

Introduction

—

The Slugs and the Tiles

Imagine a mosaic completely constructed from two sorts of tiles, triangles and pieces of pie. The mosaic may of course have many different shapes in it, so long as those shapes are constructed from the basic ones: circles, figure-eights, half-moons, rectangles, rhombuses, and so on. Now imagine a population of intelligent slugs who live on the mosaic and are cognitively sophisticated but whose perceptual access to it is limited (never mind why) to two shape-detecting systems: the first scans the mosaic and detects triangles, the second scans it and detects circles. Given their access to the mosaic, it would be natural for these slugs to think that, at least so far as the tiles of the mosaic are concerned, it was constituted *only* by triangles and circles—of course this is a mistake, but it would be a natural one in the situation.

Now imagine further that among these slugs there is a group of philosophers called slugists who begin to wonder about the nature of the mosaic and to formulate a synoptic vision of what it is and what it is like. For reasons that again need not concern us, many of the slugists take it that the circle-detecting system cannot tell them about the mosaic in its most fundamental nature, and in consequence that circles must not be fundamental or primitive—the circular, they say, supervenes on the non-circular. In particular, the slugists think, the basic element of the mosaic is the triangle—what the mosaic consists of, one of them says, is a vast array of triangular tiles set out in a medium, with every other shape being a construction from these. The slugists now face what they call *the problem of the circle*: the circle-detecting system tells them that there are circles, and yet there is apparently no place for circles in a mosaic completely constructed by triangles.

Why is there apparently no place for circles in a mosaic completely constructed by triangles? Some slugists insist that this is simply obvious, but even they admit that this insistence is, as they put it, weak from a polemical point of view. So two distinct lines of argument are presented to establish, as far as is possible, that the circular does not supervene on the non-circular. The first argument proceeds from the premises, first, that it

is conceivable that there is a mosaic identical to the actual mosaic in all non-circular respects, but for which circles are completely absent, and second, that what is conceivable is possible. The second argument proceeds from a thought-experiment wherein a super-slug is presumed to know everything non-circular about the mosaic but who is not permitted to exercise her circle-detecting system. The super-slug would never know whether the mosaic contained any circles or not. Both arguments conclude that circles are something over and above the rest of the mosaic. As one of the slugists puts it, circles are a special creation—when God was creating the mosaic, he created first the non-circular aspects, and then he had to create the circles in a quite independent act of creation.

Different slugists respond to these arguments in different ways. Some respond that the circle-detecting system is not to be trusted, and that in fact there are no circles, or at least no circles-as-we-traditionally-conceive-them. Others agree that circles are indeed primitive items, and postulate laws linking them and triangles. A third group says that there must be some sort of elusive and extremely non-obvious analysis of circles into triangles. And a fourth says that circles and triangles are necessarily connected in a manner that is irreducibly a posteriori or synthetic in nature.

For us who are neither slugs nor slugists it is clear that all these responses are mistaken, and the arguments that give rise to them fallacious. It is *not* possible that there be a duplicate mosaic in the sense we have defined. A mosaic exactly like the actual mosaic in non-circular respects *would* be exactly like it both in terms of triangles *and* pieces of pie—but such a mosaic would *also* be exactly like the actual mosaic in terms of circles. In other words, circles are *not* a special creation of God, and we stand in need of a plausible diagnosis of what has gone wrong in arguments leading to the non-supervenience of the circular on the non-circular.

What has gone wrong? The obvious and correct answer is that, in view of their perceptual limitation, the slugs have only a selective access to the mosaic. They think that the non-circular truths of the mosaic are completely a matter of triangles. In that situation it is understandable that they find a mosaic without circles possible, and within the confines of this deficit their reasoning is impeccable; understandable and impeccable, but wrong.

The Epistemic View

The story of the slugs and the tiles—the slugs and slugists, though not the tiles, are due to Jackson 1982; see also chapter 5 below—allows us to introduce the main idea to be debated and defended in this book: the epistemic solution to the problem of consciousness, or experience, in philosophy of mind; more briefly, *the epistemic view*. The core of this view is the hypothesis that our epistemic position with respect to the empirical world is analogous to that of the slugs to the mosaic. If our epistemic position is analogous, the correct response to the problem of conscious experience in philosophy of mind is likewise analogous to the correct response to the problem of the circle.

Like the slugists, contemporary philosophers have formulated a synoptic view of what the empirical world is and what it is like. For reasons we will come to, many philosophers take it that introspection cannot tell us about the world in its most fundamental nature, and in consequence that experiences must not be fundamental or primitive—the experiential, they say, supervenes on the non-experiential. In particular, many philosophers think, a complete account of the nature of the world could in principle be given in non-experiential terms. Contemporary philosophy now faces a problem of consciousness, or conscious experience or—as I will usually call it—*the problem of experience*: introspection tells us that there are conscious experiences, and yet there is apparently no place for experiences in a world that is fundamentally non-experiential.

Why is there apparently no place for experiences in a world that is fundamentally non-experiential? Some insist that this is obvious, but even they agree that this is weak from a polemical point of view. And so two distinct lines of argument are presented to establish that the experiential does not supervene on the non-experiential. The first proceeds from the premises, first, that it is conceivable that there is a world identical to the actual world in all non-experiential respects, but different from it in some experiential respects, and second that what is conceivable is possible. The second proceeds from a thought-experiment wherein a super-scientist is presumed to know everything about the non-experiential but has had no relevant experiences herself. The super-scientist would not know the whole truth about these experiences. Both arguments conclude that the experiential is something over and above the non-experiential, a conclusion inconsistent

with the dominant world-view. The problem of experience is the problem of resolving, or learning to live with, this inconsistency.

The problem is simple, but is remarkably persistent and difficult, and survives a great many variations in terminology, formulation and emphasis. How are we to solve it? Well, here is an idea: our position really *is* analogous to that of the slugs. Then the mistake in our thinking when confronted with the problem of experience would be analogous to the mistake in theirs when confronted with the problem of the circle. In particular, ignorance would explain two things: why we find so persuasive the two arguments that lead us to suppose that the experiential fails to supervene; and why it is nevertheless mistaken to follow them through to their conclusion. In short, if our position is analogous to that of the slugs, the arguments that generate the problem of experience would be mistaken, and the problem itself would be solved.

The Ignorance Hypothesis

Of course, to suggest that our epistemic position is analogous to that of the slugs is not to suggest that the two positions are identical in every potential respect. Here are some ways in which the example might be misleading: (i) the slugs are ignorant of what is intuitively an element of the mosaic, and so of the world in which they live. Theirs is a basic-level ignorance, as I will put it later on (see chapters 4 and 7). By contrast, our own ignorance might well take a non-basic form, i.e., we may be ignorant of truths that supervene on the basic truths rather than being themselves basic truths. (ii) The truths of which the slugs are ignorant are *by themselves* sufficient to account for all the truths about circles. By contrast—see chapter 4—the truths of which we are ignorant might account for experience only in combination with other truths, and in particular with truths we already know. (iii) The pie pieces are similar in one way to circles and in another to triangles, and this might suggest a certain form of (what philosophers often call) a dual aspect theory. By contrast, the position I advance is perfectly consistent with a version of physicalism. Finally, (iv) the perceptual limitation of the slugs will be interpreted by some as resulting in a cognitive limitation, and hence as resulting in the slugs being cognitively closed—in McGinn's (1990) phrase—with respect to the truths in question, i.e., in its being a biological

impossibility that the slugs come to know these truths. By contrast—see chapter 5—the truths of which we are ignorant may well turn out to be knowable from our point of view.

Instead of these aspects of the slugs' position, the salient thing for our purposes is this: the slugs are ignorant of a type of truth which is not itself about circles—pie pieces are not circles, after all—but which is nevertheless relevant to the nature of circles. So too, if the epistemic view is correct, we are ignorant of a type of truth which is not itself a truth about experience but which is relevant to the nature of experiences—in brief, a type of experience-relevant non-experiential truth.

Suppose we call the hypothesis that we are ignorant of a type of experience-relevant non-experiential truth *the ignorance hypothesis*. Then we may formulate the epistemic view as the conjunction of two theses. The first is a conditional thesis linking the ignorance hypothesis and the problem of experience:

E1. If the ignorance hypothesis is true, the problem of experience is solved.

The second is the categorical thesis about the antecedent of this conditional:

E2. The ignorance hypothesis is true.

The epistemic view may therefore be assessed for plausibility at both points. We may ask whether the connection between ignorance and the problem of experience is as E1 says. And we may ask whether it is the case, as E2 says, that we are ignorant in this way. In my view, both questions can be plausibly answered in the affirmative, but it is important to see first that they are different questions.

The Plausibility of the First Part

The reasons for answering both questions in the affirmative will be our concern throughout this book. However, to present the main ideas briefly, the crucial considerations in the case of E1 are, first, that the arguments that give rise to the problem of experience are in one form or another modal arguments and, second, that ignorance has an impact on modal arguments. By a 'modal argument' I mean an argument whose starting

points are various imaginable or conceivable cases, whose interim conclusions are various claims about what is and is not possible, and whose final conclusion are claims about the nature of the actual world. That there is *some* connection between ignorance and modal arguments is clear from examples. The slugs advance modal arguments that are frustrated by ignorance. To the extent that our epistemic position is analogous to theirs, our modal arguments would be likewise frustrated.

However, not only is the connection between ignorance and modal arguments intuitively clear, it also may be embedded and explained within an attractive framework of ideas for thinking about modal arguments, a framework articulated by a number of philosophers, including in particular Kripke; and this fact has an important bearing on the plausibility of the epistemic view. In *Naming and Necessity*, Kripke suggested that if we are faced with a modal argument whose conclusion we know or believe to be false, the argument might be disarmed by showing that the proponent of the argument has confused the proposition whose modal status is under discussion for a contrasting proposition that is, in the context, unobjectionable. Kripke himself thought that this method of disarming modal arguments could not be brought to bear in the special case of arguments about experience, a conclusion which to a large degree sets up the contemporary debate: on the one hand, many philosophers believe that the arguments about experience must be disarmed; but on the other hand, the standard ways of disarming the arguments apparently fail. As I will argue, an attractive feature of the epistemic view is that it avoids this dilemma entirely. For Kripke's conclusion presupposes the falsity of the ignorance hypothesis: without the hypothesis, the situation seems, as Kripke implied, completely impenetrable; with the hypothesis, the arguments are susceptible to a version of the standard analysis he outlined, and everything falls into place.

The Plausibility of the Second Part

So much for E1, the thesis about the connection between the ignorance hypothesis and the problem of experience—what about E2, that the hypothesis itself is true? Since the ignorance hypothesis is contingent, it is unreasonable here to expect considerations of a logical sort. And nor should one expect a proponent of the epistemic view to *produce* the truths of which we are supposed to be ignorant; obviously this is ruled out by the nature of

the case. However, what *can* be done, and what I attempt in what follows, is to present grounds that render the ignorance hypothesis plausible even if we cannot argue directly that it is true.

As we will see in chapters 5-7, there are a number of different grounds here. Some have their source in very general observations about humans and their place in the world, and about the sort of philosophical problem that the problem of experience is. Another has its source in some metaphysical speculation that is due in large measure to Russell's discussion in *The Analysis of Matter* and elsewhere. However, the most persuasive consideration from my point of view emerges from the idea that the intellectual situation that confronts us when we reflect on experience is not unprecedented, and the precedents strongly suggest the epistemic view. One precedent is the discussion of linguistic understanding characteristic of the 17th century; another is the discussion about chemistry central to the 19th and early 20th century. In these cases, philosophers advanced modal arguments similar to those advanced today concerning experience. In these cases the modal arguments were a symptom of ignorance; a natural inference is that our own situation is a symptom of ignorance too.

Dialectical Considerations

In addition to these plausibility considerations in favor of the ignorance hypothesis, a number of dialectical considerations will figure prominently in our discussion. The first arises from the fact that, while most philosophers would presumably not deny that we are ignorant of the nature of consciousness in some sense or other, they would also hold that this ignorance is not relevant to the distinctively philosophical aspects of the issue. They hold, or can be brought to hold, that philosophical problems are such that they “are solved, not by giving new information, but by arranging what we have always known”—the phrase is from Wittgenstein's *Philosophical Investigations*. This idea, or something like it, is remarkably influential in philosophy of mind, much more influential than the Wittgensteinian apparatus within which it first found expression. For example, in *The View from Nowhere*, Thomas Nagel holds that the origin of the problem is the contrast between the objective and the subjective standpoints. Similarly, in *The Conscious Mind*, David Chalmers holds that the origin of the problem is the contrast between truths that

concern (what he calls) structure and function and those that do not. If either of these positions is correct, the epistemic view is not simply wrong: it is quite disastrously misguided. For if these views are correct, the problem is conceptual, and as such would persist through the acquisition of the type of knowledge I think is absent. But I will argue in chapter 8 that these views are not correct, and that the problem of experience is not conceptual in origin but epistemic. Indeed, in my view it is the misclassification of the problem as conceptual that constitutes the greatest blockage to progress on the topic.

A different dialectical consideration is that positions alternative to the epistemic view are implausible. What are these alternatives? The contemporary literature is dominated by four positions, parallel to those adopted by the slugs in response to the problem of the circle—I call them eliminativism, primitivism, the a priori entailment view and the a posteriori entailment view. Eliminativism says that there are no experiences, or at least no experiences-as-we-standardly-conceive-them. Primitivism says that experiences are metaphysically distinct from other things, and postulate laws linking the experiential and the non-experiential. The a priori entailment view says that there must be some sort of elusive and extremely non-obvious analysis of the experiential into the non-experiential. And the a posteriori entailment view says that the experiential and non-experiential are related by a relation of necessity that is irreducibly a posteriori or synthetic in nature. These views are very different from one another, and raise all sorts of important and interesting questions. But I will argue in chapters 9-11 that, to the extent that they present genuine alternatives to the epistemic view—an issue which as we will see is controversial—they do not provide plausible solutions to the problem of experience.

Different Versions of the Same View

While the epistemic view is not in the standard catalogue, versions of it have certainly been suggested before and in that sense I claim no great originality in advancing it. What then is distinctive about the present discussion? Some new elements, or at least elements not set out elsewhere the way I do here, have already been mentioned: the spelling out of the connection between ignorance and modal epistemology; the claim that our present situation has important historical precedent; and the critique of positions which deny the ignorance hypothesis, either as an account of the problem itself or as an account of a

solution to the problem. But there are three further features that differentiate our own contribution from others.

The first is an emphasis on the distinctive nature of the philosophical problem of experience. Philosophers who appeal to ignorance are sometimes viewed as, and sometime are, skeptical of their subject. The position adopted here is not skeptical in this sense; on the contrary, the first part of the book is a defense of a philosophical problem. In fact, the difference between empirical and philosophical questions is important for our position. For suppose the only interesting problem or problems in the area were empirical. Then our position would hardly merit formulation; *obviously* we are empirically ignorant with respect to experience. On the other hand, if there *are* distinctively philosophical problems in the area our position is not only interesting, but is widely held to be false. For many contemporary philosophers, the solution to the philosophical questions of experience lies not in ignorance of some relevant facts but rather in correcting a mistaken conception or picture of experience, or else a mistaken conception or picture of our thought and talk about experience. I think these suggestions are unpersuasive, but it is important to see at the outset that they are different from our own. Indeed, from our point of view, the problem is not a mistaken conception of experience so much as a mistaken conception of the *problem* of experience; in particular, it is a mistake to think of it as at root conceptual.

The second distinctive feature is an emphasis on the particularity of the problem of experience. Many contemporary philosophers emphasize the similarity among the problems they discuss. Problems about conscious experience, the self, morality, modality, semantics, and probability are all presented as having a similar structure, and even as reflecting a single underlying problem, which has a 'value case', a 'probability case' and so on. Of course in one sense this is unobjectionable; one may certainly learn something about a problem by comparing it to others, and we ourselves will do a certain amount of comparing and contrasting in what follows. But the approach also tends to suppress the differences among problems, and in consequence suppresses the possibility of the epistemic view. For when one looks sideways at many other philosophical problems, a counterpart of the epistemic view is often implausible. If these other problems are a guide to the experience problem, the inevitable conclusion is that the epistemic view is implausible in this case too. But in my view this conclusion is quite mistaken: the

implausibility of the appeal to ignorance in other cases is evidence, not of the falsity of the epistemic view, but of the difference between the problem of experience and other philosophical problems.

The third distinctive feature of our discussion has to do with the fact that, in the past, appeals to ignorance have been made within the context of relatively radical views about the nature and content of the ignorance in question. For example, some view the ignorance hypothesis as essentially involving an appeal to new physics or to quantum theory or at least as involving a speculation about the future course of science. Some view our ignorance as a consequence of the scientific or intellectual paradigm that we are currently in, and that in consequence we cannot remedy it without revising the notion of the material world or the notion of experience or both. Some view our ignorance as a consequence of the cognitive structures that are part of our genetic endowment, and that as such we are cognitively closed with respect to the solution to the problem. And some view our ignorance as a consequence, not of historical or psychological facts, but of the a priori structure of empirical inquiry, as following from the fact (if it is a fact) that empirical inquiry does not acquaint us with the intrinsic nature of matter. Whatever the plausibility of these views considered alone, they will play no essential role in the version of the epistemic view defended here. My view is that a hypothesis about our current epistemic situation is the best explanation for the distinctively philosophical predicament we are confronted with when we think about experience. No quantum mechanical theory of consciousness follows from this, no speculations about the future course of science, and nothing at all about the limits of thought, regardless of whether those limits are of an historical, psychological or a priori nature.

The effects of our emphasis on the uniqueness of the problem of experience, and of divorcing the epistemic view from its radical context, will, I hope, be felt throughout our discussion. But one effect should be stated immediately. The fact that our view has implausible counterparts in other cases, and has often been associated with radical views, encourages the impression that a view of this kind is itself radical and implausible, and this encourages the further impression that it might legitimately be set aside—if you *already* think a view implausible, it does not take much persuasion to ignore it. I think this greatly underestimates the epistemic view. Not only is the view not radical, it is the

obvious and common sense position to adopt in the circumstances—at any rate, so I hope to have convinced you by the end of the book.

Outline of the Argument

So that is the epistemic view, or at least an initial sketch of the view. How is our defense of it to proceed?

Part I (chapters 1-3) is concerned with the identification of the problem for which the epistemic view is a solution. In chapter 1, I identify the subject matter of the problem: experience as presented by the *phenomenal conception*. According to this conception, experiences are psychological events whose defining characteristic is that there is something it is like to undergo them—whose defining characteristic is their *phenomenal character*, as it sometimes put. I argue that the phenomenal conception is a minimal conception in that it is neutral on many aspects of experience that are thought to be philosophically puzzling.

In chapter 2, I identify the problem to be discussed about this subject matter: *the logical problem*. The logical problem is best presented as the conjunction of three theses which are individually plausible but jointly inconsistent: the thesis that there are experiential truths; the thesis that that if there are experiential truths, such truths are entailed by non-experiential truths; and the thesis that if there are experiential truths, such truths are not entailed by non-experiential truths. I distinguish the logical problem from two other problems with which it is apt to be confused, the empirical problem and the traditional mind-body problem.

In chapter 3, I discuss the skeptical challenge that it is mistaken to suppose that there is a genuinely philosophical problem of experience at all. According to one version of the challenge, the problem should be rejected on methodological grounds: it is mistaken to reason from conceivability to possibility at all. My response is that this reasoning is ubiquitous in philosophy, and thus to the extent that there is a problem here it is everyone's rather than mine. According to another version of the challenge, the problem should be rejected on conceptual grounds: the thesis of physicalism certainly made sense at a particular moment in the history of science, but that moment is long gone and contemporary interpretations of it do not permit the questions typical of philosophy of

mind to be legitimately raised. My response is that physicalism and related concepts play an illustrative or inessential role, rather than an essential role, in the logical problem, and that, once this is appreciated, the basis for skepticism evaporates.

Part II (chapters 4-8) sets out and defends our solution to the problem. In chapter 4 I take up E1, the thesis that if the ignorance hypothesis is true, the logical problem is solved. My defense of this thesis begins from the idea that there are a number of distinctive ways that a modal argument can go wrong—a number of standard mistakes, as I call them. The philosophical challenge presented by the modal arguments that are constitutive of the logical problem is that it is hard to see in these cases that we are making any of these standard mistakes. On the other hand, if the ignorance hypothesis is true, it becomes immediately plausible to suppose that we are making one of (or a combination of) these mistakes. This suggests in turn that E1 is true: if the ignorance hypothesis is true, the logical problem is solved.

In chapters 5-7, I turn from E1 to E2, and discuss the case for believing the ignorance hypothesis in the first place. Chapter 5 starts by emphasizing our ignorance in general, our ignorance of experience in particular, the way in which the assumption of ignorance explains a philosophical predicament that otherwise remains elusive, and the way in which the ignorance hypothesis is suggested by the basic structure of the logical problem. Chapter 6 takes up Russell's idea that reflection on the nature of scientific inquiry prompts a version of the ignorance hypothesis, an idea much discussed in contemporary philosophy. Chapter 7 proceeds by pointing out that the epistemic view is known to be correct for older philosophical problems which are structurally analogous to the logical problem, problems which concern, not the phenomenal character of experience, but the intellectual character of thought and the chemical features of physical objects.

In chapter 8, I respond to objections, concentrating on two in particular. According to the first, the epistemic view is mistaken because if the problem at issue is understood correctly, we *are* in possession of the relevant truths; according to the second, the epistemic view is mistaken because it has a range of alarming side effects, some of which have been suggested by proponents of the view themselves. These objections raise important questions, but I will suggest that reflection on them strengthens, rather than weakens, our position.

Part III (chapters 9-11) is concerned with alternatives to our position, and with securing the suggestion that a comparative analysis favors the epistemic view. In chapter 9, I discuss a posteriori entailment, the view that the supervenience of the experiential on the non-experiential is irreducibly a posteriori. The main problem for this view emerges when we notice that the philosopher who did most to make it prominent, Kripke, also considered and rejected it. In my view, the lesson of Kripke's discussion on this point is that the mere idea of a posteriori entailment does not solve the problem of experience, and therefore that a proponent of the a posteriori entailment view is obliged to add further material. On the other hand, an examination of what this further material might be yields the result that either the a posteriori entailment view has no answer to the arguments, or else collapses into the epistemic view.

In chapter 10, I advance an analogous argument in the case of the a priori entailment view. On the one hand, the mere idea of priori entailment will not answer the argument—you need to add further material. On the other hand, an examination of what this further material might be yields the result that either the a priori entailment view has no answer to the arguments, or else collapses into the epistemic view.

In chapter 11, I conclude the overall discussion by doing three things: I review the reasons for rejecting eliminativism and primitivism; I criticize an independent argument for these views that would, if successful, also undermine our own position—this is an argument founded on the idea, known as 'revelation', that understanding experience means knowing its essence; finally, I summarize and make explicit some general morals.