

There Are No Phenomenal Concepts

Derek Ball

04/08/09

Abstract

It has long been widely agreed that some concepts can be possessed only by those who have undergone a certain type of phenomenal experience. Orthodoxy among contemporary philosophers of mind has it that these phenomenal concepts provide the key to understanding many disputes between physicalists and their opponents, and in particular offer an explanation of Mary's predicament in the situation exploited by Frank Jackson's knowledge argument. I reject the orthodox view; I deny that there are phenomenal concepts. My arguments exploit the sort of considerations that are typically used to motivate externalism about mental content. Although physicalists often appeal to phenomenal concepts to defend their view against the knowledge argument, I argue that this is a mistake. The knowledge argument depends on phenomenal concepts; if there are no phenomenal concepts, then the knowledge argument fails.

Anti-physicalist philosophers of mind since Descartes have sought to ground a metaphysical distinction between mind and body by appealing to epistemic facts. (Recent examples include the knowledge argument (Jackson 1982, 1986) and arguments which exploit the explanatory gap, including modal arguments based on conceivability or rational intuition (Kripke 1972; Bealer 1994; Chalmers 1996), or on semantic features associated with the necessary a posteriori (White 1986, 2006).) Many physicalists now maintain that these arguments can be defused by a proper understanding of the concepts that we

Forthcoming in *Mind*.

use to think about the phenomenal characters of our experiences (Loar 1997, 2003; Harman 1990; Papineau 1993, 2002, 2007; Sturgeon 1994; Lycan 1996; Hill 1997; Balog 1999; Tye 2000, 2003; Perry 2001; Block 2003, 2006; Byrne 2006; Levin 2007). Although there is considerable disagreement about the nature of these concepts, it is widely held that that there exist *phenomenal concepts* that play a special role in our awareness of our experiences; according to David Papineau, ‘Phenomenal concepts are common coin among nearly all contemporary philosophers working on consciousness’ (Papineau 2007, p. 111).

In this paper, I argue that *there are no phenomenal concepts*, and therefore that this defence of physicalism cannot succeed. Nonetheless, my view should be attractive to physicalists. Although physicalists have typically developed theories of phenomenal concepts in order to resist the knowledge argument, I will show that the rejection of phenomenal concepts is itself a response to the knowledge argument: *if there are no phenomenal concepts, then the knowledge argument fails*.

The paper is divided into three sections. In section 1, I address some terminological issues and provide a precise statement of the relationship between phenomenal experience and possession of phenomenal concepts. Section 2 explores the role that phenomenal concepts play in the knowledge argument, including the way in which physicalists have used phenomenal concepts to respond to the argument. Section 3 is devoted to my case against phenomenal concepts.

1 Phenomenal concepts

1.1 Terminology

I will be using the terms ‘concept’ and ‘content’ as follows. *Contents* are the objects of *de dicto* propositional attitudes such as beliefs, desires, and thoughts. Beliefs and other propositional attitudes are mental representations, and they have an internal structure: they are composed of mental representations called *concepts*. For reasons that will become clear below, phenomenal concept theorists are committed to individuating concepts and contents in a fine-grained (Fregean) way. On a fine-grained conception, in normal cases two contents are distinct if one can rationally hold conflicting attitudes toward them at the same time. For example, on this view the con-

tent of my belief that Hesperus is a planet is not the same as the content of my belief that Phosphorus is a planet, since a person could rationally believe one while failing to believe the other.

I will follow phenomenal concepts theorists in assuming that concepts and contents are fine-grained. I will also assume (following (Burge 1979) and others) that concepts and contents must be individuated in a way that conforms broadly to everyday propositional attitude attributions. It follows from this assumption that concepts are *public*: it is possible for distinct thinkers to possess concepts of the same type.

1.2 Phenomenal concepts

Phenomenal concepts are typically held to be concepts that can be formed only by introspection of phenomenal experiences. The basic idea is familiar from the British empiricists; for example, phenomenal concepts are not unlike Locke's 'simple ideas' which the mind can obtain only by 'sensation' and Hume's 'ideas' that are but feeble copies of 'impressions'. In particular, phenomenal concept theorists typically hold that one can possess a given phenomenal concept only if one has undergone a certain type of experience; they thus hold with Locke that at least with respect to phenomenal concepts, it is impossible for anyone to 'frame the idea of a scent he had never smelt' (Locke 1690, p. 100 (Sect. II.2)), and similarly with Hume that 'A blind man can form no notion of colours; a deaf man of sounds' (Hume 1748, p. 98 (Sect. II)).

Can we give a precise description of this purported phenomenon? The following suggestion may appear promising:

The concept C is a phenomenal concept only if:

1. There is some phenomenal experience type e , such that C refers to e .
2. A person can possess C only if she has had an experience of type e .

As it stands, however, this proposal is unsatisfactory, since some phenomenal concepts theorists have held that in unusual circumstances it may be possible to come to possess phenomenal concepts without having had the relevant experiences. For example, Tye (2000, p. 27) holds that one might come

to possess phenomenal concepts due to neurosurgery or a miracle. It might also be held that it is possible to have a phenomenal concept of the missing shade of blue (Block 2003; Loar 1991, Sect. 4), or that a newly materialized Swampman has phenomenal concepts, or that Martians are born with innate phenomenal concepts.

These cases can be accommodated by relaxing the proposed necessary condition as follows:

The concept C is a phenomenal concept only if:

1. There is some phenomenal experience type e , such that C refers to e .
2. *Under normal circumstances, a human being can possess C only if she has had an experience of type e .*¹

But even this weakened condition is unsatisfactory, because some phenomenal concept theorists deny that phenomenal concepts refer to types of experience (Harman 1990; Tye 2000, 2003). Instead, these theorists hold that phenomenal concepts apply to experienced properties of objects. For example, on this view, the phenomenal concept RED_P refers to a property of ripe tomatoes and fire engines, rather than to a property of the experiences which are typically caused by these objects.

I will not attempt to adjudicate this dispute here. Instead, I will take my target to be the following condition, which is neutral between the two views:

The Phenomenal Concepts Criterion (PCC) The concept C is a phenomenal concept only if:

1. There is some phenomenal experience type e , and some property p , such that experience tokens fall under e in virtue of their relation to p .
2. C refers to p .
3. Under normal circumstances, a human being can possess C only if she has had an experience of type e .

¹A referee raises the possibility that the neurosurgery envisioned by Tye might become the norm. But since we are looking for a necessary condition rather than a definition or analysis, we can restrict our attention to current, actual norms.

Those who hold that phenomenal concepts refer to intrinsic qualities of experiences will take the relation mentioned in clause (1) to be instantiation. Those who hold that phenomenal concepts refer to experienced properties of objects will probably take the relation mentioned in clause (1) to be some species of representation. My arguments will apply to either view. Similarly, I intend my usage of terms like ‘experience of red’ and names of concepts like ‘RED’ to be neutral between the two views.

Many phenomenal concept theorists explicitly endorse (PCC) (or a similar criterion).² For example:

- ‘one cannot possess a predicative phenomenal concept unless one has actually undergone token states to which it applies.’ (Tye 1995, p. 169)
- ‘The crucial feature of phenomenal concepts, for the purposes of this argument, is that they are experience-dependent: the concept’s acquisition depends on its possessor having previously undergone the experience it refers to.’ (Papineau 2007, pp. 126-7)
- ‘since [phenomenal] concepts are canonically linked to subjective experience, one fails to wield such concepts if one fails to have such experience.’ (Sturgeon 1994, p. 231)
- ‘A phenomenal concept is individuated with respect to fundamental uses that involve the actual occurrence of phenomenal properties. In these fundamental uses, an actually occurring experience is used to think about that very experience. No one could have a phenomenal concept if they could not in some way relate the concept to such fundamental uses in which the subject actually has an instance of the phenomenal quality.’ (Block 2006, Sect. 1)
- ‘what it is to have a given concept is functionally determined. Someone has the appropriate concept of something’s being red if and only if the person has available a concept that functions in the appropriate way. [...] In this case, connections to perceptual input are crucial. If the concept is to function in such a way that the person has the full concept of something’s being red, the person must be disposed to form representations involving that concept as the natural and immediate

²In addition, (PCC) is taken to be an essential aspect of the theory of phenomenal concepts by Byrne (2006) and Stoljar (2005).

consequence of seeing something red. Since the blind person lacks any concept of this sort, the blind person lacks the full concept of something's being red.' (Harman 1990, p. 670)

Although some phenomenal concept theorists do not explicitly endorse (PCC), the most popular theories of phenomenal concepts straightforwardly commit their adherents to it. According to one view, phenomenal concepts are *recognitional concepts* whose reference is fixed by their users' abilities to recognize phenomenal qualities in experience (Loar 1997). Since humans normally have no ability to recognize phenomenal qualities that they have not experienced, those who hold that phenomenal concepts are recognitional concepts are committed to (PCC). On another view, phenomenal concepts refer via a sample of an experienced property (Block 2006). This 'quotational' view is also committed to (PCC), since in normal cases such a sample could do its job only by being experienced.

Moreover, there is reason to suppose that (PCC) will be an essential aspect of any theory of phenomenal concepts. Phenomenal concepts are meant to provide an explanation of the situation exploited by Frank Jackson's (1982; 1986) famous knowledge argument. (Indeed, this was the primary use of phenomenal concepts in Loar's (1997) and Harman's (1990) seminal discussions.) Jackson imagined a brilliant scientist named Mary who possesses all of the physical information about colour vision, but who has been imprisoned since birth in a black-and-white room and thus has never experienced colour. Jackson argued that Mary could not know what it is like to experience red while she is in the room, and that she learns something when she leaves. Phenomenal concept theorists typically accept this aspect of Jackson's position, and claim that when Mary leaves her room, she gains a phenomenal concept of red that she could not possess in her room. This explains her ability to think new thoughts. (I explain this strategy in more detail in Sect. 2.2 below.) Clearly, this strategy can succeed only given (PCC). I therefore conclude that phenomenal concept theorists are committed to (PCC).

2 Phenomenal concepts and the knowledge argument

2.1 Why the knowledge argument needs phenomenal concepts

Jackson claimed that Mary's new knowledge refutes physicalism: since Mary knows all of the physical facts in her room, any knowledge that she gains when she leaves must be knowledge of non-physical facts.³ Following Martine Nida-Rümelin (2002), we can explicitly formulate Jackson's argument in the following way:

Premiss P1 Mary has complete physical knowledge about human colour vision before her release.

Consequence C1 Therefore, Mary knows all the physical facts about human colour vision before her release.

Premiss P2 There is some (kind of) knowledge concerning facts about human colour vision that Mary could not have before her release.

Consequence C2 Therefore (from (P2)), there are some facts about human colour vision that Mary could not know before her release.

Consequence C3 Therefore (from (C1) and (C2)), there are non-physical facts about human colour vision.⁴

Physicalists must deny (C3). Many contemporary physicalists attempt to use phenomenal concepts to effect this denial. But anti-physicalists also need phenomenal concepts: without phenomenal concepts, the knowledge

³Other proponents of the knowledge argument include Robinson (1993a,b), Chalmers (1996, 2004), Alter (1998), Nida-Rümelin (1998), Gertler (1999), and Holman (2006). There are also a number of would-be physicalists who admit that they know of no good response to the knowledge argument, including Graham and Horgan (2001) and Hellie (2004). For a history of related arguments, see Stoljar and Nagasawa 2004.

⁴An alternative version of the knowledge argument, presented by Jackson in later work (e.g. Jackson (1995)) and developed by Chalmers (1996, 2004), attempts to use Mary's situation to show that the phenomenal facts are not apriori entailed by the physical facts. I will not discuss this version of the knowledge argument here, since the physicalist views on which I am focusing accept its conclusion. Nonetheless, many of the points I make in this section could be adapted to this version of the knowledge argument.

argument fails. Specifically, I will argue that phenomenal concepts offer the only explanation of (P2) on which it is at all plausible both that (P1) is true and that (C2) follows from (P2).

If (P2) is true, then there is some belief content B that Mary cannot know in her room. If Mary cannot know that B in her room, then one of the following must be the case: either (i) in her room, Mary can entertain B , wonder whether B , possibly even believe that B , but this belief would not be knowledge; or (ii) Mary cannot even entertain B until she leaves her room. (ii) is the most natural account of Mary's situation, as it is usually described. But we can imagine cases which are better described by (i). I will now argue that no version of the Knowledge Argument can be developed on the basis of a situation like (i), because no situation like (i) can be described in which (P1) is true.

The distinction between (i) and (ii) is made vivid by Nida-Rümelin's (1996; 1998) Marianna case.⁵ Like Mary, Marianna is a brilliant scientist who knows all of the physical facts about colour vision, but has grown up imprisoned in a colourless room. Unlike Mary, Marianna is not released directly into the outside world. Instead, she is first led into an empty room which is decorated with random splashes of coloured paint. The coloured splashes are not labelled, and Marianna is neither given information about their chemical composition or reflectance properties, nor allowed access to scientific equipment with which she could obtain this information.

There is some temptation to think that there are facts that Marianna still could not know: she has seen colours, but it may be that she could not know which colour is red and which is green. Whether or not this is so, it seems that Marianna has made epistemic progress. Marianna can at least ask questions that Mary could not. For example, Marianna might introspectively attend to an experience of red and ask herself whether *this* is what it's like to see red. Call the content of Marianna's query Q . It is plausible that Mary, on the other hand, cannot wonder whether Q , since she cannot attend to an experience of red in introspection.

If we suppose that Marianna could not come to know that Q in her new room, then it is clear that Marianna is in a situation of type (i) with respect to Q . It may seem possible to develop a version of the knowledge argument on the basis of type (i) situations. In fact, however, it is by no means obvious that a type (i) situation based version of the knowledge argument can be

⁵A similar case is developed by Warner (1986).

developed in which the analogue of (P1) is true. This is because it does not seem possible to develop a type (i) situation in which the protagonist knows all of the physical facts about colour vision and all of the physical facts about her environment. A protagonist who knew these facts could deduce the relevant phenomenal truths. For example, it is only plausible that Marianna could not come to know Q because she is denied access to physical information about her environment: in particular, she knows neither the chemical composition nor the reflectance of the coloured splashes. If she did know these physical facts, she could infer Q . But if no type (i) situation can be developed in which the protagonist knows all of the physical facts, then no knowledge argument can be developed on the basis of a type (i) situation.⁶

It appears, then, that the knowledge argument must appeal to situations of type (ii). So there must be some belief content B that Mary cannot even entertain in her room. I can think of only three possible explanations of Mary's inability to entertain B in her room:

1. B involves some concept that Mary cannot come to possess in her room.
2. Although Mary possesses all of the concepts necessary to think that B , for some reason she cannot combine these concepts in the appropriate way while she is in her room.
3. Although Mary possesses all of the concepts necessary to think that B , some of these concepts are *indexical*: they express B in some contexts but not in others. Every context in which Mary could use these concepts in her room is such that these concepts do not express B .

(2) can be ruled out immediately: Mary is a properly-functioning cognitive agent, and there is no reason that she could not exercise her concepts freely. (3) is more plausible. But (3) provides no comfort to the anti-physicalist, since if (3) is correct, then (C2) does not follow from (P2); there is no obvious reason that the indexical concepts in question should not refer to physical entities.

⁶Rob Koons suggests the following modification of the argument: Q is in some natural sense a *general* fact. Marianna knows all the general physical facts, but she does not know Q . Therefore Q must be a non-physical general fact. But if my contention that Marianna could infer Q if she knew all of the particular physical facts is granted, the claim that Q is physicalistically unacceptable seems unwarranted.

The only remaining option is (1). But (1) amounts to the claim that there are concepts that satisfy (PCC); that is, that there are phenomenal concepts. I conclude, therefore, that the proponent of the knowledge argument must accept phenomenal concepts.⁷

2.2 The physicalist response

I have argued that the proponent of the knowledge argument must accept the following claims:

1. (P2) is true.
2. (1) is true because there is some mental content B that Mary can entertain only after she leaves her room (i.e., because Mary is in a type (ii) situation).
3. (2) is true because B involves a phenomenal concept (which satisfies (PCC)).

Physicalist phenomenal concept theorists agree with all of these claims. These physicalists attempt to block the knowledge argument by giving accounts of phenomenal concepts and phenomenal knowledge according to which (C2) does not follow from (P2). According to these accounts, we should distinguish between fine-grained *contents* and coarse-grained *propositions* and *facts*. On this view, there can be many contents that express the same proposition and are made true by the same fact. For example, the content of the belief that Mark Twain is a writer is distinct from the content of the belief that Samuel Clemens is a writer, even though the fact that Mark Twain is a writer = the fact that Samuel Clemens is a writer.

Given this distinction, it is possible for the physicalist to admit that Mary comes to know a new content while denying that she comes to know a new fact. For example, suppose that RED_p is a phenomenal concept and RED is a non-phenomenal concept, and that RED and RED_p have the same extension. Then the physicalist can say that Mary already knows that tomatoes typically

⁷ Note that I am not arguing that the dualist needs phenomenal concepts to support the claim that there is some thought that Mary cannot entertain in her room; this claim is independently plausible. Instead, I have tried to show that the dualist needs phenomenal concepts because phenomenal concepts provide the only explanation of this claim that supports dualism.

cause experiences of red, and upon leaving her room, learns that tomatoes typically cause experiences of red_p. These are different contents, but they are made true by the same fact. Thus although there is knowledge that Mary lacks in her room, no fact is such that Mary doesn't know it in her room. Thus (C2) does not follow from (P2), and the knowledge argument fails.

Many physicalists have found this account attractive, especially since phenomenal concepts are held to provide a response to other arguments based on the explanatory gap (Loar (2003), Tye (2000, Ch. 2)). In what follows, I will argue that phenomenal concepts are in fact highly problematic, and that physicalists are better served by denying that such concepts exist. I begin by considering some well-known arguments for externalism about mental content.

3 Concept possession

3.1 Concepts and conceptions

Concepts are mental representations which are components of thoughts. A *conception* is the collection of beliefs associated with a concept. Thus my concept of dogs is the mental representation DOG, and my conception of dogs is of furry, four-legged animals, of a certain shape and size, typically kept as household pets, and so forth.⁸

The following principle is a consequence of this view of concepts (and an extremely natural principle in any case):

Concept Possession (CP) *s* possesses the concept *C* if *s* is able to exercise *C* in her thoughts.

Examples made famous by Hilary Putnam (1970; 1975a) and Tyler Burge (1979; 1986) suggest that it is possible for a thinker to possess a concept *C*, even though her conception of *c* is extremely vague or inaccurate. For example, Putnam claimed that he could not distinguish beech trees from elm trees. Although he did not have a rich conception of beeches and elms, it is very plausible that he could still exercise the concepts BEECH and ELM

⁸Some theories collapse the distinction between concept and conception. The arguments of Putnam and Burge that I discuss below show that these theories are highly problematic. And there are further reasons to think that these theories are implausible; see (e.g.) Fodor 1998.

in thought. For example, he could have believed that beeches are trees, or that there are elms in the park. By (CP), then, Putnam possessed the concepts BEECH and ELM. Similarly, Burge argued that a thinker who has a belief that he would express with the English sentence, ‘I have arthritis in my thigh,’ could exercise the concept ARTHRITIS in his thoughts. Such a thinker would have a radically false conception of ARTHRITIS, but could nonetheless possess the concept.

Anyone who wishes to deny that it is possible for a person with a vague or inaccurate conception of *c* to possess the concept *C* must provide an alternative account of these cases. For example, on such an account the belief that Putnam would express with the sentence, ‘Beeches are trees,’ cannot be the belief that beeches are trees, since the belief that beeches are trees involves the concept BEECH. Similarly, on such an account the belief that Burge’s Arthritis-man would express with the sentence, ‘Arthritis can occur in the thigh,’ cannot be the belief that arthritis can occur in the thigh, since this belief involves the concept ARTHRITIS.

In cases involving very young children, the extremely confused, or the insane, we may upon reflection decide that the content of the subject’s belief is indeterminate, or even that the subject has no relevant belief at all. But this is entirely implausible in the cases at issue. Surely Putnam expresses *some* determinate belief when he says, ‘Beeches are trees.’ Similarly, it is difficult to deny that Arthritis-man expresses some determinate belief when he says, ‘Arthritis can occur in the thigh.’

Burge therefore suggests that his opponent’s most promising strategy is to hold that Arthritis-man has a true object-level belief that does not involve the concept ARTHRITIS (e.g., the belief that tharthritis can occur in the thigh, where tharthritis is a disease that can occur in joints and in thighs), in combination with the false metalinguistic belief that the sentence ‘Arthritis can occur in the thigh’ is typically used to express the belief that tharthritis can occur in the thigh.

Burge gives three⁹ highly persuasive arguments against this strategy.¹⁰ (I will focus on Burge’s ARTHRITIS case, but the arguments would apply *mutatis mutandis* to Putnam’s BEECH/ELM case.)

⁹In fact, Burge gives several other arguments that I ignore because they are not easily adaptable to phenomenal concepts.

¹⁰Several recent philosophers have endorsed something like this strategy (e.g. Crane (1991); Donnellan (1993)). Somewhat surprisingly, to my knowledge none of these philosophers has engaged Burge’s arguments.

1. If the strategy is to succeed, it must establish that Arthritis-man does not possess the concept ARTHRITIS. But it does not establish this. The claim that Arthritis-man possesses the concept THARTHRTIS and has a false metalinguistic belief is entirely compatible with the claim that Arthritis-man also possesses the concept ARTHRITIS.
2. If Arthritis-man possesses an idiosyncratic concept rather than the concept ARTHRITIS, then he cannot share beliefs with others. But this result is strange. For example, Arthritis-man's doctor might believe that Arthritis-man has arthritis in his ankles. Arthritis-man might agree. But if Arthritis-man does not possess the concept ARTHRITIS, it is hard to see how this could be so: the doctor believes that Arthritis-man has arthritis in his ankles, while Arthritis-man believes that he has tharthrtis in his ankles.

It is particularly difficult to see how Arthritis-man could share beliefs with thinkers who do not speak English. At least Arthritis-man and his doctor both believe that he has the disease called 'arthritis' in his ankles. A non-English speaker would have no such belief; indeed, it would seem that Arthritis-man and a non-English speaker could have no arthritis-related mental contents in common. But again, this is difficult to accept. It is very plausible that Arthritis-man and a monolingual Chinese doctor might both believe that arthritis is a disease.

3. Most importantly, the strategy makes wildly inaccurate predictions about Arthritis-man's reasoning processes when his usage is corrected. For example, if the strategy were correct, one would expect that when told Arthritis-man is told by his doctor that arthritis cannot occur in the thigh, he would object that his belief was true, and that the doctor is just quibbling about the meanings of words. But this is absurd: Arthritis-man's belief is false even by his own lights.

I find these arguments decisive. Unless there is some other way to avoid attributing the concept ARTHRITIS to Burge's character (and similarly, the concept BEECH to Putnam), we must admit that it is possible for a person to possess a concept in spite of a limited or inaccurate conception.

Arthritis-man's own non-metalinguistic beliefs and abilities do not fix the extension of his concept ARTHRITIS. His conception is too inaccurate for that. In virtue of what, then, does Arthritis-man possess the concept

ARTHRITIS? A natural view is that he possesses this concept in virtue of his interactions with his linguistic community. He allows the extension of his word, and of his concept, to be determined by the practices of the community as a whole.

Concepts that can be possessed in virtue of such community interaction are often called *deferential*.¹¹ One mark of deferentiality is the possibility of cases (like the ‘arthritis’ case) in which a subject possesses a concept despite an inaccurate conception. In such a case, it is unclear what would make it the case that the subject possesses the concept in question if not interactions with her linguistic community.

3.2 Phenomenal concepts and concept possession

The phenomenal concept theorist is committed to the claim that there are phenomenal colour concepts that Mary cannot possess in her room. In this section, I argue that this claim cannot be sustained, and thus that the phenomenal concept theory is false. My strategy is to describe two scenarios in which Mary possesses the concepts we use to think about our experiences in her room. In the first scenario, she possesses these concepts by interacting with normal perceivers linguistically. I defend this scenario in three stages. First, I argue that the concepts we use to think about our experiences can be expressed in language. Second, I argue that these concepts are deferential in the sense described above. Finally, I defend the claim that Mary could possess these concepts by advancing analogues of Burge’s arguments. I then describe the second scenario, in which Mary possesses the concepts we use to think about our experiences in virtue of her scientific theorizing, despite lacking a linguistic community.

3.2.1 Preliminary data

One thing that Mary learns can be expressed by the following English sentence:

- (1) That is what it’s like to see red. (Where ‘that’ refers to an experience of red, to which Mary is attending in introspection, or to some feature of such an experience.)

¹¹I intend my usage of this term to leave open the possibility that all, or nearly all, concepts are deferential.

The phenomenal concept theorist must claim that some word or phrase in this sentence expresses a phenomenal concept. (I consider the claim that phenomenal concepts cannot be expressed by words below.) There appear to be three options: ‘that’, ‘red’, and ‘what it’s like to see red’, and there is no consensus among phenomenal concept theorists as to which option is correct. I will remain neutral on this issue. But I will point out one constraint on an acceptable answer: if ‘that’ expresses an indexical concept (e.g., a concept that refers to whatever its user is introspectively attending to), then ‘that’ does not express a phenomenal concept. This is because Mary could possess such an indexical concept in her room. Thus I take it that even if it is the word ‘that’ which expresses a phenomenal concept, the concept in question could be expressed by a non-indexical word (though I return to this issue below.)

In order to have a precise target, I will suppose that the phenomenal concept theorist holds that the word ‘red’ in (1) expresses a phenomenal concept: the concept RED_P . My arguments are easily applicable *mutatis mutandis* if some other word is held to express a phenomenal concept.

Consider the following English sentences:

- (2) Ripe tomatoes typically cause experiences of red.
- (3) What it’s like to see red resembles what it’s like to see black more than it resembles what it’s like to hear a trumpet playing middle C.
- (4) If x is a number then x is not what it’s like to see red. (Stoljar 2005, p. 479)
- (5) That is not what it’s like to see red. (Where ‘that’ refers to some experience which the speaker is introspecting, or to some feature of such an experience.)
- (6) Seeing red is a phenomenal state.

If the word ‘red’ expresses the phenomenal concept RED_P in (1), then it is very plausible that ‘red’ expresses RED_P when (2)–(6) are used by a person with normal colour vision and a history of colour experiences to express her beliefs.

Mary could also use (2)–(6) to express beliefs before she leaves her room. But if RED_P is a phenomenal concept, then (by (PCC)) Mary cannot possess RED_P . The phenomenal concept theorist must therefore claim that Mary

possesses some other, non-phenomenal concept RED, along with the metalinguistic belief that the word ‘red’ expresses the concept RED. The following sort of argument might be offered in defence of this claim:

1. Mary’s concept RED lacks some features that my concept RED_P has: for example, RED was not caused by experiences of red, is not linked to images of red, and does not enable Mary to recognize red objects.
2. Therefore, there is no significant concept type of which RED and RED_P are both tokens.

Given the conclusion of the previous section, it is clear that this argument is invalid. Consider an analogous argument:

1. Putnam’s concept BEECH lacks some features that my concept BEECH_P has: for example, Putnam’s concept was not caused by experiences of beeches, is not linked to images of beeches, and does not enable Putnam to recognize beeches.
2. Therefore, there is no significant concept type of which Putnam’s concept BEECH and my concept BEECH_P are both tokens.

3.2.2 Linguistic expression

The BEECH argument is invalid: its premiss is true, but its conclusion is false. It might be claimed that RED_P is individuated in a distinctive way that makes the argument valid. In the absence of further argument, this claim would be obviously *ad hoc*. But it might appear that further argument is at hand; in particular, it might be claimed that since phenomenal concepts depend on the nature of a thinker’s phenomenal experience, they are not expressible by words in a public language. For example, Ned Block writes:

*[S]uppose that spectrum inversion is rife: there is no spectrum of the vast majority, despite functional similarity in relevant respects. How could I justify the claim that the qualitative content of my experience as of red things is *looking red*, whereas yours is, say, *looking blue*? Any argument I could use against you could equally well be used by you against me. And it will not do to say we are both right, each with our own sense of ‘red’. ‘Red’ is a public language word, not a word that means one thing in your mouth and another thing in mine. (Block 1990, p. 678)*

Block's view is a problem for the claim that the relevant concepts are deferential, since a concept is deferential only if it is expressible in a public language. If phenomenal concepts are not expressed by words in public language, then it is not obvious that Burge's arguments apply to them.

There are reasons to resist the claim that phenomenal concepts cannot be expressed in public language. First, as I pointed out above, it is a datum that what Mary learns is expressed by the sentence, 'That is what it is like to see red.' Since the phenomenal concept theorist holds that what Mary learns involves a phenomenal concept, it is hard to see how the claim that phenomenal concepts are not expressible in public language could be sustained. Perhaps the phenomenal concept theorist would hold that the phenomenal concept can only be expressed by an indexical word (even though the concept itself is not indexical), and that Burge's and Putnam's arguments do not apply to concepts expressed by the relevant sort of indexical word. But what could motivate this view? Surely someone who can express a particular non-indexical concept by the word 'that' could also express the same concept with some other word; for example, she might simply stipulate that 'red-p' is to express the concept she just expressed by 'that'. Others could learn the word 'red-p' and Burge's arguments would apply.

Indeed, much of the phenomenal concepts literature depends on the possibility of introducing a word in this way. For example, Block himself writes:

Consider a specific phenomenal property, Q, e.g. the property of feeling like the pain I am having right now. [...] 'Q' in my terminology is very different from 'Q_R' in Perry's terminology since 'Q_R' is a term that Mary understands in the black and white room. 'Q' by contrast is meant (by me even if not by Perry and Smart) as the verbal expression of a phenomenal concept. (Block 2006, p. 252)

'Q' is not a word of ordinary English; but this fact is irrelevant to the success of Burge-style arguments. Consider an analogue of Mary who has never experienced pain. Block must claim that Pain-Mary could not come to understand 'Q' and possess the concept it expresses; but surely she could come to have beliefs that she would express with sentences like 'Q is painful', and surely she might come to defer to a relevant community (i.e., to Block himself). Very plausibly, we would attribute to her the belief that Q is painful. But as we have seen, this is sufficient for concept possession in typical cases.

3.2.3 Deference

Thus we should conclude that the claim that phenomenal concepts cannot be expressed linguistically is implausible. But then it is simple to show that RED_P is deferential in the sense described above; that is, to show that the concept can be possessed even by those who have an inaccurate conception, in virtue of their relations to their linguistic communities. Several sorts of case can be used to show this. Most straightforwardly, some people over-extend their colour concepts. For example, some people are disposed to judge that some orange objects are red (when they are not in fact red). They are disposed to apply their red-concepts to prototypically red objects, and to rescind their idiosyncratic judgements when confronted with the objections of fellow speakers. We are inclined to attribute beliefs that involve the concept RED to such speakers, and rightly so: the case is very much like the ‘arthritis’ case, and similar arguments apply (Burge 1979, p. 110).

Some proponents of phenomenal concepts might object that while this style of argument is successful as regards concepts that are correctly applied to properties of physical objects, it fails as regards phenomenal concepts, since these are correctly applied to intrinsic qualities of experiences. But there are even more striking cases of inaccurate conceptions regarding such concepts. For example, some philosophers hold that there are no experiences. If they are consistent, these philosophers will withhold application of concepts that are correctly applied to qualities of experiences. But we would still regard them as possessing the concepts in question: after all, they disagree with us.

Similar arguments can be developed as regards other candidate phenomenal concepts. For example, consider the concept PAIN. It is possible to possess this concept despite having an inaccurate conception of pain. For example, the belief that nausea is a type of pain is surprisingly widespread. Those subject to this belief can agree and disagree with others, and in many cases would be willing to change their views in response to correction. We attribute to them beliefs using the word ‘pain’, and it is very plausible that we should regard them as possessing the PAIN, the very same concept as normal speakers and scientific experts.

The phenomenal concept theorist must hold that Mary cannot possess phenomenal concepts in her room. But if the concepts we use to think about our experiences can be possessed in virtue of one’s relations to one’s linguistic community, then there is a route by which Mary could come to possess them.

This conclusion is confirmed by Burge-style arguments:

First Argument: Suppose that before she leaves her room, Mary thinks a thought that she would express with (7), and after she leaves her room, she thinks a thought that she would express with (8):

(7) I don't know what it's like to see red.

(8) I know what it's like to see red.

The thought expressed by (7) is the negation of the thought expressed by (8). But the phenomenal concept theorist cannot admit this, since on the phenomenal concept theorist's view, the thoughts expressed by (7) and (8) involve distinct concepts.

It might be objected that despite appearances, (7) is not the negation of (8); instead, they bear the same relation as the following pair:

(9) I don't know that Hesperus is bright.

(10) I know that Phosphorus is bright.

But the objection fails. After leaving her room, Mary might truly think:

(11) I used to wonder what it's like to see red, but now I know.

By contrast, imagine that I once possessed only the concept HESPERUS (but not the concept PHOSPHORUS). I often wondered whether Hesperus is bright. I have since come to possess the concept PHOSPHORUS and have learned that Hesperus is Phosphorus, and that Phosphorus is bright. Given that we are individuating concepts in a fine-grained way, this situation cannot be reported (*de dicto*) as 12:

(12) I used to wonder whether Phosphorus is bright, but now I know.

Thus there is a disanalogy between the phenomenal concept case and traditional Frege cases. In general, it is hard to see how (11) could be true if what Mary knows involves a concept that she did not possess in her room, since the complement clause ellipsis in (11) eliminates the possibility that 'what it's like to see red' is ambiguous in the two attributions.

Second Argument (Compare argument 2 from Burge, above.) The beliefs that Mary expresses with (2) – (6) could be shared by thinkers with a normal history of colour experience. But if normal thinkers possess only the concept

RED_P , and not the concept RED, then they cannot share Mary's beliefs. The phenomenal concept theorist must therefore hold that normal perceivers possess both concepts. On this view, there are two thoughts that normal perceivers can express with sentences like (2) – (6). But introspection renders this implausible. We are not aware of having two sets of thoughts. (An analogous manoeuvre would be laughable in the ARTHRITIS case.)

There are further reasons to doubt that we should admit distinctions among necessarily coextensive concepts that are inaccessible to introspection. One of the most significant motivations for the claim that mental contents are conceptually structured is that such structures 'enable one to capture a thinker's ability to relate different thoughts to one another according to rational inferential patterns' (Burge 1993, p. 309). But distinct, necessarily coextensive concepts that we cannot distinguish would not aid us in reasoning. It is difficult to see what such concepts could add to our understanding of a thinker's rationality.

(Further, even if normal perceivers did possess two red-concepts, this would not show that Mary possesses only one. Compare argument 1 from Burge, above.)

Third Argument (Compare argument 3 from Burge, above.) Finally, the phenomenal concept theorist's view does not adequately capture the thought processes that Mary undergoes when she leaves her room. According to the strategy, on leaving her room Mary gains not only knowledge of what it's like to see red, but a huge number of new pieces of knowledge. For example, she gains new knowledge corresponding to (2) – (6). But this is only the tip of the iceberg; in her room, Mary has exceedingly well-developed views about colours and colour experience. On the phenomenal concept view, she will thus gain an extremely large number of new beliefs when she leaves her room. But this is implausible. For example, Mary would not be surprised to discover that what it is like to see red_p is not a number. But according to the phenomenal concept strategy, this is something that she learns.

It might be objected that Mary could know in her room that there is some concept that normal perceivers express with the word 'red', and that the referent of this concept is not a number. This would explain why she is not surprised. But this strategy would overgeneralize: Mary could know in her room that there is something it is like to experience the referent of the concept that normal perceivers express with the word 'red', but she would still find the thought that *that* is what it is like to see red surprising.

In response to each of my arguments, a phenomenal concept theorist

might simply bite the bullet and hold that the concepts in question are more fine grained than we would have supposed. But the same move could be made in Burge’s ‘arthritis’ case. (Burge’s arguments do not show that his opponent’s view is *incoherent*, but that it is *implausible*.) Since the arguments in the two cases are so closely analogous, they stand or fall together. Thus anyone who accepts the Burge’s position in the ‘arthritis’ case should deny that there are phenomenal concepts.

I have focused here on Mary, and on phenomenal colour concepts. But it is clear that a similar story could be devised for any concept that is held to satisfy (PCC), and that similar arguments would apply. I therefore conclude that no concept satisfies (PCC): the phenomenal concepts theory has implausible consequences and should be rejected. Mary’s conception of red is impoverished in her room, since it is not associated with phenomenal images, memories, and so forth. But there are good reasons to think that Mary possesses the same concepts as the rest of us.

3.3 Possession without deference

According to a widely held view, Burge’s Arthritis-man possesses the concept ARTHRITIS in a very different way than his doctor: Arthritis-man possesses the concept only *deferentially* or *partially*, while the doctor possesses it *non-deferentially* and *fully*. Since I have argued that similar considerations apply to Mary’s possession of colour/experience concepts, it might be objected that Mary’s possession of the relevant concepts is merely deferential, and thus importantly different from relevant speakers. In particular, it might be suggested that the key idea of the phenomenal concept strategy can be captured by the following modification of (PCC):¹²

The Phenomenal Concepts Criterion* (PCC*) The concept *C* is a phenomenal concept only if:

1. There is some phenomenal experience type *e*, and some property *p*, such that experience tokens fall under *e* in virtue of their relation to *p*.
2. *C* refers to *p*.

¹²This response was suggested in correspondence by David Chalmers, and by an anonymous referee.

3. Under normal circumstances, a human being can *non-differentially* or *fully* possess *C* only if she has had an experience of type *e*.

In section 2.1, above, I distinguished type (i) situations (in which Mary can entertain but not know the relevant content) from type (ii) situations (in which Mary cannot even entertain the relevant content). (PCC*) seems to raise the possibility of new versions of the knowledge argument based on each sort of case. I will argue that neither argument succeeds, beginning with the argument based on a type (ii) situation.

On a plausible understanding of (PCC*), Mary can possess phenomenal concepts only given the presence of a suitable linguistic community. But then it may seem that we can run the knowledge argument simply by imagining Lonely Mary, who is like Mary except that she lacks a linguistic community, and has developed all of her scientific knowledge independently. On the view in question, Lonely Mary could not possess phenomenal concepts in her room. But, it might be argued, she could still know all of the physical facts. Since there would be contents involving phenomenal concepts that she could not entertain, and thus that she could not know, the knowledge argument would go through.

This argument fails for two reasons. First, the proposed view does not explain the data. There seem to be facts that non-lonely Mary does not know, and it seems that she would make a discovery on leaving the room. The fact that Lonely Mary lacks certain concepts would offer no account of non-lonely Mary's knowledge. And it seems likely that whatever account we give of non-lonely Mary could be applied to Lonely Mary as well, once she has gained the relevant concepts. So it is hard to see how Lonely Mary's alleged conceptual gap is relevant to the phenomenon of knowing what it is like.

Second, the claim that Lonely Mary could not possess the relevant concepts in her room is simply false. It is widely accepted that in general, concepts can be introduced by stipulation, especially in the context of a scientific theory. For example, we could imagine a person living in complete isolation from other beings who introduces the concept WATER by fixing its reference with a description like 'the substance composed of aggregates of H₂O molecules' (Burge (1982), McLaughlin and Tye (1998, pp. 300–2), Ball (2007)). This person might share our water-beliefs and other attitudes, even if she does not know enough chemistry to determine what water looks like.

Similarly, we can imagine Lonely Mary working to develop a detailed theory of the neural correlates of consciousness. She knows that brain state R would correlate to some phenomenal state, even though she has never been in brain state R. Burge-style arguments of the sort I developed above suggest that she could introduce the putatively phenomenal concept Q by stipulating that Q is to refer to what it's like to be in R. This could not be via deference, since by stipulation there is no one to whom she could defer. Thus it is not the case that a person who has not undergone colour experiences can possess colour-experience concepts only deferentially, since she could introduce such concepts by stipulation in the course of her scientific theorizing.

Of course, we could imagine that Lonely Mary fails to make such stipulations. But it is unclear that the physicalist should grant that in this scenario Lonely Mary would know all of the physical facts. It is true that she could know all of the facts of microphysics. But compare a thinker who knows all of the facts of microphysics but for whatever reason fails to possess the concept TABLE. Such a thinker would fail to know the table facts. But surely this is irrelevant to whether tables are physicalistically acceptable. In general, it is unclear that such artificial restrictions on the sorts of concepts that a thinker possesses provide an acceptable starting point for delineating the physical and the non-physical.

Thus the knowledge argument cannot be defended on the basis of (PCC*) and a type-(ii) situation. What about a type-(i) situation? I argued above that no successful version of the knowledge argument could be based on a type-(i) situation like the Marianna case. But (PCC*) may seem to make available another sort of type-(i) case. It might be held that while in her room, Mary can entertain the relevant contents; but since she can only possess the relevant concepts deferentially or partially (where this is understood to include the sort of stipulative introduction discussed in the foregoing paragraphs), she cannot entertain these contents in a way that would enable her to know them.

The putative phenomenon that this argument appeals to is unprecedented in the sort of cases typically appealed to in discussions of 'deference' and 'partial understanding'. For example, although Burge's Arthritis-man fails to believe that arthritis is a disease of the joints, there is nothing that *prevents* him from coming to know this fact. He is perfectly capable of learning that arthritis is a disease of the joints by testimony. The proponent of this version of the knowledge argument would have to hold that some feature of Mary's grasp of the relevant contents prevents her from gaining knowledge

by testimony, even though she grasps the content expressed by the relevant sentences, knows the source to be reliable, and so forth. Similarly, the proponent of this version of the knowledge argument would have to hold that although Lonely Mary knows most of her own scientific theories, there is a certain part of or consequence of these theories that her grasp of the relevant concepts prevents her from coming to know. This would involve an *ad hoc* denial of widely accepted closure principles for knowledge.

These positions are coherent but highly unattractive. As before, the phenomenal concept theorist could staunchly bite the bullet; but as before, such a response is no more appealing than the internalist biting the bullet in response to Burge's original arguments. I therefore conclude that (PCC*) neither reinstates a new version of the knowledge argument, nor provides an adequate account of Mary's epistemic situation.

4 Conclusion

I have defended the following claims:

1. There are no phenomenal concepts.
2. If there are no phenomenal concepts, then the knowledge argument fails.

A thorough defence of physicalism against the knowledge argument would require an account of Mary's epistemic situation.¹³ But that is a topic for another paper.¹⁴

Arché
The University of St Andrews

DEREK BALL

¹³Many options remain open to the opponent of phenomenal concepts. She could hold that Mary learns nothing (Dennett (1991)). She could hold that Mary gains some non-factual sort of knowledge, such as knowledge how (Lewis 1996) or knowledge by acquaintance (Conee 1994). Or she could hold that the knowledge Mary gains is indexical (Horgan 1984; Perry 2001). I hope to discuss some other alternatives in future work.

¹⁴I wish to thank John Bengson, Josh Dever, Cory Juhl, Rob Koons, Bryan Pickel, David Sosa, several anonymous referees, and the editor of *Mind* for useful comments. I am especially indebted to David Chalmers, Mark Sainsbury, and Michael Tye, each of whom was extremely generous both with written comments and with helpful discussion.

17-19 College Street
St Andrews
Fife, KY16 9AL
Scotland
db71@st-andrews.ac.uk

References

- Alter, Torin 1998: 'A Limited Defense of the Knowledge Argument'. *Philosophical Studies*, 90, pp. 35–56.
- Alter, Torin and Sven Walter, (eds) 2007: *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*. New York: Oxford University Press.
- Ball, Derek 2007: 'Twin-Earth Externalism and Concept Possession'. *Australasian Journal of Philosophy*, 85, pp. 457–72.
- Balog, Katalin 1999: 'Conceivability, Possibility, and the Mind-Body Problem'. *Philosophical Review*, 108, pp. 497–528.
- Bealer, George 1994: 'Mental Properties'. *Journal of Philosophy*, 94, pp. 185–208.
- Block, Ned 1990: 'Inverted Earth', in Block et al. 1997, pp. 677–94. Originally published in *Philosophical Perspectives*, 4, 1990.
- Block, Ned 2003: 'Consciousness', in Nadel 2003, pp. 760–70.
- Block, Ned 2006: 'Max Black's Objection to Mind-Body Identity', in Alter and Walter 2007, pp. 249–306. Originally published in *Oxford Studies in Metaphysics*, 2, 2006.
- Block, Ned, Owen Flanagan, and Güven Güzeldere, (eds) 1997: *Consciousness: Philosophical Debates*. Cambridge, MA: The MIT Press.
- Burge, Tyler 1979: 'Individualism and the Mental', in Ludlow and Martin 1998, pp. 21–84. Originally published in *Midwest Studies in Philosophy*, 4, 1979.
- Burge, Tyler 1982: 'Other Bodies', in Woodfield 1982, pp. 97–120.

- Burge, Tyler 1986: 'Intellectual Norms and the Foundations of Mind'. *Journal of Philosophy*, 83, pp. 697–720.
- Burge, Tyler 1993: 'Concepts, Definitions, and Meaning'. *Metaphilosophy*, 24, pp. 309–25.
- Byrne, Alex 2006: 'Review of *There's Something About Mary*'. *Notre Dame Philosophical Reviews*, 1/20/2006.
- Chalmers, David 1996: *The Conscious Mind*. New York: Oxford University Press.
- Chalmers, David 2004: 'Phenomenal Concepts and the Knowledge Argument', in Ludlow et al. 2004, pp. 269–98.
- Conee, Earl 1994: 'Phenomenal Knowledge'. *Australasian Journal of Philosophy*, 72, pp. 136–50.
- Crane, Tim 1991: 'All the Difference in the World'. *Philosophical Quarterly*, 41, pp. 1–25.
- Davidson, Donald and Gilbert Harman, (eds) 1972: *Semantics of Natural Language*. Dordrecht: Reidel.
- Dennett, Daniel C. 1991: *Consciousness Explained*. New York: Back Bay Books.
- Donnellan, Keith S. 1993: 'There Is a Word for That Kind of Thing: An Investigation of Two Thought Experiments'. *Philosophical Perspectives*, 7, pp. 155–71.
- Fodor, Jerry 1998: *Concepts: Where Cognitive Science Went Wrong*. New York: Oxford University Press.
- Gertler, Brie 1999: 'A Defense of the Knowledge Argument'. *Philosophical Studies*, 96, pp. 317–36.
- Graham, George and Terence Horgan 2001: 'Mary Mary, Quite Contrary'. *Philosophical Studies*, 99, pp. 59–87.
- Guttenplan, Samuel, (ed) 1996: *A Companion to the Philosophy of Mind*. New York: Blackwell.

- Harman, Gilbert 1990: 'The Intrinsic Qualities of Experience', in Block et al. 1997, pp. 363–76. Originally published in *Philosophical Perspectives*, 4, 1990.
- Hellie, Benj 2004: 'Inexpressible Truths and the Allure of the Knowledge Argument', in Ludlow et al. 2004, pp. 333–64.
- Hill, Christopher S. 1997: 'Imaginability, Conceivability, Possibility and the Mind-Body Problem'. *Philosophical Studies*, 87, pp. 61–85.
- Holman, Emmett L. 2006: 'Dualism and Secondary Quality Eliminativism: Putting a New Spin on the Knowledge Argument'. *Philosophical Studies*, 128, pp. 229–56.
- Horgan, Terence 1984: 'Jackson on Physical Information and Qualia', in Ludlow et al. 2004, pp. 301–8. Originally published in *Philosophical Quarterly*, 34, 1984.
- Hume, David 1748: *An Enquiry concerning Human Understanding*. New York: Oxford University Press, 1999, Tom L. Beauchamp (ed).
- Jackson, Frank 1982: 'Epiphenomenal Qualia'. *Philosophical Quarterly*, 32, pp. 127–36.
- Jackson, Frank 1986: 'What Mary Didn't Know', in Block et al. 1997, pp. 567–71. Originally published in *Journal of Philosophy*, 83, 1986.
- Jackson, Frank 1995: 'Postscript', in Ludlow et al. 2004, pp. 409–16. Originally published in Moser and Trout 1995.
- Kripke, Saul A. 1972: *Naming and Necessity*. Cambridge, MA: Harvard University Press, 1980. Originally published in Davidson and Harman 1972.
- Levin, Janet 2007: 'What is a Phenomenal Concept?', in Alter and Walter 2007, pp. 87–110.
- Lewis, David 1996: 'Reduction of Mind', in Guttenplan 1996, pp. 412–30.
- Loar, Brian 1991: 'Personal References', in Villanueva 1991, pp. 117–33.
- Loar, Brian 1997: 'Phenomenal States', in Block et al. 1997, pp. 597–616. Adapted from a version in *Philosophical Perspectives*, 4, 1990.

- Loar, Brian 2003: 'Qualia, Properties, Modality'. *Philosophical Perspectives*, 13, pp. 113–29.
- Locke, John 1690: *An Essay Concerning Human Understanding*. London: Fontana/Collins, 1964, A.D. Woozley (ed).
- Ludlow, Peter and Norah Martin, (eds) 1998: *Individualism and Self-Knowledge*. Stanford: CSLI.
- Ludlow, Peter, Yugin Nagasawa, and Daniel Stoljar, (eds) 2004: *There's Something About Mary: Essays on Phenomenal Consciousness and Frank Jackson's Knowledge Argument*. Cambridge, MA: The MIT Press.
- Lycan, William 1996: *Consciousness and Experience*. Cambridge, MA: The MIT Press.
- McLaughlin, Brian and Michael Tye 1998: 'Externalism, Twin Earth, and Self-Knowledge', in Wright et al. 1998, pp. 285–320.
- Metzinger, Thomas, (ed) 1996: *Conscious Experience*. Exeter: Imprint Academic.
- Moser, Paul K. and J.D. Trout, (eds) 1995: *Contemporary Materialism: A Reader*. New York: Routledge.
- Nadel, Lynn, (ed) 2003: *Encyclopedia of Cognitive Science*. New York: Nature Publishing Group.
- Nida-Rümelin, Martine 1996: 'What Mary Couldn't Know', in Metzinger 1996, pp. 219–41.
- Nida-Rümelin, Martine 1998: 'On Belief About Experiences: An Epistemological Distinction Applied to the Knowledge Argument against Physicalism'. *Philosophy and Phenomenological Research*, 58, pp. 51–73.
- Nida-Rümelin, Martine 2002: 'Qualia: The Knowledge Argument', in Edward N. Zalta, (ed) 'The Stanford Encyclopedia of Philosophy (Fall 2002 Edition)', <http://plato.stanford.edu/archives/fall2002/entries/qualia-knowledge/>.
- Papineau, David 1993: *Philosophical Naturalism*. Oxford: Basil Blackwell.

- Papineau, David 2002: *Thinking About Consciousness*. Oxford: Oxford University Press.
- Papineau, David 2007: 'Phenomenal and Perceptual Concepts', in Alter and Walter 2007, pp. 111–44.
- Perry, John 2001: *Knowledge, Possibility, and Consciousness*. Cambridge, MA: The MIT Press.
- Putnam, Hilary 1970: 'Is Semantics Possible?', in Putnam 1975b, pp. 139–51. Originally published in *Metaphilosophy*, 1, 1970.
- Putnam, Hilary 1975a: 'The Meaning of "Meaning"', in Putnam 1975b, pp. 215–71. Originally published in *Minnesota Studies in the Philosophy of Science*, 7, 1975.
- Putnam, Hilary 1975b: *Mind, Language, and Reality: Philosophical Papers Volume 2*. New York: Cambridge University Press.
- Robinson, Howard 1993a: 'The Anti-Materialist Strategy and the Knowledge Argument', in Robinson 1993c, pp. 159–83.
- Robinson, Howard 1993b: 'Dennett on the Knowledge Argument', in Ludlow et al. 2004, pp. 69–73. Originally published in *Analysis*, 53, 1993.
- Robinson, Howard, (ed) 1993c: *Objections to Physicalism*. Oxford: Oxford University Press.
- Stoljar, Daniel 2005: 'Physicalism and Phenomenal Concepts'. *Mind and Language*, 20, pp. 469–94.
- Stoljar, Daniel and Yujin Nagasawa 2004: 'Introduction', in Ludlow et al. 2004, pp. 1–36.
- Sturgeon, Scott 1994: 'The Epistemic View of Subjectivity'. *Journal of Philosophy*, 91, pp. 221–35.
- Tye, Michael 1995: *Ten Problems of Consciousness*. Cambridge, MA: The MIT Press.
- Tye, Michael 2000: *Consciousness, Color, and Content*. Cambridge, MA: The MIT Press.

- Tye, Michael 2003: 'A Theory of Phenomenal Concepts'. *Royal Institute of Philosophy Supplement*, 53, pp. 91–105.
- Villanueva, Enrique, (ed) 1991: *Information, Semantics and Epistemology*. Oxford: Blackwell.
- Warner, Richard 1986: 'A Challenge to Physicalism'. *Australasian Journal of Philosophy*, 64, pp. 249–65.
- White, Stephen L. 1986: 'The Curse of the Qualia', in Block et al. 1997, pp. 695–718. Originally published in *Synthese*, 68, 1986.
- White, Stephen L. 2006: 'Aposteriori Identities and the Requirements of Rationality'. *Oxford Studies in Metaphysics*, 2, pp. 91–102.
- Woodfield, Andrew, (ed) 1982: *Thought and Object: Essays on Intentionality*. Oxford: Oxford University Press.
- Wright, Crispin, Barry C. Smith, and Cynthia Macdonald, (eds) 1998: *Knowing Our Own Minds*. New York: Oxford University Press.