.
- Peacocke, C. 1976: Truth definitions and actual languages. In Evans and McDowell (eds), 1976, 162–88
- Peacocke, C. 1981: Demonstrative thought and psychological explanation. *Synthese*, 49, 187–217.
- Peacocke, C. 1983: *Sense and Content*. Oxford: Oxford University Press.
- Peacocke, C. 1986: *Thoughts: An Essay on Content*, Oxford: Blackwell.
- Peacocke, C. 1989: When is a grammar psychologically real? In A. George (ed.), *Reflections on Chomsky*, Oxford: Blackwell, 111–30.
- Peacocke, C. 1992: *A Study of Concepts*. Cambridge, MA: MIT Press.

- Perry, J. 1977: Frege on demonstratives. *Philosophical Review*, 86, 474–97. Reprinted in Perry, 1993. Also reprinted in Ludlow (ed.), 1997.
- Perry, J. 1979: The problem of the essential indexical. *Noûs*, 13, 3–21. Reprinted in Perry, 1993.
- Perry, J. 1993: *The Problem of the Essential Indexical and Other Essays*. Oxford: Oxford University Press.
- Platts, M. 1979: *Ways of Meaning: An Introduction to a Philosophy of Language*. London: Routledge and Kegan Paul (Second Edition, Cambridge, MA: MIT Press, 1997).
- Quine, W.V.O. 1960: *Word and Object*. Cambridge, MA: MIT Press.
- Sainsbury, R.M. 1991: *Logical Forms: An Introduction to Philosophical Logic*. Oxford: Blackwell (Second Edition, 2001).
- 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.
- Searle, J.R. 1979: Metaphor. In A. Ortony (ed.), *Metaphor and Thought*, Cambridge: Cambridge University Press, 92–123. Reprinted in Davis (ed.), 1991, 519–39.
- Sperber, D. and Wilson, D. 1981: Irony and the use-mention distinction. In P. Cole (ed.), *Radical Pragmatics*, London: Academic Press, 295–318. Reprinted in Davis (ed.), 1991, 550–63.
- Sperber D. and Wilson, D. 1986a: Loose talk. *Proceedings of the Aristotelian Society*, 86, 153–71. Reprinted in Davis (ed.), 1991, 540–49.
- Sperber, D. and Wilson, D. 1986b: Pragmatics and modularity. In A.M. Farley, P.T. Farley and K-E McCullough (eds), *The Chicago Linguistic Society Parasession on Pragmatics and Grammatical Theory*. Reprinted in Davis (ed.), 1991, 583–94.
- Sperber, D. and Wilson, D. 1995: *Relevance: Communication and Cognition*. Oxford: Blackwell (Second Edition).
- Stoljar, D. 2001: The deflationary theory of truth. In E. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Stanford, CA: The Metaphysics Research Lab at the Center for the Study of Language and Information, Stanford University, URL = <http://plato.stanford/archives/spr2001/entries/truth-deflationary/>.
- Strawson, P.F. 1952: *Introduction to Logical Theory*. London: Methuen.
- Strawson, P.F. 1959: *Individuals: An Essay in Descriptive Metaphysics*. London: Methuen.
- Strawson, P.F. 1964: Intention and convention in speech acts. *Philosophical Review*, 73, 439–60. Reprinted in Strawson, 1971, 149–69. Also reprinted in Davis (ed.), 1991.
- Strawson, P.F. 1970a: The asymmetry of subjects and predicates. In H.E. Kiefer and M.K. Munitz (eds), *Language, Belief and Metaphysics*, New York: State University of New York Press. Reprinted in Strawson, 1971, 96–115.
- Strawson, P.F. 1970b: Meaning and truth. Inaugural Lecture at the University of Oxford. Reprinted in Strawson 1971, 170–89. Also reprinted in Martinich (ed.), 2001.
- Strawson, P.F. 1971: *Logico-Linguistic Papers*. London: Methuen.
- Strawson, P.F. 1974a: *Subject and Predicate in Logic and Grammar*. London: Methuen.
- Strawson, P.F. 1974b: On understanding the structure of one's language. In *Freedom and Resentment and Other Essays*, London: Methuen. Reprinted in Evans and McDowell (eds), 1976, 189–98.

- Tarski, A. 1944: The semantic conception of truth and the foundations of semantics. *Philosophy and Phenomenological Research*, 4, 341–75. Reprinted in Martinich (ed.), 2001, 69–91.
- Tarski, A. 1956: The concept of truth in formalized languages. In A. Tarski, *Logic, Semantics, Metamathematics*, Oxford: Oxford University Press, 152–278.
- Taylor, B. 1976: States of affairs. In Evans and McDowell (eds), 1976, 263–84.
- Taylor, B. 1982: On the need for a meaning-theory in a theory of meaning. *Mind*, 91, 183–200.
- Taylor, B. 1985: *Modes of Occurrence: Verbs, Adverbs and Events*. Oxford: Blackwell.
- Taylor, K. 1998: *Truth and Meaning: An Introduction to the Philosophy of Language*. Oxford: Blackwell.
- Wiggins, D. 1976: The *de re* 'must': A note on the logical form of essentialist claims. In Evans and McDowell (eds), 1976, 285–312.
- Wiggins, D. 1980: What would be a substantial theory of truth? In Z. van Straaten (ed.), *Philosophical Subjects*, Oxford: Oxford University Press, 189–221.
- Wiggins, D. 1986: Verbs and adverbs, and some other modes of grammatical combination. *Proceedings of the Aristotelian Society*, 86, 273–304.
- Wiggins, D. 1997: Meaning and truth conditions: From Frege's grand design to Davidson's. In Hale and Wright (eds), 1997, 3–28.
- Wilson, D. and Sperber, D. 1981: On Grice's theory of conversation. In P. Werth (ed.), *Conversation and Discourse*, London: Croom Helm, 155–78.
- Wittgenstein, L. 1953: *Philosophical Investigations*. Oxford: Blackwell.
- Wright, C. 1986: Theories of Meaning and Speakers' Knowledge. In C. Wright, *Realism, Meaning and Truth*, Oxford: Blackwell, 204–238.