

## **Philosophy of Mind: Consciousness, Intentionality and Ignorance\***

**Daniel Stoljar, ANU**

### **Introduction**

In a striking passage at the end of his classic paper “Epiphenomenal Qualia”, Frank Jackson observed that “it is not sufficiently appreciated that physicalism is an extremely optimistic view of our epistemic powers. If it is true, we have, in very broad outline admittedly, a grasp our place in the scheme of things. Certain matters of sheer complexity defeat us...but in principle we have it all” (1982, 135). Jackson went on to suggest that, from a perspective that emphasizes that we are organic beings with an evolutionary history and limited psychological capacities this fact about physicalism renders it quite implausible.

It seems to me that Jackson’s instincts on target here, at least as regards the sort of materialism<sup>1</sup> he mostly had in mind, that is, the sort promoted by such philosophers as J.J.C.Smart (1959) and David Lewis (1983, 1994). However, what Jackson does not say is that the same thing is true of dualism. For dualism (at least in the form opposed by such philosophers as Smart and Lewis) agrees with materialism about all aspects of reality with one exception, viz., the particular states of mind having to do with consciousness. Moreover, while these states are an exception to materialism, they are not an exception to an extremely optimistic view of our powers, since the dualist typically supposes that the subjects who are in these states know what they are in what Lewis later called an “uncommonly demanding sense” (see Lewis 1995). Hence, just as it is not sufficiently appreciated that materialism is an extremely optimistic view, it is likewise not sufficiently appreciated that dualism is too. Indeed both of the traditional positions in philosophy of mind—materialism and dualism—presuppose that “in principle we have it all”. If, following Jackson, we regard that presupposition as false (or is at least as incredibly unlikely) both positions should be rejected.

What happens to the philosophy of mind if we reject both of these standard positions? I think that doing so puts us in a position to formulate new solutions to

---

\* Thanks to the editors of this volume, Barry Dainton and Howard Robinson, for comments on a previous draft.

<sup>1</sup> I will use the phrase ‘materialism’ in this paper rather than ‘physicalism,’ which is used here by Jackson. Nothing turns on this, though the history of these words is of some interest. For discussion see ch.1 of Stoljar 2010.

some of the central question of philosophy of mind, including the problem of consciousness and the problem of intentionality. The solutions do not involve what some philosophers say they want, a theory or account of consciousness or intentionality. But I think it is on reflection a mistake to expect solutions of this sort in any case, at least at the present stage of knowledge; the precise account of what consciousness or intentionality is and how they fit into the world are problems that will be solved if at all, by total science, by the science whatever it is that will be reached in the (perhaps hypothetical) end of inquiry. In the meantime, however, what philosophers certainly can do is assess whether certain lines of reasoning concerning consciousness or intentionality that lead to philosophically objectionable (or at any rate striking) conclusions are persuasive or not. When I say that rejecting the standard positions provides us with a solution to these problems it is solutions of this latter sort I have in mind.

This paper traces out how rejecting the standard positions leads to solutions of this kind. After providing some context for the issues, I start with the problem of consciousness, and briefly review my own favored epistemic solution to that problem—a solution hinted at but not developed in the passage from Jackson, and which I have set out elsewhere (see Stoljar 2006, 2006a). Then, in the bulk of the paper, I will consider how to extend this epistemic solution to the problem of intentionality. As I will explain, this is a non-trivial matter because the problem of intentionality is distinct from the problem of consciousness in not involving counterparts of the arguments distinctive of consciousness, such as the knowledge argument and the conceivability argument, and it is most obviously arguments of this style are subject to an epistemic response. So to see how to connect issues of ignorance to intentionality we will need to formulate the problem of intentionality in a more explicit way than it is usually done. In the brief concluding section of the paper, I will make a remark about a third issue, the problem of self-knowledge, which is a problem that is different from the problem of consciousness and intentionality in, but which has emerged as a key problem, perhaps the key problem, in philosophy of mind in recent years. I will suggest that while rejecting the standard positions does not solve this problem, the materials we assemble when thinking through the perspective suggested by Jackson are nevertheless important when we turn to this problem.

### **Responding to Our Inner Cartesian**

One of the key events in philosophy of mind in the last hundred years—arguably, *the* key event—was the appearance in 1949 of Oxford philosopher Gilbert Ryle’s book *The Concept of Mind*. Ryle defined himself—and many philosophers of mind following him have likewise defined themselves—in opposition to someone called the ‘Cartesian’. The connection between Ryle’s Descartes and the real Descartes is tenuous, and in any case isn’t really to the point. Rather Ryle’s Descartes acts as a foil in philosophy of mind rather as the sceptic about the external world acts as a foil in epistemology. The point is not to *refute a real person*, i.e. someone who may or may not be the famous figure of 17<sup>th</sup> century philosophy and science. The point is rather to *critically engage with an intellectual tendency* which according to Ryle greatly influences our *own* contemporary philosophical interpretations of ordinary and scientific psychology, and which stands in the way of the attempt to provide a plausible general picture of human cognitive capacities in their relation to the rest of the world.

What exactly is this tendency? Ryle’s Descartes holds two key theses. First, he holds the *metaphysical* thesis that mental phenomena and physical phenomena are wholly distinct; that is, he is a traditional dualist. Second, he holds the *epistemological* thesis that mental phenomena are wholly transparent to themselves; that is, each of us is equipped with a faculty of introspection that, if used properly, will provide us in principle with complete and infallible knowledge of the contents of our minds. Throughout *Concept*, Ryle develops a barrage of techniques and suggestions designed to undercut both theses. Nowadays many of these techniques and suggestions are unpopular, but the project he initiated—the project of accommodating our inner Cartesian, as we might put it—is still with us. Indeed, it is no exaggeration to say that the dominant focus of philosophy of mind since Ryle has been on the *first* of the two theses just distinguished.

The obvious way to resist being a traditional dualist is to become a traditional materialist, and indeed, this is the lesson that many philosophers took from Ryle. But Ryle himself thought that the traditional materialist holds a position that is almost as bad as the traditional dualist, since both hold a presupposition that is false. According to Ryle (at least as I read him<sup>2</sup>) the false presupposition at issue is that ordinary

---

<sup>2</sup> In *Concept*, Ryle says that sentences such as 'John Doe knows French' are "neither reports of observed or observable states of affairs, nor yet reports of unobserved or unobservable states of

psychological declarative sentences are fact-stating just as ordinary physical declarative sentences are; that is, the sentence 'Otto is anxious' states a fact just as 'Otto is 6 foot 2 in his stockinged feet' states a fact. The dualist thinks that 'Otto is anxious' states a fact, namely one about a realm only contingently connected to the physical. The materialist thinks it states a fact about ordinary bodies. Ryle suggests that both sides are mistaken because 'Otto is anxious'—and psychological declaratives generally—states no fact at all. In that sense both standard materialism and standard dualism are to be rejected.

However, it is difficult to make sense of the idea that 'Otto is anxious' states no fact. For what does 'state no fact' mean? A very common suggestion is that a declarative sentence states no fact just in case an assertion of that sentence does not conventionally express a belief, and instead expresses a non-belief, say a desire or intention or command. So for example the sentence 'You will do your violin practice' might conventionally express a command, rather than a belief about the future. However, if this is what it is to be fact-stating, it is hard to believe that ordinary psychological sentences are not fact-stating; after all, does 'Otto is anxious' not conventionally express the belief that Otto is anxious? What would constitute evidence that it does not? And anyway, to say that a sentence is conventionally used to express beliefs (or is not conventionally used to express beliefs) is *itself* a statement of fact about psychology, so, if it is true, Ryle's position threatens to be self-undermining.

So in my view Ryle's suggestion about things not being fact-stating does not represent a productive way to think about these issues. On the other hand, I do think he is right that both standard materialism and standard dualism presuppose something false. For they both presuppose an overly optimistic view of our powers. They both presuppose as Jackson says that 'in principle we have it all'. But how does rejecting this presupposition allow us to respond to Ryle's Cartesian? In the next section I will consider how to respond to the part of the Cartesian picture that most people think is the hardest to deal with: the problem of consciousness.

## **The Problem of Consciousness**

---

affairs.'" (1949, p. 120). I take it that what Ryle meant by this is that expressions such as 'John Doe knows French', and presumably other psychological reports, are not in the fact stating business.

One might have thought that it would be an easy matter to reject the first thesis held by Ryle's Descartes, i.e. metaphysical dualism. Surely a small dose of scientifically informed common sense—according to which humans are evolved creatures as much a part of the natural order as zebra fish or xenopus toads—is sufficient to dispel the idea that each of us is a complex of a body and soul, a picture apparently belonging more the history of religion than to contemporary philosophical thought. As it turns out, however, things are not so simple. For as a number of philosophers of mind from the 1970s on argued (Kripke 1980, Nagel 1974, Jackson 1982, Robinson 1982 and Chalmers 1996) dualist modes of thought not only can be divorced from any religious element, but can be founded on extremely compelling and simple intuitions about consciousness and then developed with considerable clarity using machinery borrowed from modern modal logic, semantics and epistemology. To accommodate Cartesianism in this sophisticated modernized form requires us to ask some searching questions not only about our conceptions of nature and the mind but also about our conception of philosophical method.

One very plausible form of reasoning here has come to be called 'the knowledge argument'. This argument may be set out in various ways, but a simple version has it as proceeding from two main premises. The first premise concerns what it is possible for a person to know; in particular, it is possible for a person to know all the physical facts as well as every fact that follows a priori from the physical facts, and yet not know what it is like to have an experience of certain type. Jackson's (1982) Mary is the best known, but not the only, illustration of this possibility.<sup>3</sup> The second premise of the knowledge argument is that if this is possible then materialism is false. The conclusion is that materialism is false, or anyway it is false if there are facts about what it is like to have certain experiences and if people know these facts.

How to respond to this argument? There are a number of existing proposals in the literature. One response, the ability response, is based on a distinction that Ryle himself developed and defended in *Concept*, namely the distinction propositional knowledge, i.e. knowledge attributed by sentences of the form 'S knows *that* such and

---

<sup>3</sup> Here is Lewis's very vivid description of the case: "Mary, a brilliant scientist, has lived from birth in a cell where everything is black and white. (Even she herself is painted all over.) She views the world on black-and-white television. By television she reads books, she joins in discussion, she watches the results of experiments done under her direction. In this way she becomes the worlds leading expert on color and color vision and the brain states produced by exposure to colors. But she doesn't know what it is like to see color. And she never will, unless she escapes from her cell" (1988, 263).

such', and know how, i.e. knowledge attributed by sentences of the form 'S knows *how to* such and such'. Armed with this distinction proponents of the ability response say that while Mary learns something when she comes out of her room, she does not learn any propositional knowledge but rather gains a sort of know-how.<sup>4</sup>

But this proposal is problematic in two ways. First, it is not clear that there is any distinction here of the kind the ability response requires (see Stanley and Williamson 2001, Stanley 2011). Sentences of the form 'S knows how to such and such' seem to be semantically similar to sentences such as 'S knows where such and such is' or 'S know who such and such is'. But these sentences are usually thought of as attributing a sort of propositional knowledge; for example, when you know where something is, you know, for some suitably described place *t*, that the thing in question is at *t*. Likewise, if you know how to do something, it seems natural to say that you know, for some suitably described way of doing something *w*, that the thing in question is done in way *w*. If any analysis along these lines is correct, the ability hypothesis looks to be in serious trouble.<sup>5</sup>

The second problem is that, even if the know how/know that distinction is granted, the ability response relies on the idea that the experience Mary has when she comes out of her room and sees color for the first time is a novel experience. The general idea is that since she has not had the experience she does not have the relevant abilities or know-how; similarly, when she does have the experience she will gain the abilities or know-how. However, it is possible to develop the knowledge argument on the basis of examples that do not involve novel experiences (see Stoljar 2005, 2006). For such examples it is implausible that Mary gains a new ability—for she already has the ability in question (having already had the experience). If so, the ability hypothesis is not a good reply to the knowledge argument.

A different sort of response appeals to a distinction between two kinds of materialism—a priori and a posteriori materialism, as I will call them here.<sup>6</sup> Suppose that we have a sentence 'S' which somehow or other captures every physical fact of

---

<sup>4</sup> For the classic defence of the ability hypothesis see Lewis 1988, and the references therein.

<sup>5</sup> I do not say that this objection is conclusive; it depends if there is a fallback position for the ability response. For some discussion of this see Stoljar forthcoming

<sup>6</sup> In Chalmers 1996, these positions are referred to as 'Type-A materialism' (= a priori materialism) and 'Type-B materialism' (= a posteriori materialism).

the world; and suppose we have a second sentence ‘S\*’ which somehow or other captured every psychological fact. Materialism of any sort is committed to the view that the material conditional formed from these sentences—that is, ‘If S then S\*’—is necessary. The a priori materialist says that the conditional is necessary a priori; the a posteriori materialist says that the conditional is necessary a posteriori. Armed with this distinction, the a posteriori materialist response says that the knowledge argument shows at most that a priori materialism is false. But this leaves it open that a posteriori materialism is true.

But this response faces problems too. First, a number of philosophers insist that a posteriori materialism is not a possible position, however attractive it looks in the abstract. For these philosophers, there are theses in philosophy of language and epistemology that entail *if* materialism is true then a priori materialism is true (see, e.g., Lewis 1994, Jackson 1998, Chalmers 1996). These theses are contested of course, and we will not go in to them here (see, e.g., Stalnaker 2003). But the fact that many philosophers think that posteriori materialism is not simply false but impossible at least shows that appealing to the necessary a posteriori is no easy response to the knowledge argument.

The second problem is that the a posteriori materialism response to the knowledge argument even if correct is likely to win the battle but not the war. For the knowledge argument is just *one* argument against materialism about consciousness. Other arguments, such as the conceivability argument, are not going to be defeated by drawing a distinction between necessary a posteriori, since those arguments focus directly on the question of whether the connection between the mental and the physical is necessary—the question of whether the connection is in addition a posteriori or not is from this point a view a sideshow. (A quick argument for this conclusion is that Kripke himself advanced something very much like the conceivability argument in *Naming and Necessity*, but also provided (in the same work) the materials for formulating a posteriori materialism. Kripke himself evidently did not think that a posteriori materialism would be able to answer the conceivability argument, which is good evidence, but of course not conclusive evidence, that a posteriori materialism will not answer the conceivability argument.)<sup>7</sup>

---

<sup>7</sup> For my own development of the issues raised in this paragraph, see Stoljar 2006, 2006a.

### **The Epistemic Response to the Consciousness Problem**

If these common responses to the knowledge argument fail, how should one respond? In my view, the best response draws on the idea suggested in the passage from Jackson above.

Suppose that there is a type of non-experiential or physical fact of which we are ignorant but that is relevant to the nature of experience. (In other work, I have called this supposition ‘the ignorance hypothesis’.) Now take the claim that it is possible that someone knows all the physical facts and not all the facts. Such a claim involves the phrase ‘all the physical facts,’ which is what philosophers of language call ‘a quantifier phrase’. How should we interpret this phrase? Well if we suppose that the ignorance hypothesis is true, and there are physical facts of which we are ignorant but which are relevant, then we have two choices: to include them in the scope of the quantifier or not. Suppose first that they are included. Then it is not clear that we are entitled to assert the possibility claim. Is plausible to say that someone can know all the facts (including one’s of which we are ignorant) and yet not know all the facts? How can we assert that if we are ignorant of some of the relevant facts? Suppose now that they are not included. Now the possibility claim looks plausible; the physical facts that we currently know do indeed seem to be such that someone could know all of them and not know some phenomenal facts. But this possibility claim does not threaten materialism because at most it shows that facts about experience come apart from *some* physical facts not from them all. Putting this together, if the ignorance hypothesis is true, the knowledge argument is unpersuasive.

This response, if correct, will tell us that the knowledge argument is unpersuasive; to that extent the response is not committed to dualism. But one might reasonably ask why the position does not commit us to materialism, and if so, what happened to the point I mentioned before, that both standard materialism and standard dualism are to be rejected, just as Ryle thought (though not for Ryle’s reasons). Isn’t the epistemic response to the argument a straightforward response on behalf of the materialism?

The answer to this question is ‘no’, or better, ‘it depends on what you mean by materialism’. One way of spelling out materialism is to construe it as advancing a particular positive view about the physical world—e.g. that contemporary physics (or something very like it) gives a complete statement of the world and what it is like. That is the sort of materialism which Lewis, for example defends, the sort I have so

far called ‘standard materialism’. The epistemic response I have suggested is not available to a materialist of that sort, precisely because a materialist of that sort supposes that in principle we have it all. But one might use the term ‘materialism’ in a non-standard way, i.e. to mean any position that supposes that the knowledge argument (and similar arguments) are unsound and that in consequence they provide us with no reason to suppose that facts about consciousness are fundamental. Using the label in that way, it is possible to say that materialism and the epistemic view are compatible. But that does not mean the epistemic view does not involve a rejection of standard materialism.

Of course the epistemic approach to the knowledge argument faces a number of challenges. Some argue that we are not ignorant in the way that the response requires; others argue that even if we are, then this will not have the effect on the arguments that the approach assumes it will. I have responded to these problems in detail elsewhere, and will not go over them here (see Stoljar 2006). Instead, my aim in what follows is to consider whether the sort of response just sketched to the problem of consciousness may be applied to other problems in philosophy of mind.<sup>8</sup>

### **The Problem of the Problem of Intentionality**

Traditionally, philosophers have distinguished the problems presented for materialism by phenomenal consciousness from those presented by another aspect of mental states, what contemporary philosophers—roughly following the 19th century philosopher, Brentano—called their intentionality. The intentionality of a mental state is its aboutness. When I think of Vienna or I believe that the computer is on the desk or I fear that the planet will get hotter, and so on, I instantiate mental states which are in a hard to define sense about Vienna, or the computer on the desk or planet Earth. The idea is that mental states and events have a property rather like signs, sentences, and gestures; that is, they are about or represent things other than themselves.

Can we extend the epistemic response just suggested from the problem of consciousness to the problem of intentionality? Before we answer this question we need to be clear about what the problem of intentionality is. For, while it is certainly

---

<sup>8</sup> In Stoljar 2006 I was very wary of extending the epistemic response to consciousness to other problems. The present paper represents a slight (but only slight) softening of that position.

traditional to talk about the problem of intentionality, it is harder than you might think to formulate the problem in any precise way.

In the light of our discussion of the problem of consciousness, a natural first thought is that there is a problem of intentionality in just the same way as there is a problem of consciousness. So for example, just as we formulated a knowledge argument for consciousness, we might formulate a counterpart ‘knowledge argument’ for intentionality. The first premise of this argument would say that it is possible for a person to know all the physical facts as well as every fact that follows a priori from the physical facts, and yet not know what people believe, (e.g.) does not know that Hillary believes that Obama is president (assuming that Hillary does believe this). The second premise is that if this is possible then materialism is false. The conclusion is that materialism is false. If there were such a persuasive argument, we could formulate the problem of intentionality, on analogy with the problem of consciousness, as the problem of saying what if anything is wrong with this argument.

But the problem is that there is no intuitive foundation to this counterpart knowledge argument. The Mary case provides us with an initially plausible case of someone who knows all the physical facts and yet does not know what it is like to see colour. (That the case is not plausible on reflection does not mean it is not initially plausible.) But there seems to be no similarly plausible case of someone who knows all the physical facts and yet does not know that Hillary believes that Obama is president. Mary herself, for example, seems perfectly capable of knowing facts of this sort. If someone told her while in the room what Hillary believes, she might well come to know that she believes that Obama is president, and there seems no aspect of Hillary’s belief of which she is ignorant.

It might be thought that while Mary may be able to come to know what Hillary believes in this way, there still seems to be gap between a complete description of the world in basic physical terms, and the existence of Hillary’s belief. For example, if Mary were to read a description of Hillary’s brain in basic physical terms, would she then be able to work out – from this description alone – what Hillary believes? If the answer to this is ‘no’—as it seems to be—then why is there not a knowledge argument for intentionality after all?<sup>9</sup>

---

<sup>9</sup> Thanks to Barry Dainton for this objection.

However, this line of thought assumes that Mary's knowledge is limited to a certain sort of physical information about Hillary's brain, for example to cellular information. It is true that from this alone one cannot work out what Hillary believes. But Mary's knowledge is not limited in this way; indeed it is an important strength of the knowledge argument in its original form that Mary is permitted to know anything that she could come to know while in the room. So for example, suppose Mary has access not just to cellular information about Hillary but also to computational, functional, behavioural and environmental information. If so, it is hard to believe that she could not come to know that Hillary believes that Obama is president, even if, as the original argument alleges, she could not come to know what it's like to see color.

Alternatively, it might be thought that there is an important class of intentional states which Mary could *not* come to know, viz., those that are also states of consciousness associated with experiencing colour. On many views, after all, such states are *themselves* intentional states and so (one might think) it is quite an easy matter to produce a plausible case in which someone knows all the physical facts and yet does not know an intentional fact—perhaps the Mary case is precisely such a case.

However, while this might be true, it does not affect the basic issue. The problem of intentionality is supposed to concern intentional states as such, not merely intentional states which are also conscious states. Indeed, it is for this reason that it is natural to attempt to formulate the problem in terms of a standing belief such as the example of Hillary's belief. And if we focus on intentional states as such (i.e., those that are not also conscious states) then it seems there is no knowledge argument about intentionality.

If the problem of intentionality is not to be explained on a direct analogy with the problem of consciousness, perhaps it comes about from the idea that intentional properties are not fundamental? Here is Jerry Fodor forcefully giving voice to this idea:

I suppose that sooner or later the physicists will complete the catalogue they've been compiling of the ultimate and irreducible properties of things. When they do, the likes of spin, charm, and charge will perhaps appear upon their list. But aboutness surely won't; intentionality simply doesn't go that deep. It's hard to see, in face of this consideration, how one can be a Realist about intentionality without also being, to some

extent or other, a Reductionist. If the semantic and the intentional are real properties, it must be in virtue of their identity with (or maybe of their supervenience on?) properties that are themselves neither intentional nor semantic. If aboutness is real, it must be really something else. (1987, p. 97)

The most obvious thing Fodor is saying in this passage is that intentionality is not a fundamental feature of the world. However, while this is plausible, it is again not sufficient to generate a problem about intentionality. For the same thing can be said about almost everything. For example, take the Mariana Trench. When the physicists complete the catalogue they've been compiling of the ultimate and irreducible properties of things, being the Mariana Trench surely won't be on that list either. The Mariana Trench goes deep but not that deep. But nobody thinks that there is a philosophical problem about the Mariana Trench. Put differently: if there is a problem of intentionality at all, it had presumably have something to do with intentionality, but the fact that intentional properties are not fundamental is a fact shared by many things.<sup>10</sup>

At this point a third suggestion about what the problem of intentionality is naturally suggests itself. Intentionality is usually associated with a number of interesting logical features, and the problem of intentionality might be thought of as the problem of coming to grips with these features. Suppose for example, that Hilary believed, not that Obama is president, but that Odin is. In that case she would have a belief about Odin rather than a belief about Obama. But off-hand this is puzzling. To have a belief about Obama might be thought of involving a relation between Obama and the person who has the belief—Hillary in this case. But the example of Odin shows that this cannot be so, or at least cannot be so in general. Hillary cannot stand in any relation to Odin for the simple reason that Odin does not exist. How then can she believe that Odin is president?

However while it is certainly challenging to explain the sense in which Hillary can have a thought about Odin, it is doubtful that we can straightforwardly appeal to this to raise the problem of intentionality, at any rate not if we are out to formulate a problem that has something to do with the contrast between materialism and dualism.

---

<sup>10</sup> For some literature in which points like this are emphasized, see Stich 1992, and Tye 1992.

For the Odin problem has nothing to do with the truth or not of materialism! Suppose for example that dualism of some very straightforward kind is true, and that Hillary was a complex of a soul and a body—it would *still* be puzzling how something that exists can apparently stand in a relation to something that does not. Similarly, suppose that it were a fundamental fact of nature that Hillary believes that Odin is president—even so, it would remain the case that the fact in question cannot be analysed as relation between Hillary and Odin, for there is no such relation.

In summary, before we confront the problem of intentionality, we had better confront the problem of the problem of intentionality, i.e. the problem of saying what the problem of intentionality is. As we have seen, is not generated by a counterpart for intentionality of arguments like the knowledge argument—for there is no such argument. It is not generated by the fact that intentionality is not a fundamental feature of reality—for while it plausibly is not fundamental, this is not a feature of intentionality in particular. And it is not generated by the fact that intentionality exhibits interesting logical features—for while these features need explanation, doing so seems to have nothing to do with the distinction between dualism and materialism. What then *is* the problem?

### **The Descriptive/Foundational Distinction**

In my view, a good way to proceed here<sup>11</sup> is to start with a distinction that is well known in the philosophy of language but is often less explicitly drawn in the philosophy of mind. This is a distinction, in Robert Stalnaker's (1997) terms, between questions of descriptive, and questions of foundational, semantics. The descriptive semantic project, as Stalnaker describes it, concerns what language we as a community speak, or if we confront the issue in an individualist framework, what language a particular individual speaks. Suppose, to fix ideas, we focus on a particular person—Karl, for example, the hero of David Lewis's (1974) paper 'Radical Interpretation'. The descriptive semantic project with respect to Karl is then the project of saying what language it is that he speaks.

This question about Karl is an empirical question. It might be that he speaks some particular dialect or idiolect of English, or it might be that he speaks some

---

<sup>11</sup> A rather different way to proceed is to connect the intentionality issue more directly with consciousness, as suggested for example in Johnston 2007 and Pautz 2010. I will not try to engage directly with these interesting ideas in what follows.

particular dialect or idiolect of Urdu. (That the answer to this descriptive semantic question may seem obvious does not mean it is not empirical.) Likewise, it might be that he speaks a language in which names are semantically equivalent to definite descriptions or that he speaks a language in which names are not semantically equivalent to definite descriptions. (That the answer to this descriptive semantic question may seem unobvious does not mean it is not empirical.)

Now if the descriptive semantic project is empirical, to solve it we need to attend to various sources of evidence that we have about Karl. How he acts in particular circumstances is surely one good source of evidence, and as are the judgments he (and we) would make about when certain sentences are true and under what conditions. It might also be that other sorts of evidence—say about what sort of creature Karl is—are also relevant. Indeed, in principle anything at all can be evidence for the hypothesis that Karl speaks a particular language; all that is required is that the evidence, together with background assumptions, makes the hypothesis more probable than it would otherwise be.

In saying that the question of what language Karl speaks is an empirical question, I am not denying that there are a priori, or at least very general, constraints on what it is for Karl or anybody to speak a language, and that these general constraints will also factor into descriptive semantics. For example, one condition mentioned by Stalnaker is this: “if the semantics is correct, then speakers must know, at least for the most part, what according to the semantics they are saying” (1997,176). As I understand it a condition like this functions to narrow down the possibilities of what language it is that Karl is speaking. A priori, after all, the possibilities are endless and the evidence that we have about Karl will surely underdetermine which language he speaks. But the assumption that Karl must be assumed to know what he is saying, at least for the most part, serves to narrow down the possible languages which—we can reasonably suppose—Karl speaks.

Suppose now that after reviewing the evidence and the relevant general constraints we agree that Karl speaks a particular language—L17 as it might be. Then we face what Stalnaker calls foundational semantic questions; these are questions about “what the facts are that give expression their semantic values, and more generally, about what makes it the case that the language spoken by a particular individual or community is a language with a particular descriptive semantics” (1997,167). Concerning Karl, then, questions at the foundational level concern what

facts about him—e.g. what psychological, physical, or behavioural facts—explain that he speaks L17 rather than something else or nothing at all. Presumably for example there are facts about Karl that make it the case that L17 is a version of English rather than what we call Urdu (if it is). And presumably too there are facts about Karl that make it the case that L17 is a language that contains names not equivalent to any definite description (again: if it is). Questions at the foundational level ask what exactly these facts are.

### **From Language to Mind**

Now, as the reference to semantics makes clear, the descriptive/foundational distinction is in the first instance a distinction in the philosophy of language. But the problem of intentionality if it is anything is a problem in the philosophy of mind. So, to connect this distinction to the problem of the problem of intentionality we would need to transpose it from the key of language to the key of mind. How is this transposition to be achieved?

One proposal might be to suggest that Karl has a language of thought, and then to apply the distinction directly to Karl's language of thought. Descriptive semantics for the language of thought from this point of view proceeds just like descriptive semantics for English or Urdu. However, formulating the problem of intentionality this way seems to me a mistake. The problem is not that Karl does not have a language of thought. Arguments such as those given by Fodor (e.g. Fodor 1987) seem to me quite compelling; they provide good empirical evidence (though of course not conclusive evidence) that Karl has a language of thought, or at any rate would if (as we are assuming) he is a human being like the rest of us. The problem is rather that the hypothesis that Karl has the language of thought is a quite specific psychological hypothesis about him; much more specific, for example, than the hypothesis that he has a mind or is the subject of intentional states at all. On the face of it, it is possible (even if not actual) that Karl could have intentional states without having a language of thought. But if that is so, we do not want to explain what the problem of intentionality is in terms of the language of thought.

A better proposal is to draw a distinction between (as I will say) descriptive psychology and foundational psychology (i.e. on direct analogy with descriptive and

foundational semantics).<sup>12</sup> Descriptive psychology is the project of saying what mental states Karl has; for example, what states or systems of states of knowledge, belief, desire, feeling, perception, imagination, memory he has. As in the case of descriptive semantics, the questions one raises concerning descriptive psychology are empirical questions, and so will need to be responsive to some sort of evidence. Once again, our evidence here will surely include how Karl acts in particular circumstances. But we also might include evidence about how similar Karl is to us, about what sort of creature he is and so on. For example, if Karl stubs his toe, and jumps around, it would be natural for us to attribute to Karl pain of the sort that we would feel if we were in the same sort of situation.

In saying that the question of what mind Karl has is an empirical question, I am as before not denying that there are a priori, or at least very general, constraints on, or theses about, what it is for Karl to have a mind, and that these general constraints will also factor into descriptive psychology. For example, it seems reasonable to suppose that if Karl has some belief states, then together with other states that will cause him to act in certain ways. But the suggestion that belief states have causal powers seems to be a claim about what beliefs are in general, rather than a specific claim about what beliefs Karl has.

Suppose now that after reviewing the evidence and the relevant general constraints we agree that Karl has a particular mind—M17 as it might be. Then, in parallel with the language case, we face foundational psychological questions; these are questions about what facts about Karl make it true that he has M17 as opposed to some other mind, or opposed to no mind at all. If one is a dualist one might well say that Karl's having M17 is a fundamental fact, whereas if one is not a dualist one will say that there are other facts about him in virtue of which he has M17. The foundational project is to say what those facts are.

### **The Intentionality Problem and Descriptive Psychology**

We have reviewed the descriptive/foundational distinction, and proposed a way to extend that distinction to the philosophy of mind. But how does this help with

---

<sup>12</sup> In his 1997, Stalnaker agrees (as I read him) that the descriptive/foundational distinction has a counterpart in the philosophy of mind, but does not go into detail about what that counterpart is, beyond noting (as I have done) that appealing directly to the language of thought is a mistake. See also Stalnaker 2004.

formulating the problem of intentionality, with ‘the problem of the problem of intentionality’ as I called it earlier?

Well, it is very common, in the light of this distinction, to say that the problem of intentionality is a problem about foundational rather than descriptive semantics. Transposing this to the key of mind, the suggestion is that the problem of intentionality is a problem about foundational rather than descriptive psychology. I think there is something right about this suggestion, but it is also misleading. The reason it is misleading is that a lot of the issues that philosophers of mind discuss when they discuss the problem of intentionality turn out to be on the descriptive side of the divide. I will give four examples.

First, take the problem of thinking about non-existent things, such as Odin, that we considered a moment ago. Suppose that the mind that we attribute to Karl—M17—includes the belief that Odin likes ravens. If so, we cannot construe this belief as involving a relation between Karl and the subject of his beliefs, i.e. Odin. For Odin does not exist; hence it is impossible for Karl to stand in a relation to him. This is certainly an aspect of the problem of intentionality, but as we noted before it has very little to do with the contrast between dualism and materialism. In the light of the descriptive/foundational distinction it seems fairly clear why: the problem is one of descriptive psychology. To see this, notice that a common way to solve this problem is to say that the mind that Karl has (i.e. M17) must somehow involve a relation between Karl and an abstract object, for example the property of being the king of the Gods, who plucked out one eye to gain infinite wisdom, who has an eight-legged horse etc. This abstract object exists but is not instantiated, i.e. because Odin does not exist. We might also want to say that if Karl has beliefs about Obama this too involves a relation to an abstract object, it is just that in this case the abstract object is instantiated, i.e. in Obama. Making these assumptions about the mind that Karl has raises further issues—how a concrete object can stand in relation to an abstract object, for example. I will not go into that here. The point is that the hypothesis that Karl stands in a relation to various abstract objects seems to be something we arrive at through theorizing about what sort of mind he has, and so through descriptive psychology.

Second, take the dispute in the philosophy of perception over whether perceptual states are relations to concrete objects, as emphasized by disjunctivists, or

whether they involve representational states of some sort.<sup>13</sup> This is an aspect of a question about intentionality too, but it again it is a question in descriptive psychology. For example, if we attribute M17 to Karl, this will certainly involve some facts about perception, and about how perception relates to belief and so forth. What is the nature of perception? The representationalist says that perceptual states are in some ways akin to belief states in that they involve a certain kind of representational state. Suppose for example that M17 is a mind that involves a certain kind of representational state of the sort mentioned by representationalists, while M92 involves no such state, and simply says that Karl bears a phenomenological relation to his surroundings. Both hypotheses are plausibly compatible with various sorts of data, e.g. behavioural data and introspective data, but they are different from each other. The disjunctivist thinks that the problem with saying that Karl has M17 is that there is no M17 to have, i.e. for there are no perceptual representational states at all. The representationalist, by contrast, thinks that is not so, and that M17 is a possible mind. If so, there is no problem with saying that Karl has M17.

A third example concerns principles of charity, as discussed famously, for example, by Donald Davidson (see Davidson 1974). According to him, when we attribute M17 to Karl, we should be driven by the a priori principle that most of Karl's beliefs are true, and presumably that most of Karl's perceptual states are veridical. Suppose the hypothesis that Karl has M17 entails that most of his beliefs are true, while the hypothesis that Karl has M45 entails that most of his beliefs are false. Davidson's principle of charity is that it is constitutive of the nature of belief that most of a person's beliefs are true. So while M17 might be equivalent to M45 in respect of behavioural evidence, it is rational for us to adopt the hypothesis that M17 is Karl's. Other philosophers disagree with Davidson here, arguing that we have no a priori reason to favor M17 over M45. I don't want to engage this dispute but to note only that it is a disagreement about descriptive psychology.

As a final example, take the dispute about naturalness. There is a problem famously posed by Kripke (following Wittgenstein) about whether Karl—to adapt the issue to our own discussion—is adding or quadding, where to quad two numbers is to produce their sum up to some limit, and then to produce 5 thereafter (see Kripke

---

<sup>13</sup> For the literature on disjunctivism, see the papers in Byrne and Logue 2009, and the paper by Paul Snowdon in this volume.

1982). In his discussion of these matters, Lewis says that that in this case it one should adopt that view that Karl is adding, rather than quadding, because (to put it roughly) this is the most natural rule that Karl could be following (see Lewis 1983). As I understand it, it is conceded by Lewis that our evidence either from introspection or from behaviour could not discriminate these hypotheses; the fact that one hypothesis is natural is suggested as a further constraint that could. Lewis's suggestion is controversial, but for the moment I am not interested in assessing it. Instead, I am interested in noting that it seems to be a part of descriptive psychology. In particular, Lewis seems to be suggesting that in developing our theory of Karl we would need to be driven by the a priori constraint that Karl's mental states attitudes are likely to be natural ones, other things being equal.

In sum, it turns out that a lot of the questions that philosophers discuss when they discuss intentionality are questions in descriptive psychology. This point is important because it shows that the point about the problem of intentionality we mentioned at the beginning of this section—that it concerns foundational rather than descriptive questions—is at best half right. But it is also important for another reason, and this has to do with the question we raised earlier, viz., whether the epistemic response to the knowledge argument might be extended to the problem of intentionality. If by 'the problem of intentionality' we mean what we might call the descriptive problem of intentionality—i.e. the (complex) problem of saying what mind Karl has—then it would seem that there is no easy extension of the epistemic response to the problem of intentionality. This is not to say that we have nothing to learn about what mind Karl or anyone has—on the contrary, the questions here (as we have seen) are empirical, and with respect to those questions, the best policy is surely "tolerance and the experimental spirit," as Quine famously said. Nevertheless, it is not as if when we engage in the project of descriptive psychology we are concerned to assess arguments like the knowledge argument whose persuasiveness depends on all the facts being in; rather the issues have a different shape entirely.

### **The Intentionality Problem and Foundational Psychology**

We have seen that if the problem of intentionality is interpreted as part of the project of descriptive psychology, then the epistemic response to the problem of consciousness is of only marginal relevance to it. But suppose the problem is interpreted instead as part of the project foundational psychology. At the foundational

level, we face the question of in virtue of what (if anything) Karl has M17. To focus on a specific mental state, suppose that, as part of having M17, Karl believes that Obama is president. With respect to this belief, the foundational question we need to focus on is this: in virtue of what does Karl have this belief?

In my view, it is at this point that the considerations we marshalled in the course of developing the epistemic approach to the problem of consciousness have a role to play we turn to thinking about intentionality. The reason is that it is possible to sketch an answer to the question just posed, and this answer is more plausible than it would otherwise be if considered in the light of the epistemic approach.

The answer I have in mind is a version of the well-known Lewis-Armstrong argument for the identity theory (see Lewis 1966, Armstrong 1968). Transposed to our discussion, the first premise of this argument is that when Karl believes that Obama is president he is in a state that plays a particular theoretical role—i.e. it is a state that in Karl produces other states, and is produced in such and such circumstances, and produces such and such actions etc. We might summarize this by saying that according to the first premise the belief that Obama is president in Karl is that state which satisfies role R. The second premise of the argument is that there is some physical state of Karl that satisfies role R. The conclusion drawn from these two premises is that the belief that Obama is president is that physical state. If that argument is sound, it would be fair to say that we would have answered the foundational question about Karl, namely by saying in virtue of what he believes that Obama is president, namely in virtue of being in that physical state.

However, while this argument would (if sound) answer the foundational question, it raises a number of complicated and difficult issues. The first concerns the first premise of the argument. Lewis's defense<sup>14</sup> of this premise involves the suggestion that the premise is not simply true, but true by definition; that is, Lewis thinks that the state of believing that Obama is president may be defined as the state whatever it is that plays the relevant role, something that follows simply from an understanding of the terms. Moreover, according to Lewis, the definition in question (a) constitutes a reductive definition in the sense that the role itself may be spelled out

---

<sup>14</sup> Armstrong seems to me more equivocal on the matter of issue of whether the first premise of the argument discussed in the text is true by definition, and so I will concentrate on Lewis here.

using no psychological vocabulary at all, and (b) is tacitly known by us, somewhat in the way that we know the syntactic rules of our native language (see Lewis 1994).

But, in view of the controversy surrounding the possibility of reductive definitions in philosophy of any sort, this defense of the first premise makes the overall argument seem less plausible than it otherwise might be. Moreover, it is an assumption that is not required by the soundness of the argument: for a valid argument to be sound all that is required is that the premises are true, not that one of them is true by definition. In the light of this, a very natural suggestion is that Lewis-Armstrong argument is much more plausible if assumption that its first premise is analytic is dropped.

However, before we agree to this suggestion, we need to confront the reasons Lewis has for supposing that the first premise is analytic. I think there are a number of considerations motivating Lewis at this point but perhaps the main one (and the one I will concentrate on) is that when he advances this argument, Lewis is concerned to defend, not simply the identity of mental states with physical states, but a certain sort of a priori materialism—i.e. the position we contrasted with a posteriori materialism earlier. In particular, he is interested in the idea that, if the first premise of the argument is true by definition then the second premise of the argument will provide a physical statement that a priori entails the conclusion. And this is exactly as a priori materialism requires.

However, in the light of the epistemic approach to the problem of consciousness mentioned above, it should be clear that is not necessary to defend a priori materialism in this way. For, in the light of that approach, it is possible to separate out two distinct claims: the first claim is that the physical facts, whatever they are, a priori entail the psychological facts; the second claim is that it is possible to define the psychological facts in terms of the physical facts that we currently understand. The epistemic approach is not opposed to the first claim (though it is not committed to it either); that is, it is not inconsistent with that approach that the physical facts a priori entail the psychological facts. But it *is* opposed to the second claim, i.e. because according to it we are ignorant of some of the physical premises required in the entailment. Hence if the picture associated with the epistemic response is coherent, then the first claim of the two just distinguished may be true even if the second is not.

How does this distinction make it more plausible to deny that the first premise of the Lewis-Armstrong argument is true by definition? Well, as we just saw, Lewis's reason for supposing that the first premise is true by definition is that *if* this is so, *then* a priori materialism will be true. And Lewis is undoubtedly correct in asserting this conditional claim. On the other hand, in the light the distinction just made, the reverse conditional is not true: a priori materialism might be true, *even if* the first premise of the argument is not true by definition—indeed, that is possibility made salient by the epistemic view. But this allows us to agree with Lewis that a priori materialism is true but disagree with him that the first premise of the argument is true by definition. And this in turn makes the argument much more plausible than it would otherwise be.

The assumption that its first premise is true by definition is one controversial feature of the Lewis-Armstrong argument. Another concerns its second premise, the suggestion that there is some physical state of Karl that plays the relevant role. The usual way to motivate this premise is to say that materialism is true, and hence that there must be some state which plays the role (if the role is played at all). However, if the materialism at issue here is the sort we referred to earlier as 'standard materialism' this premise seems implausible. For if standard materialism is true, then the premise says that there is some physical state *of a type currently known* of Karl that plays role R, and so, is his believing that Obama is president. But this seems to greatly overstate the current level of understanding that we have into matters of this sort. In some cases, it is plausible to think that the relevant sciences here—i.e. cognitive psychology and neuroscience—have progressed to the point where they might identify some computational state of the brain with which particular mental states might be identified. But in many cases this is not plausible: "The current situation in cognitive science is light years from being satisfactory. Perhaps somebody will fix it eventually; but not, I should think, in the foreseeable future, and not with the tools that we currently have at hand" (Fodor 2000, 5). In short, if the second premise is understood in the light of standard materialism, one might well reject it on empirical grounds. In turn, however, to reject it on empirical grounds is to give up the idea that there is any physical state in virtue of which Karl believes that Obama is president; at this point, the dualist alternative seems the only option.

However, in the light of the epistemic approach to the problem of consciousness mentioned, it should be clear that is not necessary to defend the second

premise of the argument by appealing to standard materialism. For suppose instead we operate with non-standard materialism, suppose the argument is set against the backdrop of non-standard materialism, i.e., against the view that tolerates the idea that we are missing certain types of facts which are relevant to the nature of mind. Then we can think of the Armstrong-Lewis argument as setting out a strategy for solving the problem of intentionality, not as an argument that we currently have the materials to complete. In summary, the perspective suggested by the passage from Jackson from which we began suggests that the premises of the Lewis-Armstrong argument can be defended in a different way from that suggested by Lewis. In turn, doing that provides us with a better answer than we might otherwise have to questions distinctive of foundational psychology.

### **The Problem of Self-Knowledge**

We noted at the outset that Ryle's Cartesian holds two theses. The first is the metaphysical thesis that mental phenomena and physical phenomena are distinct. The second is the epistemological thesis that the mind is transparent to itself, i.e. that we have an introspective faculty which if used correctly can in principle illuminate all aspects of our mind. We have been concentrating on problems and arguments involved in the assessment of the first thesis, suggesting that the epistemic view is sufficient to answer arguments associated with consciousness, and helps out with the arguments associated with intentionality.

Turning to the second thesis, as in the case of metaphysical dualism, one might have thought that a small dose of scientifically informed common sense would be sufficient to reject it as well. Certainly it is a common feature of our intellectual culture that people are in many ways (as Wilson has put it recently) *strangers to themselves* (Cf. Wilson 2002). Social psychologists (not to mention many modern novels) routinely tell us that we often quite wrong about our own basic motives, desires and character traits. Similarly cognitive psychologists and neural scientists portray the human mind as a congeries of different sub-systems operating independently of each other and on principles that are largely unknown to us (e.g. Fodor 1983). From this point of view, it is difficult to believe the picture of the mind as an arena in which in principle nothing is hidden.

As in the case of metaphysical dualism, however, things are not so simple. For what has emerged particularly in recent discussions (e.g., Alston 1971, Shoemaker

1994, Moran 2001, Wright 2000, Byrne 2005) is that while we certainly do not have privileged access to *all* of our mental states, it is nevertheless the case that at least for *some* mental states our first person knowledge is quite *different* in character from (even if not always *better* than) the knowledge that one might have of the conscious states of others or indeed of other things quite generally. Moreover, this modernized form of privileged access has proved difficult to formulate precisely and leads to a number of puzzles and questions, just as Ryle thought, puzzles which have emerged in a somewhat piecemeal form in the literature over the last few decades of philosophical writing.

How might one explain the sense in which self-knowledge is different from other knowledge? It is obviously too late in the paper to give this question adequate attention. But I think that the discussion we have been having about how think to about the problem of intentionality permits us to make a remark about how to think about the self-knowledge problem too. For the problem of self-knowledge is fruitfully thought of as a problem of descriptive psychology, rather like the problems about non-existence, representationalism about perception, principles of charity, and naturalness that we considered earlier. We noted before that, in the case of descriptive semantics, it seemed reasonable as an a priori principle that if some descriptive semantic theory of the language that Karl speaks is correct, then Karl must know, at least for the most part, what according to the semantics he is saying. A parallel suggestion, though suitably modified, might be true in the case, not of language but of mind: if some descriptive psychological theory of Karl is correct, he must be able to know, at least for the most part and to the extent that he is rational, what mental states he is in according to the theory. From this point of view Karl, for example, cannot view what mental states he is in (or what his words mean) as a subject matter that he may or may not take an interest in—in contrast, say to Russian literature, which surely is a topic he may or may not take an interest in. Rather it is a subject matter about which if he is rational he can be assumed to have a certain sort of potential expertise. Spelling out what this expertise amounts too is a difficult matter, and will need to be left for another occasion. The point for us is that it is a project in descriptive psychology.

## References:

- ALSTON, WILLIAM P 1971 'Varieties of privileged access,' *American Philosophical Quarterly* 8 (July): 223-41
- ARMSTRONG, D.M. 1968 *A Materialist Theory of the Mind* (London: Routledge)
- BYRNE, A. 2005 'Introspection,' *Philosophical Topics* 33: 79-104.
- BYRNE, A AND LOGUE, H 2009 *Disjunctivism*, MIT Readers in Contemporary Philosophy, MIT Press
- CHALMERS, D. 1996. *The Conscious Mind*. New York: Oxford University Press
- DAVIDSON, D. 1973 'Radical Interpretation,' *Dialectica* 27: 313-28.
- FODOR, J. 1983. *The Modularity of Mind* Cambridge, MA: MIT Press
- FODOR, J. 1987. *Psychosemantics*. Cambridge, MA: MIT Press
- FODOR J. 2000. *The Mind Doesn't Work that Way*. Cambridge, MA: MIT Press
- JACKSON, F. 1983 'Epiphenomenal Qualia' *Philosophical Quarterly* 32:127-136
- JACKSON, F. 1998 *From Metaphysics to Ethics* Oxford: Oxford University Press
- JOHNSTON, M. 2007 'Objective Mind and the Objectivity of our Minds', *Philosophy and Phenomenological Research*, Volume 75, Issue 2, pages 233–268, September 2007
- KRIPKE, S. 1980. *Naming and Necessity* (Cambridge, MA: Harvard University Press)
- KRIPKE, S. 1982. *Wittgenstein on Rules and Private Language* (Cambridge, MA: Harvard University Press)
- LEWIS, D. 1966. 'An Argument for the Identity Theory,' *The Journal of Philosophy*, 63: 17-25.
- LEWIS, D. 1972. 'Psychophysical and Theoretical Identifications,' *Australasian Journal of Philosophy*, 50: 249-258, repr. in his *Papers in Metaphysics and Epistemology* (Cambridge: Cambridge University Press, 1999), 248-261. All references are to the reprinted version.
- LEWIS, D. 1974. 'Radical Interpretation', *Synthese* 23:331-44
- LEWIS, D. 1983. 'New Work for a Theory of Universals' *Australasian Journal of Philosophy*, 61: 343-377, repr. in his *Papers in Metaphysics and Epistemology* (Cambridge: Cambridge University Press, 1999), 8-55. All references are to the reprinted version.
- LEWIS, D. 1988. 'What Experience Teaches', *Proceedings of the Russellian Society*, 13: 29-57, repr. in his *Papers in Metaphysics and Epistemology* (Cambridge: Cambridge University Press, 1999), 325-31. All references are to the reprinted version. References are to the reprinted version.

- LEWIS, D. 1994. 'Reduction of Mind', in S. Guttenplan ed., *A Companion to the Philosophy of Mind* (Oxford: Blackwell, 1994), 412-31. repr. in his *Papers in Metaphysics and Epistemology* (Cambridge: Cambridge University Press, 1999), 291-324. All references are to the reprinted version.
- LEWIS, D. 1995. 'Should a Materialist Believe in Qualia?' *Australasian Journal of Philosophy*, 73: 140-144, repr. in his *Papers in Metaphysics and Epistemology* (Cambridge: Cambridge University Press, 1999), 325-31. All references are to the reprinted version.
- MORAN, R. 2001 *Authority and Estrangement*. Princeton University Press
- NAGEL, T. 1974 'What it is like to be a bat?', *Philosophical Review*, 83: 435-50
- PAUTZ, A. 2010 'A Simple View of Consciousness' In G. Bealer and R. Koons (eds.) *The Waning of Materialism*, pp. 25-67. Oxford: Oxford University Press.
- ROBINSON, H. 1982 *Matter and Sense* Cambridge University Press.
- RYLE, G. 1949. *The Concept of Mind*. Routledge and Kegan Paul. Reprinted 1963; references are to the reprinted version.
- SHOEMAKER, S 1994 'Self-knowledge and "inner sense"'. *Philosophy and Phenomenological Research* 54:249-314
- SMART, J.J.C. 1959. 'Sensations and Brain Processes'. *Philosophical Review* 68: 141-56
- SNOWDON, P. 2013 'Disjunctivism', The volume
- STANLEY, J AND WILLIAMSON T, 2001. 'Knowing How' *Journal of Philosophy* 98:411-444
- STANLEY, J. 2011. *Know How* (Oxford: Oxford University Press)
- STALNAKER, R. 1997. 'Reference and Necessity', in Wright and Hale (eds) *Blackwell Companion to the Philosophy of Language* (Oxford: Basil Blackwell). Reprinted in Stalnaker, R. 2003 *Ways a World Might Be* (Oxford: Clarendon Press). References are to the reprinted version.
- STALNAKER, R. 2004. 'Lewis on Intentionality', *Australasian Journal of Philosophy*, Vol.82, No.1, pp: 199-212.
- STALNAKER, R. 2008. 'Conceptual Truth and Metaphysical Necessity' in Stalnaker, R. 2003 *Ways a World Might Be* (Oxford: Clarendon Press).
- STICH, S. 1992 'What Is a Theory of Mental Representation' *Mind*, 101, 1992.
- STOLJAR, D. 2005. 'Physicalism and Phenomenal Concepts' *Mind and Language* 20 (5) 469-494

STOLJAR, D. 2006. *Ignorance and Imagination* (New York: Oxford University Press)

STOLJAR, D. 2006a. 'Actors and Zombies' in Alex Byrne and Judith Jarvis Thomson (eds) *Content and Modality: Themes from the Philosophy of Robert Stalnaker* (Oxford: Oxford University Press) pp. 1-17

STOLJAR, D. 2009. 'The Argument from Revelation' in Robert Nola and David Braddon Mitchell (eds) *Conceptual Analysis and Philosophical Naturalism* (Cambridge, MA: MIT Press) pp. 113-138

STOLJAR, D. 2010. *Physicalism* (London: Routledge)

STOLJAR, D. Forthcoming 'Lewis on Materialism and Experience'. In Loewer, B and Schaffer, J (eds) *Blackwell Companion to David Lewis*.

TYE, M 1992 "Naturalism and The Mental" *Mind*, 101, 1992.

WILSON, T. 2002. *Strangers to Ourselves* Cambridge, MA: Harvard University Press.

WRIGHT, C. 2000 Self-knowledge: The Wittgensteinian legacy. In C. Wright, B. Smith & C. Macdonald (eds.), *Knowing Our Own Minds*. Oxford University Press