Contents

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.

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 META-PHYSICS (chapter 2) we move around this triangle (figure 3.1).

Thus, for example, we might try to give a philosophical account of some distinctions in reality – say, between OBJECTS AND PROPERTIES (pp. 726–7), or between particulars and universals – in terms of differences among 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

(2) 'if S knows *p*, then *it is impossible* for him to be wrong about *p*'. The mistake turns on thinking that the correct wide scope reading (1) of 'it is impossible' licenses the narrow scope reading (2) which constitutes infallibilism.

An infallibilist account makes the definition of knowledge look simple: S knows *p* if his belief in it is infallibly justified. But this definition renders the notion of knowledge too restrictive, for it says that S can justifiably believe *p* only when the possibility of *p*'s falsity is excluded. Yet it appears to be a commonplace of epistemic experience that one can have the very best evidence for believing something and yet be wrong (as the account of scepticism given below is at pains to show), which is to say that fallibilism seems the only account of justification adequate to the facts of epistemic life. We need therefore to see whether fallibilist theories of justification can give us an adequate account of knowledge.

The problem for fallibilist accounts is precisely the one illustrated by the Rolls Royce example above, and others similar to it (so-called 'Gettier examples', introduced in Gettier 1963), namely, that one's justification for believing *p* does not connect with the truth of *p* in the right way, and perhaps not at all. What is required is an account that will suitably connect S's justification both with his belief that *p* and with *p*'s truth.

What is needed is a clear picture of 'justified belief'. If one can identify what justifies a belief, one has gone all or most of the way to saying what justification is; and en route one will have displayed the right connection between justification, on the one hand, and belief and truth on the other. In this connection there are several standard species of theory.

Foundationalism

One class of theories of justification employs the metaphor of an edifice. Most of our ordinary beliefs require support from others; we justify a given belief by appealing to another or others on which it rests. But if the chain of justifying beliefs were to regress without terminating in a belief that is in some way independently secure, thereby providing a foundation for the others, we would seem to lack justification for any belief in the chain. It appears necessary therefore that there should be beliefs which do not need justification, or which are in some way self-justifying, to serve as an epistemic underpinning.

On this view a justified belief is one which either is, or is supported by, a foundational belief. The next steps therefore are to make clear the notion of a 'foundation' and to explain how foundational beliefs 'support' non-foundational ones. Some way of understanding foundationalism without reliance on constructional metaphors is needed.

It is not enough barely to state that a foundational belief is a belief that requires no justification, for there must be a reason why this is the case. What makes a belief independent or self-standing in the required way? It is standardly claimed that such beliefs justify themselves, or are self-evident, or are indefeasible or incorrigible. These are not the same things. A belief might be self-justifying without being self-evident (it might take hard work to see that it justifies itself). Indefeasibility means that no further evidence or other, competing, beliefs, can render a given belief insecure. Yet this is a property that the belief might have independently of whether or not it is self-justifying. And so on. But what these characterizations are intended to convey is the idea that a certain immunity from doubt, error or revision attaches to the beliefs in question.

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 (1956) Convention T 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 redescription of utterances of sentences by a member 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.

obliges us to find an alternative. If an X cannot be found to support a bridge across the sceptical gap, the option is to try closing it – or more accurately, to show that there is no gap at all. Both the quest for X and the closing of the gap have constituted major epistemological endeavours against scepticism in modern philosophy. Some of these endeavours, in brief, are as follows.

Descartes's immediate successors were, as mentioned, unpersuaded by his attempt to bridge the gap by invoking a good divinity to serve as X. LOCKE (chapter 29), without much fanfare, employed a weaker version of the Cartesian expedient by saying that we can ignore sceptical threats to the causal story because 'the light that is set up in us shines bright enough for all our purposes'. From Locke's point of view it does not matter whether the inner light is set up by God or nature; the point is that there is something – X, the inner light which could be, perhaps, reason, empirical intuition or native trust in the reliability of the senses – that gives us grounds for accepting our ordinary knowledge-acquiring means as adequate.

Others, not content with such unsatisfactory moves, look for X elsewhere, and claim to find it in some version of foundationalism, the thought – sketched above – that our epistemic system has a basis in special beliefs that are in some way self-justifying or self-evident and which, in conjunction with the evidence we ordinarily employ in making knowledge claims, secures them against scepticism. As we saw earlier, a chief ground for rejecting such theories is alleged to be that none of them identifies satisfactory candidates for 'foundations'. But one stimulating way of making something like a foundationalist case is offered by Kant, whose attempt prompted others.

2.7 Transcendental arguments

KANT (1724–1804) (chapter 32) regarded failure to refute scepticism as a 'scandal' to philosophy, and offered his *Critique of Pure Reason* (1929) as a solution. His thesis is that our minds are so constituted that they impose a framework of interpretative concepts upon our sensory input, among them those of the objectivity and causal interconnectedness of what we perceive. Application of these concepts transforms mere passive receipt of sensory data into EXPERIENCE (pp. 726–33) properly so called. Our faculties are such that when raw data comes under the interpreting activity of our concepts, they have already had spatial and temporal form conferred on them by the nature of our sensory capacities; all our experience, considered as relating to what is outside us, is experience of a spatially structured world, and all our experience, considered as relating to its received character in our minds, is of a temporally structured world. Upon the spatio-temporal data thus brought before our minds we impose the categories, that is, the concepts that make experience possible by giving it its determinate character. And here is Kant's point: if the sceptic asks us to justify our claims to knowledge, we do so by setting out these facts about how experience is constituted.

Kant claimed HUME (1711–76) (chapter 31) as his inspiration for these ideas, because Hume had argued that although we cannot refute scepticism – reason was not, he claimed, up to the task – we should not be troubled, for human nature is so constituted that we simply cannot help having the beliefs that scepticism challenges us to justify. Those beliefs include, for example, that there is an external world, that causal relations hold between events in the world, that inductive reasoning is reliable, and so

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: 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 counter-examples 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: 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