

Persons and their Underpinnings

MARTIN DAVIES

Abstract: I defend a conception of the relationship between the personal and subpersonal levels as *interaction without reduction*. There are downward inferences from the personal to the subpersonal level but we find upward explanatory gaps when we try to construct illuminating accounts of personal-level conditions using just subpersonal-level notions. This conception faces several serious challenges but the objection that I consider in this paper says that, when theories support downward inferences from the personal to the subpersonal level, this is the product of an unacceptably mechanistic view of persons. According to this objection, if we were to focus on persons as conscious rational thinkers and agents then the support for putative downward inferences would be undermined. I consider and reject developments of this objection in response to two arguments for downward inferences.

Positions in the philosophy of mind and language vary along many dimensions. Towards one end of one of these dimensions is a position according to which the only questions about mind and language that are susceptible of rational investigation are questions that belong to the sciences of mind and language including cognitive science and cognitive neuroscience.¹ The occupants of this position accept that the sciences of cognition, like any other sciences, include analytical and methodological issues; but they maintain that these issues are interwoven with the more overtly empirical questions that the science investigates, and do not provide the subject matter for a distinct discipline. According to this position, the business of the philosophy of mind and language is simply to hand all the substantive questions over to cognitive science.

Towards the opposite end of this dimension is a position that maintains that the philosophy of mind and language offers a distinctive methodology – conceptual analysis, broadly conceived – for investigating a class of substantive questions about the notions that figure in our everyday thinking about mental and linguistic matters. The occupants of this position allow that cognitive science is all very interesting in its way; but they insist that it has little or nothing to contribute to the philosopher's project of plotting the contours of our conceptual scheme. Those who occupy this second position might call the first *scientism*, or specifically *cognitive scientism*. Those who occupy that first position might call the second *philosophical isolationism*.

This account of the extreme positions on a scientism–isolationism dimension is crude and is not intended as a characterisation of the claims and programmes of any real theorists. But these extreme positions concerning the relationship between the philosophy of mind and language and cognitive science go naturally with correspondingly extreme positions about the relationship between two classes of descriptions. On the one hand, there are our everyday descriptions of ourselves as conscious and self-conscious, thinking, feeling, acting, talking persons. On the other hand, there are the kinds of descriptions that figure in scientific theories about human beings; in particular, there are the descriptions that are used in neuroscience and in information-processing psychology.

The two classes of descriptions are sometimes said to belong at the *personal* level and the *subpersonal* level respectively. The term 'the subpersonal level' is potentially misleading in a number of ways. For one thing, it would be wrong to suppose that there is

¹ There are important questions about the relationship between information-processing psychology and neuroscience, but these are not addressed in this paper. (See Stone and Davies, 1999.) In general, when I speak of cognitive science I intend to include cognitive neuroscience.

just one level of description that is relevant to information-processing psychology. For another thing, it would be wrong to suppose that information-processing psychology is exclusively concerned with units that are smaller than a whole person; the distinction between levels of description is not fundamentally a distinction between levels of aggregation. Instead of ‘the subpersonal level’, it would probably be better to talk of *system* or *machine* levels. But having noted these reservations, I shall continue to use the familiar term.

1. Interaction Without Reduction

Cognitive scientism and philosophical isolationism concern the relationship between the philosophy of mind and language, on the one hand, and cognitive science, on the other. According to cognitive scientism, all the substantive questions about mind and language – questions about consciousness, thought, agency and meaning – will be answered by the sciences of cognition. So descriptions of ourselves as conscious, thinking, acting, talking persons must, to the extent that they can figure in substantive questions, fall within range of the explanations offered by those sciences. What goes along naturally with cognitive scientism is, then, a view of the relationship between the personal and subpersonal levels as one of *reduction*.² According to this view of the inter-level relation, the only descriptions of persons that have any point from the perspective of rational enquiry are descriptions that are cast in the terms of the sciences of cognition.

According to philosophical isolationism, there are substantive questions about mind and language that have to be answered by the distinctive methods of philosophical investigation. When it comes to answering these questions, cognitive science has little or nothing to contribute. Going along naturally with philosophical isolationism is a view of the relationship between the personal and subpersonal levels as one of *independence*. According to this view of the inter-level relation, the correctness of the descriptions of persons that are of primary interest to philosophy is not answerable to empirical discoveries about the anatomy of the brain or about how information-processing systems work. This view might also be said to be a variety of *anti-psychologism*.

I do not accept either of these extreme views of the relationship between the personal and the subpersonal levels. As against both endpoints on the spectrum of possible views, I favour an intermediate conception of the inter-level relation as *interaction without reduction*. According to this conception, there are downward inferences, of a relatively a priori kind, from the personal to the subpersonal level. Philosophical theorising itself reveals that personal-level conditions (such as consciousness, thought, agency or understanding) impose subpersonal-level requirements. (This is what I mean by inter-level ‘interaction’.) But we find upward explanatory gaps when we try to construct illuminating accounts of these personal-level conditions using just subpersonal-level notions such as those that figure in neuroscience or information-processing psychology. (This is why it is ‘without reduction’.)³

² The term ‘reduction’ is used in more than one way. For present purposes, reductionism is primarily a doctrine about explanation or elucidation. It might seem that cognitive scientism, as I have characterised it, definitely requires a reductive conception of personal-level descriptions. I have spoken more cautiously of ‘what goes along naturally with’ because it is imaginable that there might be a science of human beings that uses descriptions that are not reducible to the subpersonal-level descriptions that figure in information-processing psychology and cognitive neuroscience.

³ The idea of an upward explanatory gap is familiar from discussions of the personal-level condition of phenomenal consciousness. The interaction-without-reduction view is not intended to rule out the

Though this conception of the inter-level relation strikes me as attractive, it certainly faces serious challenges. One of these is posed by *the problem of armchair knowledge*. We know some personal-level facts about ourselves in a special first-personal way. This first-personal knowledge is a kind of armchair knowledge; its status as knowledge does not depend on my carrying out any empirical investigation of the world around me or of the machinery within my skull. For example, I can know, without rising from the armchair, that I am currently thinking about interaction without reduction. But a priori knowledge of conditional statements that are underwritten by philosophical theory is also a kind of armchair knowledge. So, without rising from the armchair, I may know that I meet some personal-level condition, and also know that if I meet that personal-level condition then the information processing that is going on inside my head must meet certain requirements. But then, still without stirring myself from the armchair, I can put the two pieces of knowledge together and draw the obvious conclusion about subpersonal-level information processing. Yet, we do not ordinarily reckon that knowledge about how the information-processing system in my head works should be available just by introspecting and giving thought. Rather, this kind of knowledge is delivered by scientific investigation involving careful experiments and observations. Cognitive science involves getting out of the armchair. The advocate of the interaction-without-reduction picture needs to have an answer to this puzzle, but I shall not be offering any solution here.⁴ Instead, I shall be concerned with a different kind of objection to the picture.

The imagined objector says that, although some philosophical theories may support downward inferences from the personal to the subpersonal level, this is the product of an unacceptably mechanistic view of persons. It is no surprise, this critic says, that we can derive subpersonal-level requirements for personal-level conditions if we regard persons merely as ‘natural objects’, as Michael Dummett puts it (1991, p. 92). If we were to focus instead on persons as conscious rational thinkers and agents⁵ then the support for putative downward inferences would be undermined. According to this critic, interaction without reduction is not really a stable combination. Rather, inter-level interaction goes along with a depersonalising reductionism; and avoiding reductionism involves espousing inter-level independence.

It is, of course, difficult to evaluate this line of objection until we see it developed in response to some specific claims of inter-level interaction. But there is already something that can be said to explain the critic’s basic strategy. Suppose that we start out by insisting on a view of persons that would make the anti-reductionist idea of upward explanatory gaps plausible. Persons enjoy conscious experience and are capable of conscious thought. Their thinking, both theoretical and practical, is subject to norms and they arrive at judgements and decisions on the basis of reasons. Also, let us add, there is a style of explanation of these judgements and decisions in terms of reasons that is *sui generis*; such explanations are not to be assimilated to explanations that proceed by subsuming events under covering laws about how the world happens to work.⁶ (Whether

possibility that there are many descriptions of persons that present no obstacle to explanation in subpersonal-level terms.

⁴ For attempts at a solution, see Davies, 1998a, 1998b, forthcoming.

⁵ This is what Dummett himself recommends. See below, Section 7.

⁶ Thus, in an oft-quoted passage, McDowell says (1985, p. 389): ‘[T]he concepts of the propositional attitudes have their proper home in explanations of a special sort: explanations in which things are made intelligible by being revealed to be, or to approximate to being, as they rationally ought to be. This is to be

or not this claim about the distinctiveness of rationalising explanation is ultimately correct, it is dialectically useful to bring it in at this point so as to highlight, or exaggerate, the difference between the personal and the subpersonal levels.) Given this rich and distinctive way of talking about persons it may be – so the critic suggests – that such requirements as there are on personal-level conditions can be seen to be met already at the personal level. There is *no need to descend* to the subpersonal level of neuroscience and information-processing psychology.

In this paper, I shall consider developments of the ‘no need to descend’ objection in response to two examples of downward inferences from the personal to the subpersonal level. The first example concerns the underpinnings of inferential transitions in thought; the second focuses on the requirements for linguistic understanding.

2. First Example: Inference and the Language of Thought Hypothesis

Our first example of a downward inference is provided by an argument that I have given elsewhere and will briefly review here. It is an argument from a personal-level premise about thinkers performing inferential transitions in virtue of their form to a subpersonal-level conclusion about cognitive machinery, namely, the truth of Jerry Fodor’s language of thought hypothesis.⁷

The starting point for the argument is the assumption that occurrent judgments and thinkings are constituted or underpinned by tokenings of physical configurations. These configurations belong to kinds that figure in the science of information-processing psychology. This assumption is what Fodor (1987, p. 135) calls *intentional realism*. It is opposed to the view that talk about mental states and mental acts is merely a device for generating predictions about behaviour and also to the view that statements about a person’s mental states and acts are made true by the existence of patterns in that person’s behaviour.

Starting from this assumption that there is quite a close relationship between everyday mental events and events that would be of interest to cognitive science, the argument proceeds in two steps. The first step says that the idea of a thinker performing inferential transitions in virtue of their form is to be cashed out, at least in part, in terms of tacit knowledge of a rule of inference.⁸ The second step then establishes the completely general point that tacit knowledge of a rule requires syntactic articulation in the physical configurations that are the inputs to the transitions that are described by the rule. Putting the two steps together, we arrive at the conclusion that the physical configurations corresponding to occurrent thoughts that figure in inferences have syntactic structure. Thus, an internal language of thought (syntactic structure in physical configurations) is a subpersonal-level requirement for the personal-level condition of performing inferential transitions in virtue of their form.

contrasted with a style of explanation in which one makes things intelligible by representing their coming into being as a particular instance of how things generally tend to happen.’

⁷ See Davies (1991, 1992) for this argument. In fact, what is argued for is that conceptualised thought requires the truth of the language of thought hypothesis. A background assumption is that possessing concepts involves commitments to certain patterns of inference (Evans, 1982; Peacocke, 1992). It is proposed that conceptualised thought involves performing certain inferences in virtue of their form. Thereafter, the argument runs in the way outlined in the text. For Fodor’s own arguments in support of the language of thought hypothesis, see Fodor (1975, 1987).

⁸ This first step is a kind of inference to the best philosophical explanation. For the notion of tacit knowledge that is in play here, see e.g. Davies (1987).

There are many things that might be said about this argument. Someone might say that the argument pays insufficient attention to Wittgenstein's remark that peripheral order might proceed out of central chaos⁹ and that it is wrong to slide from the fact that personal-level psychological descriptions have a certain structure to an assumption that there is a matching structure in the brain. But, this is not a telling objection since, quite apart from the fact that the language of thought hypothesis is not about the brain as such, there is no gratuitous slide to an unwarranted assumption. What is offered is, after all, an *argument*. But it should be acknowledged that the argument will not move anyone who does not even accept its starting point, the idea of tokenings of physical configurations that underpin mental acts of judgement or thought.¹⁰

Our focus, though, will be on a particular reply to the argument, a reply that develops the 'no need to descend' objection. The reply starts from the point that persons are conscious rational thinkers. Epistemological considerations are advanced to support the view that a conscious rational thinker is consciously aware, specifically, of the *forms* of inferential transitions. Then it is claimed that, even granting intentional realism, the argument for the existence of syntactic structure encoding form at the subpersonal level is undermined by the fact that there is already conscious awareness of form at the personal level. If the form of the inference is already transparent to the thinker then why, the critic asks, does it need to be encoded in the physical configurations *as well*?

In the face of this reply we could adopt either of two strategies for defending our first example of a downward inference from the personal to the subpersonal level. We could reject the idea that a personal-level requirement of *awareness of logical form* emerges from consideration of what it takes to be a conscious rational thinker. Alternatively, we could accept, or even embrace, the idea about awareness of logical form but maintain that even so the critic's objection misses its mark. I shall be adopting that second strategy.

3. Conscious Reasoning and Awareness of Form

Let us consider a thinker, Bruce, who believes that *A or B* and also believes that *not-A*. It is likely that Bruce also believes that *B*, or will come to believe it if the question whether *B* arises. Under what conditions can we see this transition in thought as a rational one? Two answers suggest themselves immediately.

One is that Bruce's first two beliefs should actually constitute a reason for believing the third thing. Call this: the reason condition. We would show that the reason condition is met by pointing out that the first two propositions entail the third; the argument with the first two beliefs as premises and the third belief as conclusion instantiates a valid form. More generally, we show that the reason condition is met, and thereby show how a transition in thought at least could be a rational one, by conducting an investigation with an abstract subject matter. We plot the contours of the abstract space of reasons.

An investigation of the abstract space of reasons reveals that, if a thinker such as Bruce believes that *A or B* and also that *not-A*, then the right thing for him to think in addition is that *B*. Those first two things add up to a reason for someone to think that *B*. However, if they are to be Bruce's reason, then something more must be true. Bruce's

⁹ Zettel (1981), p. 106.

¹⁰ This starting point is close to the assumption of propositional modularity that is discussed by Ramsey, Stich and Garon (1990). In my view, we are committed to this assumption by some of our everyday practices of mental talk and explanation. Spelling out the nature of this commitment would, of course, involve downward inferences of just the kind that are under discussion in the present paper.

believing, or coming to believe, the first two things must cause him to believe the third thing. Call this: the causal condition.¹¹

The reason condition and the causal condition do not exhaust the requirements for Bruce's transition in thought to be a rational one. The problem, of course, is to connect the two conditions in the right way, so that, intuitively and roughly, it is because believing that *A or B* and that *not-A* is a reason to believe that *B* that Bruce's believing those two things causes him to believe the third thing.

Any solution to the problem of how to connect reason and causation would have to meet two constraints. The first constraint is that it should not require Bruce to use *((A or B) and not-A) entails B* as an additional premise in his reasoning; such a requirement would be regressive. The second constraint is that a solution should not require that Bruce can conceptualise, or spell out, or theorise about the valid form that his inference instantiates. For Bruce can be a rational thinker without being a logic student. Needless to say, no general solution to the problem of connecting reason and causation is going to be offered here. But it is plausible that one idea figuring in a solution will be that Bruce should perform the transition in thought *because it is of that form*.

The two constraints limit the ways in which this idea can be developed. Performing the transition because it is of that form cannot, in general, be a matter of invoking an additional premise. Nor can performing the transition because it is of that form require Bruce to give a theoretical account of the form of inference in question. According to the first step in the argument reviewed in the previous section, it is required that Bruce should have tacit knowledge of the rule of inference. In principle, a person can have tacit knowledge of a rule licensing transitions that instantiate a certain pattern, yet lack any awareness that those transitions do instantiate that common pattern. This is as it should be for the general cognitive scientific notion of tacit knowledge. But, if the idea of performing an inference because it is of a certain form is to help us connect causation and reason, then a total lack of awareness of the form is not what is wanted.

What I shall consider, in the remainder of this section, is the natural and plausible view that when Bruce performs his inferential transition, he should, in some way, be aware of the form of the inference. One way to motivate this view is by appealing to epistemological considerations, since one aspect of being a conscious rational thinker is being a knowledgeable reasoner. If Bruce's transition in thought is to be a rational one then it should be the kind of transition that could, for example, yield Bruce knowledge that *B* if he started out from knowledge that *A or B* and that *not-A*? The suggested view about awareness of logical form is intended to be in the spirit of recent epistemological work by Bill Brewer, who remarks (1995, p. 242):

[I]f my following the argument is really to extend my knowledge, then my understanding of it must give me some appreciation of why I am right in believing its conclusion. I have to have some grip on how I thereby know the conclusion, and my belief should be guided by this understanding. A disposition to take beliefs on board in parallel with the steps of the argument is on its own insufficient for the argument to provide me with genuine knowledge.

¹¹ However distinctive rationalising explanation may be, the correctness of a rationalising explanation imposes some conditions on the causal history of the judgement or decision that is explained; the explanation provides information about that causal history. So, rationalising explanation is still a kind of causal explanation (Lewis, 1986).

Brewer's general project is to defend an epistemological position that he calls *non-reflective internalism*. The idea of epistemic internalism is that a knower should be in a position to know how he is right. So epistemic internalism is opposed to pure reliabilism about knowledge. It might be thought that knowing how one is right must be a matter of knowing an overarching justificatory principle. Reflective (or second-order) internalism construes the internalist requirement in this way. But reflective internalism is potentially regressive; and, in the case of knowledge by deductive inference, the regress that would be involved in offering a justificatory principle as an additional premise is all too obvious (hence the first constraint two paragraphs back). Non-reflective (or first-order) internalism says that the internalist requirement can be met in some other way.

It is a far from trivial matter to provide a non-reflective internalist account of knowledge in any particular domain and Brewer does not offer a detailed account of knowledge by inference. But in the case of perceptual knowledge he argues that external objects are presented in experience as being spatially, and therefore epistemically, accessible to the subject (1996, p. 269):

[P]erceptual experience constitutes a *reason* for the beliefs a person correctly forms on its basis, by providing him directly with knowledge of how he is thereby right about the way things are in the world around him. This is not dependent upon his knowledge of any general rule to the effect that certain appearances . . . are reliable indicators of certain worldly states of affairs, which is supposed to serve as the major premise of an *argument* in defence of these beliefs. It is rather that the egocentric spatial content of perception displays the things around him in such a way as to make evident the epistemic accessibility in that very experience of the way they are.

The idea here is that a proper account of perceptual experience itself will make clear how the internalist requirement for perceptual knowledge can be met without reflective (or second-order) knowledge of a justificatory principle.¹² In a similar way, the suggestion that Bruce should be aware of his thoughts, and the transition between them, as instantiating a particular form of argument (here, disjunctive syllogism) is intended as a contribution towards an adequate account of knowledge by deductive inference. It is intended to help us avoid, on the one hand, pure reliabilism and, on the other hand, the potentially regressive requirement that the reasoner should know an overarching justificatory principle.

Suppose that we accept that there is a personal-level requirement of awareness of logical form. We have regarded this as being in addition to the subpersonal-level requirement of the truth of the language of thought hypothesis. But our imagined critic takes the argument in favour of the personal-level requirement to illustrate the idea that, if we focus on persons as conscious and rational rather than merely as natural objects, then putative arguments for downward inferences will be undermined. According to the critic, once we adopt a sufficiently rich way of talking about persons as such we can see

¹² Brewer discusses perceptual knowledge in much greater detail in *Perception and Reason* (1999; see especially Chapter 6). There the project is to account for the fact that 'perceptual experiences provide reasons for empirical beliefs' (1999, p. 18). Both pure reliabilist and 'second-order' accounts of perceptual knowledge are rejected, and the possibility of a 'first-order' account is argued to turn on the fact that (p. 204): '[An experiencing subject] understands that his current apprehension that things are thus and so is in part due to the very fact that they are. His grasping the content that that is thus is in part due to the fact that that *is* thus. He therefore recognizes the relevant content *as* his apprehension of the facts, his *epistemic openness* to the way things mind-independently are out there.'

that there is *no need to descend* to the subpersonal level in order to meet requirements for personal-level conditions. The putative argument for a subpersonal-level requirement of syntactic structure in physical configurations is undermined once we accept that the form of the inference must be transparent to the thinker.

But I claim that the critic's objection misses its mark. The argument for the downward inference begins from a premise about thinkers performing inferential transitions because of their form or, as we put it in Section 2, in virtue of their form. The notion of 'in virtue of its form' is, at least in part, causal: the form of an inference is to figure in the causal explanation of a thinker's performing the inferential transition in thought.¹³ The two steps of the argument cash this causal condition out in terms of tacit knowledge and then make use of a general connection between tacit knowledge and syntactic structure. The view that has been motivated by epistemological considerations is that the thinker should be aware of the form of the inference. This is certainly a personal-level requirement. But clearly, a thinker's being aware of the form of an inference that he or she performs is not yet sufficient for the form's figuring in the causal explanation of the thinker's performing that inference. So the epistemologically motivated view that a thinker should be aware of the form of his or her inferences cannot undermine the argument that there are subpersonal-level requirements for the personal-level condition of performing inferential transitions in virtue of their form.¹⁴

Our imagined critic began by being suspicious of the conception of the relationship between the personal and subpersonal levels as one of interaction without reduction, downward inferences but upward explanatory gaps. According to the critic, the downward inferences depend on a pernicious reductionism. If the vocabulary for talking about conscious rational persons is enriched in line with the idea of upward explanatory gaps then the arguments for downward inferences are undermined. If this critic were right, then it would be easier to maintain the conception of the inter-level relation as one of independence and correspondingly easier to maintain the isolationist position that the findings of information-processing psychology have no special relevance for the philosopher of mind and language. In this section, we have developed, and then rejected, the critic's line of objection applied to our first example of inter-level interaction, which concerned the underpinnings of inferential transitions in thought. Next we turn to our second example which is about the requirements for linguistic understanding.

4. Second Example: Meaning and Tacit Knowledge

The notion of tacit knowledge of a rule of inference figured in the first step of the argument in our first example (Section 2) and tacit knowledge has certainly been a key concept in the development of cognitive science especially, of course, within theoretical

¹³ The form is to figure in the causal explanation even though the thinker may not be able to conceptualise, or spell out, or theorise about the form, and *a fortiori* may not offer an explicit account of the form as part of his or her reason for making the transition in thought.

¹⁴ The subpersonal-level requirement was not primarily motivated by epistemological considerations, but the presence of a mechanism that systematically mediates inferential transitions of a particular valid form could figure in a reliabilist account of why inferences of that form are knowledge yielding. Such a mechanism could ground '[a] disposition to take beliefs on board in parallel with the steps of the argument' (Brewer, 1995, p. 242). However, this 'is on its own insufficient for the argument to provide me with genuine knowledge' (ibid.). Something more is needed, 'some conscious understanding of why one is right in one's conclusions' (ibid.). But, quite apart from the independent motivation for the subpersonal-level requirement, it is not plausible that this something more does away with the need for a reliable mechanism.

linguistics. Someone who favours inter-level independence or anti-psychologism might allow that the notion of tacit knowledge is coherent and useful for empirical science, particularly for the science of language. But, according to that view of the relationship between the personal and subpersonal levels, empirical discoveries about tacit knowledge have no special relevance to the personal-level descriptions offered by the philosopher of language.¹⁵ The argument that we shall now consider is intended to show that this claim about inter-level independence is wrong. There are quite specific connections, underwritten by relatively a priori considerations of philosophical theory, between knowledge of language and tacit knowledge of semantic rules.

Ordinary speakers of English are credited with speaking a language in which a sentence that everyone takes to mean one thing may really mean another thing, and in which sentences that no one uses to mean anything nevertheless have determinate meanings. Why are they credited with this?

Ordinary speakers of English who use the sentence ‘No head injury is too trivial to be ignored’ typically use it to communicate the message that no head injury is trivial *enough* to be ignored (Bennett, 1976, p. 17). In fact, there may well have been a time (before the sentence was introduced as a philosophical example) when all those who used the sentence used it in that way. What the sentence is typically used to communicate is not what it means in English. So why do we say that ordinary speakers are speakers of English rather than of English’, an abstract language in which ‘No head injury is too trivial to be ignored’ is assigned the meaning that ordinary speakers take it to have? Call this: the problem of *meaning despite use*.

Each ordinary speaker of English uses only a finite number of sentences in a lifetime. Many sentences of English have never been used and will never be used by any speaker. So why do we say that ordinary speakers are speakers of English rather than of English*, an abstract language that simply omits the sentences of English that no one will ever use? Why do we say that they speak English rather than some other extension of English*? Call this: the problem of *meaning without use*.

As we address these two problems, a plausible starting point is that we could not justify those attributions to ordinary speakers if they merely had phrasebook (whole sentence by whole sentence) mastery of a large collection of English sentences (Schiffer, 1993, p. 237). If this is correct then our attributions depend on the fact that ordinary speakers have a structured rather than unstructured mastery of the sentences that they do use. Ordinary speakers may treat a few sentences as unstructured units, but this cannot be the general pattern. Ordinary speakers must have a structured mastery of the sentences that they use, but this cannot be a matter of knowing (in the everyday sense of conscious propositional knowledge) the semantic rules of English. Some linguists or semantic theorists may know those rules, but ordinary speakers surely lack such knowledge. According to the downward inference that is to be our second example, structured

¹⁵ An empirical discovery about tacit knowledge is, of course, potentially evidentially relevant to the empirical attribution to a person of knowledge of language. We must allow this on pain of illegitimately circumscribing the domain of evidence that is potentially relevant to an empirical hypothesis. See Fodor (1981, p. 199): ‘[A]ny facts about the use of language, or about how it is learned, or about the neurology of speaker/hearers, or, for that matter, about the weather on Mars, could, in principle, be relevant to the choice between competing linguistic theories.’ But someone might say – in the spirit of independence – that discoveries about tacit knowledge are potentially relevant to personal-level attributions of knowledge of language only in the way that any discovery – perhaps a discovery about the orbits of Jupiter’s moons – might, in principle, be relevant.

mastery is a matter of speakers' judgements about the meanings of sentences that they do use having a certain kind of explanation in terms of tacit knowledge of semantic rules.

The questions about meaning despite use and meaning without use are familiar from discussions of Paul Grice's (1989) programme for analysing the notion of literal meaning in terms of speaker's meaning plus convention. The proposal that the only way for a Gricean to deal with the meaning-without-use problem is by an appeal to the notion of tacit knowledge is a relatively familiar one. Thus, for example, Brian 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'. If this is taken as a conclusion merely about Grice's programme then the result is bound to have limited interest, for it is now widely accepted that Grice's bold analytical programme cannot be carried through (Schiffer, 1987, Chapter 9). But the general line of thought is not specific to Grice's programme, and the problems of meaning without use and meaning despite use need to be addressed in any case. Stephen Schiffer (1993; responding to Lewis, 1992) discusses the first of these problems – the meaning-without-use problem – and reaches three conclusions about it.

Schiffer's first conclusion is that a strictly Gricean, or intention-based semantics, solution to the problem is unsatisfactory because it is too intellectualist and so imposes conditions that are not really necessary for using a language (Schiffer, 1993, pp. 241–2; see also 1987). The second conclusion is that it would also be a mistake to impose a Chomskyan condition of the kind that Loar appears to suggest. The reason Schiffer gives is that an internally represented grammar is not required for structured mastery of a language (p. 243). A Chomskyan solution, like a Gricean solution, imposes a condition that is not really necessary.

But it is Schiffer's third conclusion that provides the basis for our downward inference (*ibid.*, p. 244):

Our discussion has suggested that there must be something involved in the processing of L utterances that somehow determines L [rather than any of the infinitely many other languages that include the used fragment]. . . . I think it is right that an account of the actual-language relation should appeal to mechanisms of language processing in order to nail down the unused parts of used languages, but it must suitably leave open the nature of those mechanisms.

If Schiffer is right then a solution to the meaning-without-use problem needs to appeal to empirical facts about information processing. There are subpersonal-level requirements for the personal-level condition of knowing and speaking English.

It might be thought that there is an alternative way of determining English as the actual language of ordinary speakers, rather than just the used fragment English* or any of the infinitely many other languages that agree with English on the meanings of the sentences in English*. A semantic theorist can fit a set of semantic rules to the finite set of sentences in English* and then use those rules to extrapolate to the meanings of the sentences that are never used. In essence, this is the strategy of fitting a curve to a finite set of data points. The theorist who is engaged in curve fitting extrapolates and interpolates and may even undertake a little curve smoothing as well. So it may seem that curve fitting offers a solution to the problem of meaning despite use as well as to the problem of meaning without use, and does so without any speculation about subpersonal-level mechanisms.

However, the plausible starting point for our reflection on the two problems yields a decisive objection to the curve-fitting alternative. The curve-fitting strategy focuses on

sentences rather than on speakers; indeed, a semantic theorist can fit a set of semantic rules to the sentences of a purely formal language. So the strategy would be applicable even if the speakers under investigation had mere phrasebook mastery of the sentences that they happened to use. But our starting point was the thought that our description of ordinary speakers as speaking English, rather than English* or English', would not be justified if they had only phrasebook mastery. If this is correct then something more than curve fitting is required and Schiffer's proposal seems compelling. What is needed is that the curve-fitting procedures of extrapolation, interpolation and smoothing should follow the contours of the processing that actually goes on in speakers' heads.

Schiffer's argument supports an inference from the personal-level notion of knowledge of language to subpersonal-level 'mechanisms of language processing', but he does not explicitly commit himself to subpersonal-level tacit knowledge of a semantic theory. The more cautious result is already enough for our purposes. We can consider how the 'no need to descend' objection would be developed in response to this example of a downward inference from the personal to the subpersonal level. But, in fact, it is legitimate for us to go beyond Schiffer's explicit conclusion and claim that a solution to the meaning-without-use and meaning-despite-use problems does require an appeal to tacit knowledge of semantic rules.

Schiffer rejects any commitment to an internally represented grammar. (This is his second conclusion.) But it is plausible that the idea of internal representation that is rejected goes beyond what would be minimally required for tacit knowledge of semantic rules. It is certainly correct that any account of tacit knowledge must allow for the possibility of a difference between tacit knowledge of one set of rules and tacit knowledge of another set of rules even though the two sets of rules determine the same input-output relation in extension. The idea of tacit knowledge as explicit internal representation, with the rules written out in an inner code, meets this requirement. But a thinner notion of tacit knowledge, cast in terms of causal-explanatory structure, can meet it as well (Evans, 1981; Davies, 1987). Very roughly, the idea would be that internal processing machinery might have a causal-explanatory structure that mirrors the logical-derivational structure of one set of rules but not of another set of rules even though the two sets determine the same relation in extension. Given this thinner notion of tacit knowledge, it is very plausible that if the structure of language processing is adequate to provide solutions to the meaning-without-use and meaning-despite-use problems, then the processing machinery embodies tacit knowledge of semantic rules. These rules would determine meanings for whole sentences, including sentences that are never used, by assigning semantic properties to words and to ways of putting words together: they would be the rules of a compositional semantic theory.

Schiffer himself offers a putative example of internal machinery that would be suitable for language processing yet without involving internal representation of semantic rules. In his example, the language-processing machinery embodies an internally represented *translation manual* from English to Mentalese, a Fodorian language of thought. However, it is not clear that an internally represented translation manual would really be adequate by itself to support the information processing that would have to be involved in understanding a language. In general, an internal representation of a rule, or any other embodiment of tacit knowledge of a rule, supports transitions between inputs and outputs that are themselves representations. For example, an internal representation of the rule linking displacement and elapsed time for bodies falling under gravity would support transitions between inputs that are representations of times and outputs that are

representations of displacements. So we would expect that an internal representation of a translation rule linking English and Mentalese expressions would support transitions between inputs that are representations of English expressions and outputs that are representations of Mentalese expressions. But, understanding a language involves hearing sentences and having thoughts and, if the language of thought hypothesis is correct, then having thoughts involves tokening sentences of Mentalese, not just tokening representations of sentences of Mentalese. So what is needed for understanding a language is machinery that supports transitions from inputs that are representations of English expressions to outputs that are Mentalese expressions; that is, to outputs that are representations of worldly objects, properties, and states of affairs. This is just the kind of machinery that is apt to embody tacit knowledge of semantic rules for English, for these are rules that link English expressions with worldly entities.

Internally represented translation rules are not sufficient, by themselves, for understanding language; they need to be augmented in some way. The resulting processing system might not internally represent semantic rules in the sense of explicitly encoding them. But it does seem inevitable that it will embody at least the thinner kind of tacit knowledge of semantic rules.¹⁶ The issues here deserve a much more detailed discussion. But it seems reasonable to go beyond what Schiffer explicitly says and to claim that a solution to the meaning-without-use and meaning-despite-use problems requires an appeal, not just to the mechanisms of language processing, but also to tacit knowledge of compositional semantic theories.¹⁷

In response to Schiffer's actual argument or the slight elaboration of it that we have just proposed, our imagined critic develops the 'no need to descend' objection. The critic argues, on epistemological grounds, for a personal-level requirement of *awareness of linguistic structure* and claims that this undermines the argument for a downward inference. The critic urges that, if the structure of a sentence is already available to the language user in conscious awareness, then that structure does not need to be encoded in information-processing machinery or represented in a tacitly known semantic theory *as well*. As in the case of our first example, we can respond to the critic either by rejecting the personal-level requirement or by accepting it and maintaining that even so the objection misses its mark. As before, I shall be adopting the second strategy.

5. Knowledge of Language and Awareness of Structure

Let us consider an example of understanding a sentence and let us stress, even labour, the parallels with our discussion of the example of deductive reasoning.¹⁸ Suppose that Narelle knows that s_1 means that p_1, \dots , and that s_n means that p_n , and judges that s means that p , where s is a hitherto unconsidered sentence built up from words and ways of putting words together that already occur in s_1, \dots, s_n . Under what conditions can we

¹⁶ See Peacocke (1986) for a related argument.

¹⁷ I am grateful to Brian Loar and Stephen Schiffer for helpful discussion. Schiffer has suggested that the difference between us over tacit knowledge is mainly verbal.

¹⁸ In order to stress the parallels, we focus on knowledge of meaning that is based on an inference from premises about what other sentences mean. This is not the typical case of linguistic understanding; usually a language user simply hears a sentence with its meaning. A non-reflective internalist epistemology would suggest a condition of awareness of structure in the case of this perceptual model of understanding too; and once again, the epistemologically motivated suggestion would not undermine arguments for subpersonal-level requirements.

see this transition as a rational one that could yield Narelle knowledge about the meaning of s ?

One condition – the reason condition – is that Narelle's n premises should constitute a reason for her to judge that s means that p . It should be possible to proceed by rational inductive means from those premises to that conclusion. As in the case of conscious reasoning, we would show that the reason condition is met by conducting an investigation with an abstract subject matter. It would be an exercise in formal semantics.¹⁹

As before, the reason condition is not, by itself, sufficient to guarantee that Narelle's actual transition in thought is a rational, and so knowledge yielding, one. There is also a causal condition; and, again as before, the reason condition and the causal condition have to be connected in the right way. The problem of how exactly this connection is to be made will not be solved here; but one constraint on a solution is that it should not require Narelle to have any explicit conceptualised mastery of a semantic theory for her language. For Narelle can be a knowledgeable language user without being a philosopher or linguist. The natural suggestion to make – following once again the lines of Brewer's (1995, 1996) non-reflective internalism in epistemology – is that, in some way, Narelle should nevertheless be aware of s 's structure, of how s is built up from words and constructions occurring in other sentences whose meanings she knows. In this way, Narelle could have some appreciation of how it is that she is able to know the meaning of a sentence that she has never come across before.

If we follow this suggestion about awareness of linguistic structure, then we shall arrive at a personal-level requirement on conscious, rational, knowledgeable language users. Our imagined critic proposes to use this line of thought to undermine our argument for a downward inference from knowledge of language to tacit knowledge of semantic rules (or, indeed, to any requirement on the structure of the mechanisms of language processing). According to the critic, if the structure of the sentence is already available to the language user in conscious awareness then there is *no need to descend* to the subpersonal level of tacit knowledge and processing machinery. But once again, I claim that the critic's objection misses its mark.

Subpersonal-level states of tacit knowledge of semantic rules are invoked to help solve the meaning-without-use problem. It is because of the presence of this tacit knowledge in Narelle, for example, that the sentence s means that p in her language whether or not she ever considers it, hears it or uses it. If she judges that s means that p – a piece of linguistic performance that is not guaranteed by the presence of that tacit knowledge – then she judges correctly. If the personal-level requirement of awareness of linguistic structure is to undermine the argument for tacit knowledge of a compositional semantic theory then it must provide an alternative solution to the meaning-without-use problem.

We have already seen that the strategy of mere curve fitting is not adequate to solve the problem. The reason is that curve fitting would equally lead to the conclusion that s means that p in Narelle's language even if s remained unused and Narelle (and all the other members of her community) had only phrasebook mastery of the sentences that were used (Schiffer, 1993, p. 237). Appeal to awareness of linguistic structure is

¹⁹ In particular, we could show that the condition is met by exhibiting a semantic theory in which the axioms that are drawn on in the canonical proofs of the meaning-specifying theorems for s_1, \dots, s_n are already sufficient to yield a canonical proof of a meaning-specifying theorem for s , to the effect that s means that p . For these purposes, semantic theorising is a relatively a priori project; it is a project of rational reconstruction (Wright, 1986).

inadequate for essentially the same reason: awareness of linguistic structure is compatible with mere phrasebook mastery of sentences. This point is not immediately obvious because the idea of phrasebook mastery of sentences is most vivid in the case where the language user is utterly blind to the structure that a semantic theorist would uncover. But someone with phrasebook mastery need not be completely unaware that a sentence is built out of familiar words in familiar ways. It is possible to be aware of sentences as being structured even though one's knowledge of those sentences' meanings is strictly a whole sentence by whole sentence matter. The structure of which a language user is aware might not figure in the explanation of that language user's judgements as to what sentences mean.

It is consistent with the personal-level requirement of awareness of linguistic structure that Narelle should have *this* kind of phrasebook mastery of the sentences that she does use, including s_1, \dots, s_n , and that she should never use the sentence s . Indeed, it is consistent with the requirement of awareness of linguistic structure that every member of Narelle's language community should rely on phrasebook mastery and that no one should ever use the sentence s . The curve-fitting approach proclaims that, in such a case, s would mean that p in Narelle's language. But according to the plausible starting point for our reflections in Section 5 this is incorrect.

The personal-level requirement of awareness of structure does not distinguish between cases in which the unused sentence s has a determinate meaning in Narelle's language (because she has structured mastery) and cases in which it does not (because she has phrasebook mastery). So the requirement of awareness of structure does not provide an alternative solution to the meaning-without-use problem. The critic's objection misses its mark because the personal-level requirement of awareness of linguistic structure is not enough, by itself, to solve the meaning-without-use problem from which the argument for the downward inference began.

6. A Better Construal of Awareness of Structure

In response to our rejection of the critic's objection, someone might say that epistemological considerations would actually motivate a stronger requirement of awareness of structure – a requirement that is not compatible with any kind of phrasebook mastery of sentences. Let us consider that possibility.

The epistemological position that motivates the requirement of awareness of structure is opposed to pure reliabilism: reliable mechanisms are not enough for knowledge. So reliable mechanisms that draw on tacit knowledge of semantic rules in order to generate beliefs about meaning would not be sufficient for linguistic understanding. Awareness of structure is supposed to provide 'some conscious understanding of why one is right in one's conclusions' (Brewer, 1995, p. 242). But it is difficult to see how awareness of structure could perform this role if it were compatible with the absence of the reliable mechanisms. The causal mechanism that reliably generates correct beliefs and the awareness that renders the success intelligible need to be more closely integrated. So, it might with some justice be said, the structure of which the language user is aware needs to be the very structure – the structure with its semantic significance – in virtue of which the language user arrives at a judgement about meaning.

This does seem to be the kind of position that a not purely reliabilist epistemology would properly motivate, and it seems plausible as the beginning of an account of knowledge of meaning. But it does not offer the critic any prospect for an objection to putative downward inferences from the personal to the subpersonal level. The suggestion

we now have is that a language user should arrive at a judgement about the meaning of a sentence *in virtue of* the semantic structure that she is aware of the sentence as having, though without using conceptualised principles of semantic theory as premises in her conscious reasoning. This is very closely analogous to the personal-level premise of the argument for the language of thought hypothesis (Section 4). We considered a thinker who performs an inferential transition in virtue of its form and is aware of the logical form of the inference although he is not able to give any theoretical account of the form and does not use any overarching principle as an additional premise. The proposal in the first step of the argument was that the idea of performing an inferential transition in virtue of its form should be cashed out, in part, in terms of tacit knowledge of a rule of inference. In the case of linguistic understanding, the absolutely inevitable suggestion is that the truth of the ‘in virtue of’ claim once again requires the presence of subpersonal-level mechanisms that embody tacit knowledge of rules. So, on this better construal of what the personal-level requirement of awareness of linguistic structure comes to, a subpersonal-level requirement follows from it, rather than being supplanted by it.

7. Participant Understanding

We began the main line of argument in this paper by insisting on a conception of conscious rational persons that would make the idea of upward explanatory gaps plausible (Section 1). Our aim in doing so was to investigate whether this would permit developments of the ‘no need to descend’ objection that would undermine arguments for downward inferences (Sections 2 and 4) thus showing interaction without reduction to be an unstable combination. We have been led to consider some epistemologically motivated personal-level requirements of awareness of form and structure. These requirements are plausible and important, but they leave the arguments for downward inferences untouched (Sections 3 and 5) and actually provide premises for further arguments of the same kind (Section 6).

We introduced the anti-reductionist conception of conscious rational persons by way of a contrast with a mechanistic conception of persons as ‘natural objects’, in Dummett’s phrase. We have argued that the anti-reductionist conception is still consistent with the claim that there are downward inferences from the personal to the subpersonal level. But perhaps it will be said that we have not fully appreciated the contrast that Dummett intended. Here is the striking passage in which the ‘natural objects’ phrase occurs (Dummett, 1991, p. 92):

A meaning-theory should not, therefore, aspire to be a theory giving a *causal* account of linguistic utterances, in which human beings figure as natural objects, making and reacting to vocal sounds and marks on paper in accordance with certain natural laws. We have no need of such a theory. . . . [W]e describe the practice and the institutions that surround the practice, and then it becomes intelligible as an activity of rational agents. And that is all the understanding that we seek of language. . . . [I]f a causal theory were possible, it would not provide the sort of understanding that we seek.

What Dummett contrasts with a ‘causal account’, here, is the kind of description of a practice that can render it ‘intelligible as an activity of rational agents’. We seek this kind of description in the case of unfamiliar practices which, as theorists or commentators, we want to understand. Also, Dummett says, we ‘implicitly grasp’ an account of the same

kind in the case of practices that we understand and in which we ourselves participate as conscious rational agents.

If we consider Bruce's participant understanding of conscious reasoning or Narelle's participant understanding of knowledge of meaning then we notice two features. One is that Bruce regards *B* as the right thing to think and Narelle regards *p* as the correct meaning to attach to *s*. To the extent that Bruce and Narelle display an implicit grasp of principles, these seem to be normative principles of right reasoning and correct meaning, rather than empirical causal principles. The second feature of these cases is that Bruce and Narelle have some appreciation of why they are right; an appreciation afforded, not by empirical knowledge about the causal workings of a reliable mechanism at the subpersonal level, but by conscious awareness of the form of the inference and the structure of the sentence.

Suppose that we grant that Bruce's and Narelle's participant understanding involves implicit grasp of normative principles and conscious awareness of form and structure; it is not a matter of deploying a causal theory. According to Dummett, it is this same kind of understanding that we seek as philosophical theorists and commentators and to achieve this kind of understanding we have no need of a causal theory of the information processing that takes place in Bruce's or Narelle's head.

There is surely something right and important here. It is possible for us to share Bruce's and Narelle's participant understanding without engaging in any empirical investigation of their cognitive machinery. We can understand Bruce's reasoning and Narelle's judgement about meaning 'from the inside' (Heal, 1998) by engaging in mental simulation and identifying with Bruce and with Narelle in imagination (Davies and Stone, 1998).²⁰ But although it seems right to grant this possibility, it does not follow that shared participant understanding is 'all the understanding we seek'. Nor does it follow that causal theories are of no concern to philosophy; nor that the personal level is independent of the subpersonal level. In fact, it is quite consistent to combine a commitment to the importance of shared participant understanding with acceptance of subpersonal-level requirements for personal-level conditions.

Suppose that I engage in a mental simulation of Bruce. I re-enact his reasoning and conclude that *B* from the premises that *A or B* and that *not-A*. Within the scope of the simulation, I say nothing about cognitive machinery, but just offer those premises as reasons for believing that *B*. Then, stepping back from the simulation, I say that Bruce believes that *B* because he believes that *A or B* and that *not-A*; and I say that he performs the transition in thought in virtue of its form. We can allow that I find Bruce intelligible as a conscious rational person because I can engage in this shared participant understanding. But none of this makes it any less the case that, if those claims about Bruce's inference and beliefs are to be true, then certain causal conditions have to be met. The truth of those claims is not indifferent to the causal order exhibited by the subpersonal-level information processing going on in Bruce.

Suppose that I engage in a mental simulation of Narelle. I re-enact her thinking about the hitherto unconsidered sentence *s* and arrive at the judgement that *s* means that *p*. Within the scope of the simulation, I say nothing about language-processing mechanisms, but simply justify my judgement by adverting to the meanings of other sentences *s*₁, . . . , *s*_n. Stepping outside the simulation, I say that Narelle was right. Even before she

²⁰ Brewer (1995, pp. 247–8) makes the connection between non-reflective internalism and mental simulation.

considered *s* it meant that *p* in her language and if she had never considered *s* it would still have meant that *p*. But if that claim about meaning in Narelle's actual language is to be true then a causal condition has to be met; and this causal condition requires subpersonal-level tacit knowledge. I may also say that Narelle arrived at her judgement about the meaning of *s* in virtue of the semantic structure that she was aware of it as having; and that claim too depends for its truth on the causal order exhibited by Narelle's cognitive machinery.

My personal-level claims about Bruce and Narelle depend for their truth on subpersonal-level requirements being met. Yet mental simulation is supposed to provide me with knowledge about Bruce and Narelle without my having to conduct an empirical investigation to ensure that those requirements are in fact met. Similarly, of course, the truth of corresponding claims about my own beliefs, inferences and linguistic understanding would depend on facts about the information processing going on in me. Yet I can have first-personal knowledge of my own mental states without undertaking research in cognitive science. So if the possibility of self-knowledge and of knowing other minds 'from the inside' is to be preserved within the interaction-without-reduction picture then we shall have to allow that here, as elsewhere, knowledge may depend on empirical assumptions without being justificatorily based on empirical investigation.²¹

In the first six sections of this paper I argued that, even when we focus on persons as conscious and rational and take on board epistemologically motivated requirements on conscious awareness, still this does not block downward inferences from the personal to the subpersonal level. Now, in the final section, I have argued that anti-psychologism is not helped by our recognition of participant understanding that can be shared by theorists and commentators through mental simulation.

Philosophers and cognitive scientists no doubt vary in the degree of sympathy they have for the personal-level notions that have been in play here: rationalising explanation that is causal but not subsumptive, non-reflective internalism about knowledge, normative participant understanding, identification with another in imagination. It might be that these notions are most naturally associated with philosophers towards the isolationism end of the scientism-isolationism dimension that I mentioned at the beginning. But, if the argument of this paper is right, then even non-reductionist philosophers who invoke all of these notions should be ready to engage with the empirical study of mind and language and to adopt an interactive view of the relationship between the personal and the subpersonal levels of description.²²

²¹ Also, as we noted at the outset, the advocate of the interaction-without-reduction conception of the inter-level relation needs to have a solution to the problem of armchair knowledge.

²² Talks based on this material were given at the University of Adelaide, Carleton University and Cornell University, at workshops in Bonn and Turin, and at the 1998 conference of the European Society for Philosophy and Psychology in Lisbon. I am grateful to the audiences in all those places and to José Luis Bermúdez and Matthew Elton for their advice.

References

- Bennett, J. 1976: *Linguistic Behaviour*. Cambridge: Cambridge University Press.
- Brewer, B. 1995: Compulsion by reason. *Proceedings of the Aristotelian Society, Supplementary Volume* 64, 237–53.
- Brewer, B. 1996: Internalism and perceptual knowledge. *European Journal of Philosophy*, 4, 259–75.
- Brewer, B. 1999: *Perception and Reason*. Oxford: Oxford University Press.
- Davies, M. 1987: Tacit knowledge and semantic theory. Can a five per cent difference matter? *Mind*, 96, 441–62.
- Davies, M. 1991: Concepts, connectionism, and the language of thought. In W. Ramsey, S. Stich and D. Rumelhart (eds), *Philosophy and Connectionist Theory*. Hillsdale, NJ.: Lawrence Erlbaum Associates, 229–57.
- Davies, M. 1992: Aunty's own argument for the language of thought. In J. Ezquerro and J.M. Larrazabal (eds), *Cognition, Semantics and Philosophy: Proceedings of the First International Colloquium on Cognitive Science*. Dordrecht: Kluwer Academic Publishers, 235–71.
- Davies, M. 1998a: Language, thought, and the language of thought (Aunty's own argument revisited). In P. Carruthers and J. Boucher (eds), *Language and Thought*. Cambridge: Cambridge University Press, 226–47.
- Davies, M. 1998b: Externalism, architecturalism, and epistemic warrant. In C. Macdonald, B.C. Smith and C. Wright (eds), *Knowing Our Own Minds*. Oxford: Oxford University Press, 321–61.
- Davies, M. forthcoming: *Knowing What Thought Requires*. Oxford: Oxford University Press.
- Davies, M. and Stone, T. 1998: Folk psychology and mental simulation. In A. O'Hear (ed.), *Contemporary Issues in the Philosophy of Mind*: Royal Institute of Philosophy Supplement 42. Cambridge: Cambridge University Press, 53–82.
- Dummett, M. 1991: *The Logical Basis of Metaphysics*. Cambridge, MA.: Harvard University Press.
- Evans, G. 1981: Semantic theory and tacit knowledge. In S. Holtzman and C. Leich (eds), *Wittgenstein: To Follow a Rule*. London: Routledge and Kegan Paul, 118–37. Reprinted in *Collected Papers*. Oxford: Oxford University Press, 1985.
- Evans, G. 1982: *The Varieties of Reference*. Oxford: Oxford University Press.
- Fodor, J. 1975: *The Language of Thought*. New York: Crowell.
- Fodor, J. 1981: Some notes on what linguistics is about. In N. Block (ed.), *Readings in Philosophy of Psychology, Volume 2*. London: Methuen, 197–207.
- Fodor, J. 1987: *Psychosemantics*. Cambridge, MA.: MIT Press.
- Grice, H.P. 1989: *Studies in the Way of Words*. Cambridge MA.: Harvard University Press.
- Heal, J. 1998: Understanding other minds from the inside. In A. O'Hear (ed.), *Contemporary Issues in the Philosophy of Mind*: Royal Institute of Philosophy Supplement 42. Cambridge: Cambridge University Press, 83–99.
- Lewis, D. 1986: Causal explanation. In *Philosophical Papers, Volume 2*. Oxford: Oxford University Press, 214–40.
- Lewis, D. 1992: Meaning without use: Reply to Hawthorne. *Australasian Journal of Philosophy*, 70, 106–10.
- Loar, B. 1981: *Mind and Meaning*, Cambridge: Cambridge University Press.

- McDowell, J. 1985: Functionalism and anomalous monism. In E. LePore and B.P. McLaughlin (eds), *Actions and Events: Perspectives on the Philosophy of Donald Davidson*. Oxford: Basil Blackwell, 387–98.
- Peacocke, C. 1986: Explanation in computational psychology: Language, perception and level 1.5. *Mind and Language*, 1, 101–23.
- Peacocke, C. 1992: *A Study of Concepts*. Cambridge, MA.: MIT Press.
- Ramsey, W., Stich, S. and Garon, J. 1990: Connectionism, eliminativism and the future of folk psychology. In J.E. Tomberlin (ed.), *Philosophical Perspectives 4: Action Theory and Philosophy of Mind*. Atascadero, CA.: Ridgeview Publishing Company, 499–533.
- 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.
- Stone, T. and Davies, M. 1999: Autonomous psychology and the moderate neuron doctrine. *Behavioral and Brain Sciences*, 22, 849–50.
- Wittgenstein, L. 1981: *Zettel*. Oxford: Basil Blackwell.
- Wright, C. 1986: Theories of meaning and speakers' knowledge. In *Realism, Meaning and Truth*. Oxford: Basil Blackwell, 204–38.