

## SOMETHING LIKE ABILITY

Paul Noordhof

### Abstract

One diagnosis of what is wrong with the Knowledge Argument rests on the Ability Hypothesis. This couples an ability analysis of knowing what an experience is like together with a denial that phenomenal propositions exist. I argue against both components. I consider three arguments against the existence of phenomenal propositions and find them wanting. Nevertheless I deny that knowing phenomenal propositions is part of knowing what an experience is like. I provide a hybrid account of knowing what an experience is like which is the coherent expression of a single idea: knowing what an experience is like is knowing what it would be like to have the phenomenal content of the experience as the content of an experience one is currently having. I explain how my conclusions indicate that the focus of discussion should be on the alleged explanatory gap between phenomenal facts and physical facts and not on the Knowledge Argument. The latter is a poor expression of the difficulty Physicalists face.

Some opponents of Physicalism have argued that we cannot know what a particular kind of experience is like without having had it. Full physical information about the brain (or, indeed, the environment) wouldn't suffice. They claim that, if Physicalism were true, we could know what a particular kind of experience is like on the basis of the relevant physical information. Hence, Physicalism is false (This is the *Knowledge Argument*) [Jackson 1982; Jackson 1986; Nagel 1979; Robinson 1983: 4–5].<sup>1</sup>

Most friends (and one fairminded enemy) of Physicalism have tried to explain how they can agree that we cannot know what a particular kind of experience is like without having had it and yet deny that this means that Physicalism is false.<sup>2</sup> They differ over the explanation they give. One answer is the so-called *Ability Hypothesis*. It has two components. First, there is the *Ability Analysis* of knowing what an experience is like

(A1) S knows what an experience of kind E is like iff S has the ability to recognize the experience, imagine the experience, and recall the experience.

[Nemirow 1980: 475; Lewis 1983: 131; Lewis 1999a: 285–6; Lewis 1999b: 326–8; Carruthers 1986: 143–5; Nemirow 1990: 493–4; Mellor 1993: 2–6].<sup>3</sup>

<sup>1</sup> It is more accurate to describe Nagel [1979] as concluding that we don't know how Physicalism could be true rather than that Physicalism is false. This is not always how he has been taken.

<sup>2</sup> Hugh Mellor [1993] endorses the Ability Hypothesis but thinks that there are other reasons for not being a Physicalist.

<sup>3</sup> Proponents of the Ability Analysis differ slightly in matters of emphasis. Laurence Nemirow [1980; 1990] claims that the basic ability is the ability to imagine or visualize the experience. He

Second, there is the claim that

(A2) There are no propositions concerning what an experience is like.

Proponents of the Ability Hypothesis claim that there is no such thing as ‘phenomenal information’ [Lewis 1983: 131; Lewis 1999a: 286–7]. There is no fact concerning what an experience is like, not even a physical fact [Nemirow 1990: 492, 494–5; Mellor 1993: 6–9]. Call (A2) the *Denial of Phenomenal Propositions*.

The proponents of the Ability Hypothesis take the Ability Analysis and the Denial of Phenomenal Propositions as a mutually supporting package. I will argue that these elements come apart. Although the Ability Analysis is false, an analysis which draws upon elements of it is correct. This will be enough to establish that the original Knowledge Argument is unsound. The arguments that proponents of the Ability Hypothesis provide to establish that there are no phenomenal propositions don’t work. The key move is to recognize that, strange as it may seem, knowing what an experience is like does not require knowledge of phenomenal propositions. The proponents of the Ability Hypothesis have engaged in overkill.

The paper will proceed as follows. First, I shall refine the Ability Hypothesis and state the preliminary case against it. Then I shall discuss the arguments of David Lewis, Hugh Mellor and Laurence Nemirow that it must be true. I shall explain why their arguments don’t establish that there are no phenomenal propositions. In seeing what is wrong with these arguments, we shall see why the Ability Analysis component of the Ability Hypothesis needs further reformulation. Finally, I shall present my own positive account of knowing what an experience is like and trace out the consequences of this account for the assessment of Physicalism. The upshot will be that a reformulation of the Knowledge Argument can be provided which threatens to be more persuasive. However, the reformulation just makes clear that the real issue is the alleged explanatory gap between physical properties and phenomenal properties. My paper only indirectly contributes to that debate. My aim is to establish once and for all that our focus should be on that issue at the expense of the Knowledge Argument.

### I. Refining the Ability Hypothesis: Its Plausibility and the Case Against It

In order to obtain the Ability Hypothesis in its most plausible form, the Ability Analysis of knowing what an experience is like needs a bit of tightening up. For one thing, it does not seem that the ability to recall an experience is necessary so long as the ability to

<sup>3</sup> *continued*

holds that the ability to recognize something constitutes a less sophisticated, immature notion of knowledge of what something is like different from that discussed by Nagel [1979] and Jackson [1990: 497]. He holds that recollection of an experience is just a way to refresh one’s ability to imagine [494]. In his earliest version, Lewis doesn’t mention the ability to recall but adds the ability to predict one’s behaviour via imaginative experiments [1983: 131]. He suggests that gaining these abilities might be like making a template that can be applied to patterns in the future [131–2]. In his 1988 article, Lewis mentions all three abilities [1999a: 285–6]. In his most recent discussion, although he only refers to the ability to recognize and imagine, he just says that the abilities obtained *include* these two and refers without qualification back to his 1988 article [1999b: 326–7]. Mellor argues that the ability to recall is not necessary [1993: 5–6]. Carruthers similarly focuses on the ability to recognize and imagine [1986: 144–5].

imagine it is present. This is obvious in the case of complex imaginative acts like reading a musical score and imagining how it must sound, forming a picture in one's mind before transcribing it onto canvas, and so on. As Hugh Mellor points out, these acts are not just a matter of putting simple *recalled* elements together. They involve imaginatively extrapolating what results from their combination—e.g. the dissonance of a semitone between two simultaneously played notes—in order to know what an experience would be like [Mellor 1993: 5–6]. Equally, when a subject first experiences the taste of marmite (apparently the UK equivalent of vegemite), only two of the necessary conditions for having the ability to recall the experience are met, namely having the experience at some point prior to the recollection and being able to imagine the experience [Gertler 1999: 323; Raymont 1999: 120].<sup>4</sup> A third, namely that the experience must be over, is needed for the ability to be correctly attributed. It appears that recollection, at best, enables us to imagine experiences. It is not an essential element in knowing what an experience is like.

It might look as if proponents of the Ability Hypothesis abandon appeal to recollection at their peril. The fact that we can recall an experience only if we have previously had it might seem to fit rather well with the desire to explain why we can only know what an experience is like if we have experienced it. They need an innocuous explanation of the latter if they are to successfully oppose the Non-Physicalist. However, if the Ability Hypothesis relied upon this feature of recollection to provide the explanation, the approach would be trivialised. One can recall anything only if what is recalled is true.

There is a similar doubt about the role of recognition. *Re*-cognition implies previous cognition. I can only recognize things if I have come across them, a picture of them, or a description of them, some time before. Once more that seems to trivialize the idea that we need experience in order to know what something is like. Nevertheless, we need something to play a role akin to that of recognition in the analysis. Suppose I have never had a certain experience, E, but can imagine E. That might suffice for knowing what E is like if I am able to imagine E correctly. What could this correctness amount to? A reasonable suggestion is that I have imagined something correctly only if I am able to recognize that I have imagined it [Mellor 1993: 4–5]. But, in the light of my strictures against trivialization, I think we should put identification in place of recognition. We can identify things without having previously come across them.<sup>5</sup> If we can identify E as what we have imagined, we have imagined it correctly.

With these modifications the Ability Analysis now reads

S knows what an experience of kind E is like iff S has the ability to imagine the experience and identify the experience, both when it is had and when it is imagined.

The analysis suggests a particular plausible reading of the claim that we cannot know what a particular kind of experience is like without having had it. The modality of

<sup>4</sup> Many thanks to an anonymous referee for drawing Raymont's work to my attention.

<sup>5</sup> Mellor takes recognition to involve not being surprised at the nature of something or the fact that its nature matches one's expectations [1993: 4]. If Mellor is right that recognition only involves this—i.e., is rather like my identification—then this paragraph is otiose. He says that he is happy to accept 'identify' so long as I disclaim its essentialist overtones (personal communication). I hereby do. I should note that the notion of an ability to identify to which I appeal is nonfactive. In being able to identify an experience one is imagining, one needn't necessarily be having the experience in question. I try to address doubts about this notion in section III.

‘cannot’ is psychological impossibility. This may vary from subject to subject. It may be psychologically impossible for one subject to know what experiencing a certain piece of music would be like without actually experiencing it, another may pick it up from the score. The modality of the ‘cannot’ is certainly not metaphysical impossibility. It would be metaphysically possible for subjects to have these abilities without having had an experience of the appropriate kind, say by banging their heads against a wall instead [Lewis 1999a: 264–5, 288]. This is one point of advantage over those who have argued that the knowledge of what an experience is like is knowledge by acquaintance [Conee 1994: 139–42]. It would be metaphysically impossible for someone to have knowledge by acquaintance of a certain kind of experience if they had not had an experience of that kind. Whereas it seems clear that they would know what the experience was like if they were able to imagine it.

This may appear to need some qualification.<sup>6</sup> Earl Conee gives a nice illustration of the issue.

Suppose that Martha is a superlative colour interpolator. She is highly skilled at visualizing an intermediary shade that she has not experienced between pairs of shades that she has experienced. Martha happens not to have any familiarity with the shade known as cherry red. She has seen, and vividly recalls, the look of burgundy red and the look of fire engine red. At this moment, before Martha has imaginatively interpolated between those two shades, it is clear that Martha does not yet know what it is like to see something cherry red. She does not know this, although she is fully prepared to find out by exercising her imagination. Yet Martha already knows how to visualize cherry red, since she knows how to perform the imaginative interpolation between burgundy and fire engine red. Thus knowing how to visualize something cherry red at will is not sufficient for knowing what it is like to see the colour.

[Conee 1994: 138]

If Conee is right, it might seem that the mere ability to imagine an experience is not enough. One must either have had the experience at some point or manifested the ability to imagine the experience.

There are two factors at work in Conee’s discussion: interpolation and manifestation. Martha has not yet interpolated her previous experiences. Given that this is so, it seems incorrect to say that Martha presently has the ability to imagine cherry red. Rather her successful interpolation would enable her to imagine cherry red. Of course, it is a nice question when we should attribute an ability rather than the ability to acquire an ability. Part of the answer depends upon precisely how we describe the ability. For instance, we might say that I am able to give the answer to a complex arithmetical sum because I could spend some time working it out. I certainly don’t have the ability to give the answer *just like that* though. In order to give a correct analysis of knowing what an experience is like, we need to appeal to an ability to imagine *just like that* (an *immediate ability* let’s call it). Once this distinction is made, Conee’s case is not a knock down counterexample to the

<sup>6</sup> I would like to thank an anonymous referee for raising this issue.

Ability Analysis. The case rests upon the point about interpolation but need for interpolation can be taken to imply an absence of the requisite immediate ability.<sup>7</sup>

This may not assuage all worries though. Is there still a lurking feeling that, unless the ability to imagine cherry red has been *manifested*, the subject does not really know what it is like to experience cherry red? But this is hard to square with the intuition, shared by Conee, that knowledge of what an experience is like need not involve current knowledge by acquaintance [Conee 1994: 143]. What makes a subject still know what an experience is like when he or she is not now acquainted with that experience? Why isn't the loss of the acquaintance sufficient to undermine a subject's claim to know what an experience is like? The answer seems to be that the presently unmanifested ability to imagine what an experience is like is sufficient to establish that a subject knows what the experience is like *now*. If the ability need not be presently manifested, it is hard to see why it must have been manifested in the past.

This point has increased force when it is remembered that, whatever phenomenal similarities imagining and perceiving have, they are phenomenally different in certain respects. The ability to imagine and identify the perceptual experience imagined do not amount to an ability to acquaint oneself with the nature of the perceptual experience. Imagination is not a re-presentation of exactly what the perceptual experience is like. It is rather that, according to this view, a subject's knowledge of what an experience is like involves the integration of his or her powers of identification and imagination. He or she can imagine the appropriate perceptual experiences, if called for, and display his or her appreciation of what they are like by being able to identify correctly what is imagined.

If the Ability Analysis were correct, we would have a simple explanation of why we cannot know what a certain type of experience is like on the basis of physical information. There is no reason to suppose that such information will convey to us the relevant abilities. In which case, there is no reason to reject Physicalism. Unfortunately, even with these revisions, the Ability Hypothesis is implausible. Our discourse about what an experience is like appears to express propositions. When I state that this is what it is like to have an experience of a red tomato or of the smell of jasmine, it seems that I am giving propositional expression to my knowledge of what an experience is like. If the Ability Hypothesis is correct, this appearance is misleading. On that view, I am *expressing* a capacity to imagine and identify. The pieces of language used don't express propositions but rather express a capacity to imagine and identify. We might accept this in the case of first person reports about our own experiences. It is harder to accept in the case of other people's experiences. If I state that Mary has an experience like this, then, according to the Ability Hypothesis, I am expressing my ability to imagine and identify her experience. It seems preferable to take my statement as attributing something to Mary rather than myself [Jackson 1986: 292–3; Loar 1997: 607]. Equally, if I argue that (i) if this is what it is like to have an experience of red, then it is similar to what it is like to have an experience of orange, (ii) this is what it is like to have an experience of red, therefore (iii) it is similar to what it is like to have an experience of orange, it seems that I have engaged in a sound piece of reasoning. Yet it is hard to see how the Ability Hypothesis

<sup>7</sup> Note this is not to deny that abilities can be realized in different ways (see Raymont [1999: 118–20], for a nice discussion of this). It is just to claim that there is a particular way of having the relevant ability to which the proponent of the Ability Analysis will want to appeal.

would account for this. Although (ii) might be the expression of an ability, the antecedent of (i) does not seem to be so [Loar 1997: 607; Lycan 1996: 94]. In which case, the argument involves an equivocation. If there genuinely are no phenomenal propositions, then the proponents of the Ability Hypothesis would have to take ‘this is what it is like to have an experience of red’ as the expression of an ability to imagine an experience of red. Yet one would not express that ability by saying ‘*if this is what it is like to have an experience of red*, then it is similar to what it is like to have an experience of orange’. In this case, one need not even be taking oneself to exercise the ability of imagining something red. Proponents of the Ability Hypothesis would have to claim that the conditional was an expression of a higher-order ability to recognize that, from the ability to imagine and identify orange, one may extrapolate to arrive at the ability to imagine the experience currently being enjoyed. They would also have to deny that the argument was sound or extend the notion of soundness to include arguments that cannot be understood in terms of truth. As Loar remarks, the issues this raises are familiar in discussions of Emotivism [Loar 1997: 607; Geach 1964; Hale 1986; Blackburn 1988].

In the light of these two points, it seems that the Ability Hypothesis is revisionary and, thereby, in need of independent argument. In the next section, I shall examine the three arguments that have been provided and find them wanting.

## II. Arguments in Favour of the Ability Hypothesis

### A. *The Argument from Imagination*

Hugh Mellor suggests that one reason for supposing that knowledge of what an experience is like is not propositional knowledge is that we cannot know what an experience is like without being able to imagine it. He claims that this is not so in the case of knowledge about the external world, for instance, that a pillar box is red or that the gate is closed [1993: 7]. I presume that the following line of thought is at work. If there were some true proposition concerning what an experience is like, then we could know what the experience was like just by knowing the proposition. Knowing what an experience is like needn’t involve the ability to imagine. Since it must, there is no such true proposition.

This argument appears to fail. I grant that if I am not currently experiencing something red, then I can only know what the experience is like if I can imagine and identify it. However, if I am currently experiencing something red then it does not matter whether I can imagine it. I know what an experience of red is like just through having the experience and being able to identify it. Of course, in the case of experiences of red, it might seem impossible to experience something red and not to be able to imagine what the experience is like. But the apparent simplicity of the case should not obscure the general point. Unless there is some intimate connection between being conscious of an experience and being able to imagine it, nothing rules out the existence of perceivers who are unable to imagine their experiences and yet who know what it is like to have these experiences when they are having them. We are unable to imagine in all its detail an experience of a complex object, e.g., a Vermeer painting. Yet, when we are having the experience, we know exactly what it is like to experience an object of that complexity. If this is so, it is hard to deny the more extreme possibility that there might be perceptive but unimaginative creatures.

The ability to identify an experience by itself is not enough to count as knowing what an experience is like. So no weakening of the Ability Analysis in this direction will do. Suppose that, if I had a bat's echolatory experiences, I would identify them as such. As Mellor points out, it would not follow that I knew *now* what these experiences were like [1993: 4]. So it seems that proponents of the Ability Hypothesis should discard the Ability Analysis for the following hybrid analysis.

S knows what an experience of type E is like iff (i) *either* S can imagine an experience of type E *or* S is having an experience of type E, and (ii) S can identify experiences of kind E both when they are imagined and when they are had.

Having made this concession, they could argue that a subject could not know what an experience was like in any other way, in particular, through knowing a proposition concerning what that experience was like. In which case, we would still have an answer to the Knowledge Argument. The reason why physical information would not provide knowledge of what an experience is like need have nothing to do with the existence of non-physical facts.

One problem with this response is that, when we perceive that the pillar box is red, we gain knowledge of a fact [Mellor 1993: 7]. Similarly, it might be thought that when we have an experience of kind E which we can identify, we obtain knowledge of the fact that the experience is like such and such. The next two arguments try to undermine this line of reasoning.

### B. The Argument from External Objects

It does not seem that we are limited to knowing what experiences are like. We can know what objects are like in the same way. For instance, we know what thunder and lightning are like by being able to imagine the crashes and flashes and we know what St Paul's is like by being able to imagine the dome, pillars and the rest of the building or, for that matter, clapping eyes on the building. This suggests that the analysis of knowing what experiences are like given at the end of the last section can be generalised.

S knows what an X is like iff (i) *either* S has the ability to imagine an X *or* S is currently having an experience of X, and (ii) S has the ability to identify X in imagination and experience.

Mellor argues that we would not be inclined to say that there were propositions about what these things are like (hereafter let me extend my use of '*phenomenal propositions*' to cover this possibility). So if there is knowledge of what X is like, it does not follow that there is a phenomenal proposition about X. Once the connection between knowing what something is like and the existence of phenomenal propositions about it is broken in this case, the connection anywhere becomes very much open to question. This is his second argument against the existence of phenomenal propositions describing what experiences are like.

In case we are inclined to protest that there *are* propositions about what objects are like—for instance, that a pillar box is red—Mellor reminds us that, following Thomas

Nagel, we suppose that the phrase ‘what it’s like’ characterizes a distinctive feature of experience [Nagel 1979: 166–7]. So we should be wary of allowing that there are statements about what pillar boxes and St Paul’s are like [Mellor 1993: 1, 9–10]. Unless, that is, we are closet Idealists.<sup>8</sup>

However, I don’t think we should allow ourselves to be embarrassed out of the truth. ‘The pillar box is red’, ‘St Paul’s has a large Dome and pillars’, ‘There were flashes of lightning and crashes of thunder’ describe what pillar boxes and St Paul’s are like and what the evening storm was like respectively. That’s because they describe features of St Paul’s and the evening storm which would be experienced if a subject were in the appropriate circumstances. Phenomenal propositions concern possible contents of experience. We do not attribute something which is distinctive of experience to objects when we describe what they are like by noting that they are red, crashing, have domes or, indeed, are having blue experiences. But that doesn’t matter. What is distinctive of experience is that there is something that it is like to have the experience *for* the subject of experience.<sup>9</sup> There is nothing that it is like to have the property of redness (say) *for* the object that is red.

Once we have in our sights the distinction between there being something that an object is like to experience and there being something it is like to have the experience *for the subject*, this motivation for denying the existence of phenomenal propositions about the external world vanishes. Nevertheless, it might still be thought that phenomenal propositions about the world are in some sense derivative and hence not really phenomenal propositions about the world as such. One way of putting this thought would be to say that, when we state that an object in the world is like such and such, we are describing the kinds of experience that we would have if we were experiencing the object. The statement would only be saying something about the object in question because it is the normal cause of the type of experiences we would undergo.<sup>10</sup>

One thing that seems wrong with this position is that it fails to recognize that the proper way to capture at least part of what an experience is like is by describing the objects and properties in the world the experience concerns. It is not as if the experience seems to have a nature independent of this characterization. That’s why, although implicitly we are revealing the kind of experiences we would have when we describe what an object is like, we do not actually state that we would have these experiences when describing what an object in the world is like. This brings me to a second point. As an analysis of phenomenal propositions about the world, the alternative proposal is counterintuitive. As things stand, phenomenal propositions about the world seem precisely that, propositions about the world. It is one thing to say that—as I sought to do—

<sup>8</sup> It is puzzling that Mellor is prepared to allow that there is knowledge of what objects in the world are like bearing in mind he takes the phrase ‘what it’s like’ to apply only to experiences. Perhaps he is addressing his remarks to people who are misguided enough to believe that there is knowledge of what objects in the world are like (people like me). But then, we are not likely to be convinced that there are no phenomenal statements about what objects in the world are like.

<sup>9</sup> See, for example, Nagel [1979: 166–7]. It is worth noting that this involves a special sense of the phrase: *there is something it is like to have the experience for the subject of experience*. As Ted Honderich pointed out, it implicitly appeals to the notion of phenomenal consciousness [1988: 77–8]. We can ask what is it *like for an object to be striped* (not knowing what stripes are) and be told about the arrangement of colours

<sup>10</sup> I’m grateful to an anonymous referee for raising this issue.



phenomenal propositions concern only those objects and properties that show up in experience. That still makes these propositions about the world, just a subset of it. It is quite another to say that these phenomenal propositions are about experiences. Of course, some of an object's properties may be characterized in terms of the kind of experiences to which they would give rise. Colour may be an example. But that doesn't make phenomenal propositions about objects and properties really propositions about experiences. Objects may have these properties without there being any experiences at all.

Suppose then there are phenomenal propositions about the world in the sense I've specified. In which case, since one person's *modus ponens* is another person's *modus tollens*, we might expect that there are phenomenal propositions about experiences too. Yet mention of these propositions does not play a part in the analysis of knowing what something is like which I have just provided. That opens up the following possibility. The correct account of knowing what an experience is like does not involve knowledge of phenomenal propositions. This point is of some importance. It provides us with a way of dealing with the third and final argument against their existence.

### *C. The Argument from Inexpressibility*

The most popular argument in favour of the Ability Hypothesis has been the argument from inexpressibility [Mellor 1993: 7–9; Nemirow 1990: 493]. The argument runs as follows.

- (1) If there were propositions concerning what experiences are like, then we would be able to make statements about what experiences are like.
- (2) We cannot make statements concerning what experiences are like.

Therefore,

- (3) There are no propositions concerning what experiences are like.

I do not propose to contest that, if there were propositions, we would be able to express them. Instead, I want to consider in more detail the nature of phenomenal propositions and the statements that convey them (hereafter 'phenomenal statements').

Mellor seems to hold, albeit implicitly, something like the following view.

- (I) A phenomenal statement about experience should provide somebody, who has not had the experience and who cannot imagine it, with the ability to imagine it (without essential recourse to imaginative extrapolation) [Mellor 1993: 7].

I agree with him that there are no phenomenal statements in this sense. If a blind person asks me what it is like to experience red, there is nothing that I can say which will enable him or her to imagine it. However, this only establishes that we cannot make phenomenal statements if this is a plausible understanding of what such statements should provide. That is very much open to question.

The following seems just as plausible a characterization of phenomenal statements.

(II) A phenomenal statement about what A's token experience is like should provide another person, B, who has the ability to imagine experiences *of that kind*—or similar kinds so that B may extrapolate—with the ability to imagine what A's token experience is like.

According to this characterization, the following would count as phenomenal statements: I experienced the smell of coffee, I had a migraine, I experienced something red, I felt a knot in my stomach, I felt anxious, I felt a wave of nausea, I felt so angry I could barely speak, and so on. Each of these statements would enable you to imagine what one of my experiences were like. Each provides us with a description of a possible content of experience. This time, though, they concern our mental lives.

The characterization of phenomenal statements contained in (II) does not require that anybody who understands a phenomenal statement should be able to imagine what the relevant experience is like. It only requires that if a subject is able to imagine experiences of the appropriate kind, then he or she should, in understanding the statement, be able to imagine a token experience described by it. A congenitally blind person can understand and, indeed, have a justified belief that I am having an experience of red. He or she will probably think of an experience of red as the actual occupant of a certain causal role. The causal role of the experience will include the fact that it is normally caused by viewing red objects and that it is the cause of the appropriate kind of discriminatory behaviour. If I report that I am having an experience of red, he or she has no reason to distrust me, and my behaviour seems appropriate, then he or she will have a justified belief that I am having an experience of red.

If a congenitally blind person can understand a phenomenal statement just by thinking of the experience of red as the actual occupant of a certain causal role, why must sighted people also be able to imagine the token experience? I seem to be demanding more of the sighted. There are two ways in which we may think of this. On one view, in order to understand a statement one must integrate it with our basic abilities to appreciate what is being said by the statement. If a subject has the ability to imagine an experience, then this is one of the things that must be integrated. Failure to integrate indicates that the subject does not fully understand the statement. Obviously, if one does not have the ability, one has less to do. On a second view, the ability to imagine the experience is not constitutive of our understanding of the phenomenal statement concerning that experience. It is just that, if we do understand the statement by thinking of the experience as the actual occupant of a certain causal role, then our ability to imagine our experiences, and understand the role that they play in our mental lives, means that we should be able to imagine the relevant experience as well. In the absence of a fully worked out theory of understanding statements, I am inclined to take the second option. However, both strike me as having some plausibility.

Some might question whether (II) is correct. According to them, the following is more plausible.

(III) A phenomenal statement about what A's token experience is like can only be understood by another person, B, if B knows what experiences of that type are like.

Thus, on such a view, a congenitally blind person would not understand the sentence 'I saw something red'. Although I find such a view implausible for the reasons given

above, it is quite possible to endorse (III) and yet accept my overall line of argument. The crucial point is that (III), just like (II), does not require that a phenomenal statement convey what an experience is like to somebody who does not have the relevant concept. In which case, there is no reason to deny that the statement ‘I am having a visual experience of something red’ counts as a phenomenal statement.<sup>11</sup>

There are a number of reasons why the statements I have listed will not strike everybody as phenomenal statements of the relevant kinds. The most obvious is that they do not provide a subject with knowledge of what an experience of a certain type is like. As it stands, this charge is not quite right. They may provide subjects with propositional knowledge concerning what an experience of a certain type is like in much the same way that the congenitally blind person can have a justified belief that somebody is experiencing something red. They do not provide the ability to imagine the experience without essential recourse to extrapolation nor do they confer the experience upon the recipient of the statement. However, that just illustrates the point I am trying to make. Phenomenal statements need not confer knowledge of what an experience is like in this sense. They just state what we should be able to do—by specifying a possible content of experience—if we are to know what an experience is like. The same point holds for many other abilities. There are propositions concerning how we should ride a bike. Nevertheless, being able to ride a bike is not conveyed by knowledge of them.

A second concern is that the descriptions of my experiences given above are uninformative. Migraine involves certain disturbances of vision, a particular kind of headache, and so on. This is not spelt out in my statement that I have a migraine. So, it might be argued, the statement doesn’t describe what the experience is like. This charge seems mistaken. The statement that I have a migraine attributes to my experience a particular phenomenological pattern. The fact that such a statement does not spell out the pattern does not mean that it fails to provide a precise specification of what my experience is like. It is important to distinguish failing to spell out what an experience is like from failing to give a precise specification of what it is like.

An illustration of this point might help. I currently have a Brother HL-1250 printer. This statement gives a precise description of the way that the world is in a particular respect. There are all kind of details given by it. Nevertheless, I have not spelt out all the properties my printer has. When we state what an experience is like—for instance, that it is a migraine—there are all kinds of features that we have conveyed by this. So a phenomenal statement of this kind could scarcely be described as uninformative. One could only think that such a statement was uninformative, if one was tacitly committed to (I) rather than either (II) or (III). Once it is recognized that phenomenal statements don’t have to convey, to those who lack the relevant concepts, what an experience is like, the temptation to deny that ‘I am experiencing a migraine’ is informative should go.<sup>12</sup>

<sup>11</sup> I’m grateful to an anonymous referee for suggesting that something like (III) should be discussed. I would also like to clarify one other point that might cause confusion. The support for (II) or (III) in this section does not rely upon my denial, in the previous section, that phenomenal propositions about objects and properties in the world are really about experiences. Even if I agreed that they were implicitly about experiences, it still would not follow that they conveyed to somebody who had not had the relevant experiences, and was not able to imagine them, the ability to imagine the experiences. This issue was raised by an anonymous referee in further comments.

<sup>12</sup> I grateful to an anonymous referee for pressing me on this point.

A third concern is that there seems very little to say about what some experiences are like, for instance, an experience of red. If we are asked to describe more fully what it is like, we can only say that it is more similar to experiencing orange than blue, that the experience is of a warm colour, and so on. This does not exhaust what it is like to experience red. It seems that there should be many more things to say but there aren't. This might be taken to support the claim that what an experience of red is like is not a fact [Mellor 1993: 8].

The concern just expressed may just rely upon the thought that a statement concerning what something is like should convey knowledge of what it is like. If so, then it has been answered already. To the extent that there is a separate worry, it arises because experiences of red seem unanalysable, in contrast with migraines (say). We can see that the latter is made up of a number of different experiences. However, in the case of experiences of red, all we seem to have is the simple experience. If this is right, then that is no problem for the position I am defending. There does not cease to be a fact about what an experience is like because the only candidate fact is simple.

Probably the objection which will be the sticking point for many is the possibility that subjects experience things—red tomatoes (say)—in different ways. Spectrum inversion cases are just extreme instances of this according to which the way that red tomatoes look to A will be the way that green tomatoes look to B, and so on. When I state that I am having a visual experience of a red tomato, it may be argued, I cannot be stating what my experience is like because it does not convey whether I am experiencing red tomatoes in the way that you do.

The problem with this objection is that it raises a difficulty for my position only by conceding the essential issue at hand. If my experiences really are phenomenally different from yours in the way envisaged, then there are phenomenal facts. The problem is just one of stating them. One option would be to produce statements of the following form: *S is having an experience of something with that phenomenal property*. If a question is raised about whether there are statements of this form, then the first premise of the argument from inexpressibility is thrown into doubt. In any event, if there are phenomenal facts, there are propositions stating them whether or not they can be expressed. The conclusion of the argument from inexpressibility is false.

Before I turn to the further development of my own proposal about knowledge of what an experience is like and relate this to the knowledge argument, I want to mention briefly two complications which cut across the dialectic thus far. The first is that some of the phenomenal propositions I have supplied might seem to provide a recipe for determining what a subject's experience is like rather than expressing what it is like. One fairly uncontroversial example of the latter type of proposition would be: 'My feeling of pain is exactly what you would feel if you stabbed that pin in your finger.' Controversy over precisely which propositions fall into the category of recipe-providers rather than phenomenal propositions should be separated from the question of whether there are phenomenal propositions at all. I experienced the smell of coffee, I had a migraine, I visually experienced something red, I felt anxious, I felt a wave of nausea, clearly are not recipes for imagining what my experience is like. Each provides a specification of what my experience is like. One could only think otherwise if one ignored the points made above and expected something more from phenomenal statements. By contrast, 'I felt so angry I could barely speak' might fall on the recipe-providing side. For the sceptic, it

might seem that it gives, at best, a way of identifying what the experience is like by mentioning something which goes along with it. My own view is that there is a distinctive phenomenology involved in being so angry one could barely speak but if you disagree with me about this, it doesn't affect the general point.

The second complication is that one proponent of the Ability Hypothesis, David Lewis, allows that there can be facts about what experiences are like to capture the differences between two subjects whose colour experiences are inverted with respect to each other. He thinks that the phenomenal differences may arise from different intrinsic physical properties which, say, occupy the red experience, green experience role, and so on. However, he denies that there are any phenomenal propositions. This is because Lewis supposes that the Non-Physicalist must hold that phenomenal information cannot be learnt just by lessons, that which it concerns must be experienced [Lewis 1999a: 272–3]. This would not be so for the statements about intrinsic physical properties.

This characterization of the Non-Physicalist's position makes it look more implausible than it needs to be. The Non-Physicalist does not have to claim that experience is essential to obtaining phenomenal information. A Non-Physicalist could endorse the understanding of phenomenal statements I recommended above. The question is whether statements about physical properties are plausibly thought of as capturing phenomenal differences. If it is not plausible, the Non-Physicalist will claim that, if our knowledge is limited to the kind of information with which the physical sciences provide us, then someone will need to experience red (say) in order to find out the appropriate phenomenal facts. But that is the conclusion of an argument whose premises lie elsewhere. It does not depend upon any exotic understanding of the nature of phenomenal statements. Failure to appreciate this obscures the debate between the Physicalist and the Non-Physicalist over the exact force of the Knowledge Argument. I will try to make this clear in the last section of this paper. For now, I hope we can all agree that there is a perfectly legitimate sense in which there are phenomenal statements.

### III. The Role of Phenomenal Propositions in Knowing What It's Like

Even if, given what I have argued, knowledge of phenomenal propositions is not sufficient for knowledge of what an experience is like, some might feel it is necessary. While the provisional analysis I have offered does not explicitly mention phenomenal propositions, it might be argued that, if S can identify the experience he or she is currently having or imagining, then S knows the relevant phenomenal propositions. I shall explain why I don't think this is the case. In so doing, we will see that the provisional analysis offered needs some adjustment.

An essential but problematic element of the analysis is its appeal to the ability to identify one's experience. It is partly meant to capture the thought that, if a subject has an experience and yet is unable to appreciate it in any way (if indeed that is possible), then the subject cannot be attributed with knowledge of what the experience is like. Just as in the case of our ability to imagine, the ability to identify should be immediate. Suppose we are so absent-minded, our experiences are just flitting past. We need to haul ourselves away from this state in order to identify our experience. In such circumstances, we don't have the immediate ability to identify our experiences although we do have the ability to identify our experiences after a certain effort.

Michael Tye suggests that the subject should have cognised the experience and takes this to involve propositional knowledge *that one's experience is F* [1989: 140–1]. It strikes me that this is too strong in two ways. First, what matters is that subjects should be in a position immediately to obtain propositional knowledge concerning the experiences they are having *if* they seek to do so. That's enough for them to appreciate what their experience is like. Propositional encodement of what they have appreciated does not seem to be needed. Indeed, it is natural to think that it is because we know what an experience is like in the required sense that we can know propositions about the experience. Second, it does not seem necessary that subjects should be able to demonstratively identify their experience and its phenomenal properties or indeed have the concept of either of these things in order to know what their experience is like. We can know what it is like to experience a heaving sea breaking upon a rocky coast without having the kind of sophistication that enables us to focus on the phenomenal properties of the experience rather than the objects of experience in order to demonstratively identify them (for instance). Indeed, even those of us who pride ourselves on such sophistication often note that when we think about the phenomenal content of an experience, we end up focussing on the sea and the rocks. The experience is transparent.

So how should we characterize the ability to identify that we need? I suggest that, in order to know what an experience of type E is like, say the experience of the sea and rocks, we should be able to identify what an experience of that type is *of*, either on the basis of an experience of type E we are currently having or on the basis of an imaginative experience with the same content, an imagining of the sea and rocks. When we then introspect an experience and make a judgement about what we are experiencing we use the information we have as to what the experience is *of*—the sea and the rocks—and attribute to ourselves an experience of these things. In this sense, we conceive of the sea and the rocks as the phenomenal content of the experience. We refer to items in the world to describe what our experiences are like.

The proposal seems to account for what strikes me as a fairly well-entrenched belief about the range of subjects who know what their experiences are like. In attributing such knowledge to conceptually unsophisticated creatures, we don't insist that they should have grasped that their experiences have a phenomenal content or imagine not just the sea and the rocks but that they are having an experience of the sea and the rocks.<sup>13</sup> This would certainly be beyond bats and rabbits and very young children. Instead, they must merely have the experiences, or be capable of immediately imagining what they have experienced, and be able immediately to identify what it is that they are experiencing or imagining. Of course there is room for debate about precisely which creatures do know what their experiences are like. The fundamental point is that, if a creature lacks any concept of experience or the phenomenal properties that characterize it, it does not follow that they do not know what their experiences are like. If we required that somebody must know a phenomenal statement concerning what an experience is like in order to know what it is like, we threaten to deny this knowledge to creatures who lack these concepts. This is surely wrong. I am not denying that knowing a phenomenal statement about our experience may serve to characterize a conceptually sophisticated knowledge of what an

<sup>13</sup> For more on the proper way to characterize imaginative experience, see Noordhof [2002].

experience is like, but it seems to me implausible that this is the fundamental kind of knowing what an experience is like to which reference is made in discussions on this subject, and which is a pervasive part of our appreciation of our mental lives.

Questions still arise about this proposal. It might be wondered whether we appreciate what our experience is like just by being able to identify what is revealed in experience. After all, it is possible that we might experience the same things in different ways, e.g., by vision or by sonar.<sup>14</sup> There are a couple of responses to this worry. Those who are committed to Representationalism about experience may claim that, in fact, if two subjects experience things in different ways, that's because their experiences reveal a different range of properties. In which case, there is no problem. Although I think there is some plausibility to this view, I don't want to rest my account of knowing what an experience is like upon it. Instead, it is important to recognize that the ability to identify the things which we experience is keyed to the type of experience we have. My point is that by having the capacity to identify the sea and the rocks on the basis of a particular kind of experience, we know what that experience is like. That is quite compatible with supposing that other creatures might have a different kind of experience of the sea and rocks. If so, my ability to identify the sea and the rocks on the basis of experience would not be theirs.

A second worry relates to the factive character of identification. If I identify the sea on the basis of my experience, the sea's got to be there. If the sea is not there because, say, I am having an hallucination, I cannot identify the sea. Yet presumably I know what my experience is like. In response, I don't want to say that the ability to identify the sea is something I have regardless of whether I am currently having an experience of the sea. Nevertheless, when I have an hallucination of the sea, it also seems acceptable to note that I can exercise my ability to identify the sea in the sense that I can exercise what would be a sea-identifying ability if the sea were there. Let me call this sense of 'ability to identify' the non-factive sense. My account should be taken to appeal to it.<sup>15</sup>

The characterization of ability to identify I have just offered does not assume Representationalism about experience. For all that I have said, different subjects may experience the seascape in different nonrepresentational ways. All that I am assuming is that it is appropriate to suppose that *part* of the phenomenal character of an experience is that it is of objects and properties in the world. When one hallucinates the sea, it is still plausible to suppose that the experience one is undergoing may elicit the kind of identification I have mentioned concerning what it purports to be of. How else are we to make sense of subjects' claim that it seems to them as if the sea is before them? An appeal to the non-factive notion of identification is not an artifact of my proposal that we should characterize our knowledge of what an experience is like in terms of our ability to identify what the experience is of. Suppose it were insisted that the relevant ability to identify must concern phenomenal properties of experience. When we imagine the experience, we don't reproduce all the phenomenal properties of the experience in imagination. Indeed, there are different views about the extent to which they are reproduced at all [Martin 2002; Noordhof 2002]. Yet we would identify the experience we are imagining (with

<sup>14</sup> I'm grateful to an anonymous referee for raising this issue with me.

<sup>15</sup> I would like to thank an anonymous referee for raising this point.

phenomenal properties which might or might not be instantiated) rather than the phenomenal character of the imaginative experience.<sup>16</sup>

A third worry relates to the relationship between the ability to identify and knowing what an experience is like. The line of thought runs as follows. We have the ability to identify what is experienced *because* we know what the experience is like rather than vice versa. Since the explanation is this way around, I shouldn't seek to characterize knowledge of what an experience is like in terms of the ability to identify the objects of experience (amongst other things).<sup>17</sup>

It seems to me that three elements of my position enable me to defuse the concern. The first point to make is that, since I accept that there are phenomenal statements describing phenomenal facts about experience, I am not denying that there is something that it is like to have an experience. I can take this to enable a subject to identify the objects of experience. The ability is not, as it were, ungrounded. Second, once this is recognized, the issue is whether we need to take knowing what an experience is like as explanatorily more fundamental than phenomenal facts about experience *plus* the ability to identify the objects of experience on the basis of experience. It's not clear that we do. Why should the more fundamental notion be knowing what our experiences are like? Indeed, this seems to involve an unmotivated insistence that we must ascend a level and know what our experiences are like in order to be able to identify the objects of experience. Third, the intuition that there is an asymmetry—that we are able to identify the objects of experience *because* we know what our experience is like—can be explained by the fact that knowing what an experience is like involves a complex of abilities. There is some slight explanatory force in relating the exercise of an ability to the complex of which this exercise forms a part.

It might be questioned whether it is right to call our resulting notion *knowledge* of what an experience is like. Isn't it something rather more primitive? I think this fourth concern is also ungrounded. We are justified in calling it *knowing* what an experience is like because it has the following features. First, it involves the abilities to identify and imagine. Abilities constitute know-how. So it is no surprise that we should talk of knowing what an experience is like. Moreover, if someone has the ability to do such and such, that means they should reliably do such and such in the appropriate circumstances. There are good grounds for thinking that reliability about something is connected to our notion of knowledge in general. Third, and finally, our notion of knowledge of what an experience is like is connected to the notion of acquaintance. I will try to make this clearer in the next section of the paper. Before I do so though, I want to consider one final source of concern.

In a recent article, Jason Stanley and Timothy Williamson have argued that knowing how is a species of knowing that [2001]. In which case, it might be thought that, even if I have established that knowing what an experience is like is to be analysed in the way indicated, the ability to identify and imagine will involve knowledge of phenomenal statements. The first thing to say about this concern is that, at worst, I would have to concede that knowing what an experience is like requires knowledge of a phenomenal proposition concerning what the experience is of. This is not knowledge of a phenomenal

<sup>16</sup> These points are made in response to the further concerns of an anonymous referee.

<sup>17</sup> I would like to thank an anonymous referee for raising this point. I have discovered that a worry of this kind is also expressed in Gertler [1999: 323–4].



proposition about the experience. So the analysis stands. Nevertheless, I don't think I need to be even this concessive. I can agree with Stanley and Williamson that corresponding to every ability we might have, there is a proposition of the form *w is a way to do A*. The question is whether we are able to do *A* *because* we can think of the way or whether our ability to do *A* explains why we can think of the way. It seems to me that for some cases, the abilities to imagine and identify being two, the latter is true. Our ability enables us to pick out by demonstration the way that we should do something—just like that, we might say. It is not that we have prior articulation of the way. But if this is correct, then there is no problem for my approach. The analysis does not need to mention the relevant practical propositions. The connection between intentional action and employment of practical propositions is not sufficiently tight to force us to revise our verdict in these relatively basic cases of action [Stanley and Williamson 2001: 442–3].

#### IV. Knowing What It Would Be Like to Have a Content of Experience as the Content of Our Current Experience

I have canvassed a hybrid account of knowing what an experience is like. This might seem unsatisfactory. It has a cobbled together air. I suggest that what we really have in mind when we talk of knowing what an experience of a certain type *E* is like is knowing what it would be like to have a certain content of our experience as the content of an experience we are currently having. My proposal, which I now give in full, is a way of characterizing this knowledge.

*S* knows what an experience of type *E* is like iff (i) *either S* can immediately imagine the phenomenal content of an experience of type *E* *or S* is having an experience of type *E*, and (ii) *S* can immediately and non-factively identify the content of an experience of type *E*—the objects and properties it concerns—both when the content is imagined and when it is experienced.

We can now see how this is the coherent expression of a single idea. We know what it would be like to have the phenomenal content of an experience as the content of an experience we are currently having if we really are having the experience and can identify the objects and properties it concerns. So much is straightforward. However, we can also know what an experience is like through being able to imagine it. Imagination produces something phenomenally similar (although, as I have noted, not identical) to the phenomenal content of a perceptual experience. In being able to produce the relevant kind of imaginative experience and identify the objects and properties, which the perceptual experience would concern, in the imaginative experience, we reveal that we know what it would be like to actually have a perceptual experience with the phenomenal content in question.

To bring out the connection with acquaintance, I am tempted to say that the proper analysis of knowing what an experience is like is just an analysis of knowing what it would be like to be acquainted with the phenomenal content of an experience of type *E*. When I am having an experience of a heaving sea breaking on the rocks which I can identify as such, I am acquainted with the heaving sea breaking on the rocks. Being able to imagine the heaving sea breaking on the rocks is knowledge of what it would be like to

be acquainted with the sea and the rocks. Since there is a clear association of knowledge with acquaintance, we have an additional justification for my claim that my proposal is a proper analysis of *knowing* what an experience is like.

A health warning should be entered though. According to some, we are not acquainted with external objects but only sense data. When we are experiencing a table, it is still possible to doubt whether there is a table whereas it is not possible to doubt whether there is a sense datum of a table [Russell 1912: 25–7]. My use of the term ‘acquaintance’ is not so restricted. I take seriously John McDowell’s suggestion that we should disassociate the directness of acquaintance from the possibility of doubting whether there is an object with which we are acquainted [1986: 140–1]. On those occasions where we mistakenly take ourselves to be acquainted with an object, we are, in fact, being provided with knowledge of what it would be like to be acquainted with the object in question. A consequence of McDowell’s position is that we cannot always be certain whether we are receiving knowledge by acquaintance or knowledge of what it would be like to be acquainted.<sup>18</sup> This component of my approach is not essential. I mention it only to integrate it with other views of the subject matter. Those who resist this whole line can stick with my earlier formulation which does not mention acquaintance. Since perceptually experiencing something is, broadly conceived, knowing something, we do not lose the additional connection to our concept of knowledge I sought to draw out in terms of acquaintance.

In the previous section, I argued that my proposed analysis made sense of our attributions of knowledge of what an experience is like. In this section, I have argued that the analysis gives coherent expression to a single idea. Taken together, I think we conclude that we have found no reason to suppose that knowledge of what an experience is like requires knowledge of a phenomenal proposition.

#### V. Consequences for the Knowledge Argument and the Debate about Physicalism

The Physicalist and the Non-Physicalist agree that full physical information will not convey knowledge of what an experience is like. If the analysis defended here is correct, this does not imply anything damaging to the Physicalist. Since knowledge of what an experience is like involves abilities of the kind identified, it is implausible that full physical information should supply them. Moreover, since knowing what an experience is like does not involve knowledge of a phenomenal proposition, there is no pressure to say that one of the things which may have been learnt when a subject has an experience of a red tomato, say, is a phenomenal proposition. The proponents of the Ability Hypothesis engaged in overkill. They did not need to deny that there were phenomenal propositions. They presumably did so because they felt that, if there were, it would be inevitable that knowing what an experience was like would involve knowing a phenomenal proposition. I have sought to establish that this fear is misguided.

Of course, the following possibility is still open. Knowing what an experience is like is not all that a subject learns when they have an experience. They also learn a piece of phenomenal information. In which case, a reformulated knowledge argument would run as follows. In order to obtain knowledge of a phenomenal fact about a particular type of

<sup>18</sup> I’m grateful to an anonymous referee for raising this issue.

experience, someone must have experienced it. Full physical information about the brain (or, indeed, the environment) wouldn't suffice. If Physicalism were true, we could know a phenomenal fact about a particular type of experience on the basis of the relevant physical information. Hence, Physicalism is false. It strikes me that the right approach to this argument is to divide and rule. Physicalists can argue that the feeling that someone must have had the experience in order for any of us to know the phenomenal fact illicitly appeals to the following equivalence: S knows a phenomenal fact about an experience if and only if S knows what the experience is like (where S is the first person who obtained the knowledge). I have tried to establish that this is false in both directions of the biconditional. The mistake that proponents of the reformulated Knowledge Argument would be making is to take a truth about knowing what an experience is like to reveal something about the character of phenomenal propositions.

Unfortunately for those who like tidy worlds, things don't go all the Physicalist's way. The Anti-Physicalist can claim that it is plausible that someone needs to experience a red tomato in order to pass on the phenomenal information to the rest of us because it is transparently obvious that phenomenal facts are distinct from any physical facts. The response to this claim is to explain how phenomenal propositions are about physical facts after all. It may seem that they could not be. It is up to Physicalists to explain why this appearance is mistaken. So the real issue is not captured by the Knowledge Argument but rather by the discussion concerning the explanatory gap between physical properties and phenomenal properties. Let us abandon discussion of the former once and for all and focus on the latter (though not here). This involves a certain amount of preaching to the converted. Quite a number of philosophers have turned their attention to the problem of the explanatory gap at the expense of the Knowledge Argument. All I hope to have done is provide firm ground for the preferences some have already felt.<sup>19</sup>

*University of Nottingham*

Received: February 2001

Revised: January 2002

## REFERENCES

- Blackburn, Simon 1988. Attitudes and Contents, *Ethics*, 98: 501–17.  
 Carruthers, Peter 1986. *Introducing Persons*, London and Sydney: Croom Helm.  
 Conee, Earl 1994. Phenomenal Knowledge, *Australasian Journal of Philosophy*, 72: 136–50.  
 Geach, P. T. 1964. Assertion, *Philosophical Review*, 74: 449–65.  
 Gertler, Brie 1999. A Defense of the Knowledge Argument, *Philosophical Studies*, 93: 317–36.  
 Hale, Bob 1986. The Compleat Projectivist, *Philosophical Quarterly*, 36: 68–76.  
 Honderich, Ted 1988. *A Theory of Determinism*, Oxford: Clarendon Press.  
 Jackson, Frank 1982. Epiphenomenal Qualia, *Philosophical Quarterly*, 32: 127–36.  
 Jackson, Frank 1986. What Mary Didn't Know, *Journal of Philosophy*, 83: 291–5.  
 Lewis, David 1983. Postscript to 'Mad Pain and Martian Pain', in *Philosophical Papers, Volume 1*, New York: Oxford University Press: 130–2.

<sup>19</sup> My thanks to Bob Kirk and Hugh Mellor for enabling me to produce a much better paper than I would otherwise have been able to do. More recently, I would like to thank the invaluable comments of two referees for *Australasian Journal of Philosophy*. I very much appreciate the attention they gave to my paper.

- Lewis, David 1999b (1995). Should a Materialist Believe in Qualia? in *Papers in Metaphysics and Epistemology*, Cambridge: Cambridge University Press: 325–31.
- Lewis, David 1999a (1988). What Experience Teaches, in *Papers in Metaphysics and Epistemology*, Cambridge: Cambridge University Press: 262–90.
- Loar, Brian 1997 (1990). Phenomenal States, in *The Nature of Consciousness*, ed. Ned Block, Owen Flanagan, and Güven Güzeldere, Cambridge MA: MIT Press: 597–616.
- Lycan, William G. 1996. *Consciousness and Experience*, Cambridge MA: MIT Press.
- Martin, M. G. F. 2002. Transparency of Experience, *Mind and Language*, 17: 376–425.
- McDowell, John 1986. Singular Thought and the Extent of Inner Space, in *Subject, Thought, and Contest*, ed. Philip Pettit and John McDowell, Oxford: Clarendon Press: 137–68.
- Mellor, Hugh 1993. Nothing Like Experience, *Proceedings of the Aristotelian Society*, 93: 1–16.
- Nagel, Thomas 1979 (1974). What is it Like to be a Bat?, in *Mortal Questions*, Cambridge: Cambridge University Press: 165–80.
- Nemirow, Lawrence 1980. Review of *Mortal Questions* by Thomas Nagel, *Philosophical Review*, 89: 473–7.
- Nemirow, Lawrence 1990. Physicalism and the Cognitive Role of Acquaintance in *Mind and Cognition*, ed. William G. Lycan, Oxford: Basil Blackwell: 490–9.
- Noordhof, Paul 2002. Imagining Objects and Imagining Experiences, *Mind and Language*, 17 (2002), 426–55.
- Raymont, Paul 1999. The Know-How Response to Jackson's Knowledge Argument, *Journal of Philosophical Research*, 24: 113–26.
- Robinson, Howard 1982. *Matter and Sense*, Cambridge: Cambridge University Press.
- Russell, Bertrand 1912. *The Problems of Philosophy*, London: Oxford University Press.
- Stanley, Jason and Timothy Williamson 2001. Knowing How, *Journal of Philosophy*, 98: 411–44.
- Tye, Michael 1989. *The Metaphysics of Mind*, Cambridge: Cambridge University Press.