

This article was downloaded by: [University of Liverpool]

On: 01 November 2012, At: 04:34

Publisher: Routledge

Informa Ltd Registered in England and Wales Registered Number: 1072954
Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH,
UK



Australasian Journal of Philosophy

Publication details, including instructions for authors
and subscription information:

<http://www.tandfonline.com/loi/rajp20>

A Posteriori Physicalists Get Our Phenomenal Concepts Wrong

Philip Goff ^a

^a University of Hertfordshire

Version of record first published: 16 Apr 2010.

To cite this article: Philip Goff (2011): A Posteriori Physicalists Get Our Phenomenal Concepts Wrong, *Australasian Journal of Philosophy*, 89:2, 191-209

To link to this article: <http://dx.doi.org/10.1080/00048401003649617>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.tandfonline.com/page/terms-and-conditions>

This article may be used for research, teaching, and private study purposes. Any substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing, systematic supply, or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae, and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand, or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

A POSTERIORI PHYSICALISTS GET OUR PHENOMENAL CONCEPTS WRONG

Philip Goff

Dualists say plausible things about our mental concepts: there is a way of thinking of pain, in terms of how it feels, which is independent of causal role. Physicalists make attractive ontological claims: the world is wholly physical. The attraction of a *posteriori* physicalism is that it has seemed to do both: to agree with the dualist about our mental concepts, whilst retaining a physicalist ontology. In this paper I argue that, in fact, a *posteriori* physicalism departs from the dualist's intuitive picture of our phenomenal concepts in just as radical a manner as more traditional forms of physicalism. Whereas the physicalism of David Lewis and David Armstrong is counterintuitive in holding that our only way of thinking about pain is in terms of its causal role, the physicalism of David Papineau and Brian Loar departs from common sense in holding that our phenomenal concept of pain is *opaque*: thinking of pain in terms of how it feels reveals nothing of what it is for something to feel pain. The arguments of David Chalmers and Frank Jackson against a *posteriori* physicalism involve general claims about *all* concepts. In contrast, my argument makes a claim only about phenomenal concepts: phenomenal concepts are not opaque.

1. A *Posteriori* Physicalism: Practically Perfect in Every Way?

A *posteriori* physicalism promises to give everything and take nothing. The *a priori* physicalism of Lewis [1966, 1970, 1994] and Armstrong [1968] implicitly denies that there is any way of thinking about a conscious state such as pain other than in terms of its causal role in a system mediating between sensory input and behavioural output.¹ For many this is a very implausible claim about our mental concepts. In contrast, the *a posteriori* physicalist concedes to the dualist that there is a way of thinking of pain—in terms of how it feels—which bears no conceptual connection to the causal role of pain. At the same time they deny that this dualism at the level of concepts is mirrored in a dualism at the level of properties. The *a posteriori* physicalist seems to follow the dualist in saying the right thing about concepts, but the *a priori* physicalist in saying the right thing about ontology. What more could be hoped for?²

¹See also Harman [1990], Dretske [1995], Rey [1995].

²For some examples of a *posteriori* physicalism, see Levine [1983], Loar [1990, 2003], Papineau [1993a; 2002], Tye [1995], Lycan [1996], Hill [1997], Hill and McLaughlin [1999], Block and Stalnaker [1999], and Perry [2001].

Sadly, this received opinion concerning the merits of *a posteriori* physicalism is exaggerated. *A posteriori* physicalism is committed to a view about our mental concepts, which is arguably just as counter-intuitive as that of the *a priori* physicalist. Many philosophers find *a priori* physicalism problematic because it is at odds with their carefully considered judgments about their mental concepts. I hope in this paper to persuade such philosophers to come to the same judgment about *a posteriori* physicalism.

2. Phenomenal Concepts Are Not Opaque

Most human beings, when they are in pain, are able to reflect introspectively on the feeling of pain, and think of it in terms of how it feels. Call this way of conceiving of pain the ‘phenomenal conception’ of pain, and call the concept of the feeling of pain which is formed from thinking about it in this way the ‘phenomenal concept’ of pain.³ When I carefully reflect on the feeling of pain phenomenally conceived, it does not seem to me to be apparent in that conception that to feel pain is to instantiate a state which has a certain causal role. Given this, in so far as I have reason to trust my carefully considered judgments about my own concepts, I have reason to doubt *a priori* physicalism.

But when undertaking such careful reflection on the feeling of pain phenomenally conceived, it also seems to me that this conception reveals at least something of what it is for something to feel pain. That’s why I have such strong reactions when I believe that someone feels pain. That’s why I run to the other room to get the painkillers. That’s why I judge that what Dave is doing to Steve with the hammer is so deeply morally objectionable. If a phenomenal conception of the feeling of pain didn’t reveal to me anything about what it is for someone to feel pain, why would I care that people feel pain (or at least why would I care until I had done the relevant empirical work to tell me what it is for someone to feel pain)?

This is not a judgment, at least not in the first instance, about extra-conceptual reality: it is rather a judgment about the phenomenal concept of pain. Let us say that a concept *C* of a property *F* is opaque iff *C* reveals nothing of what it is (or what it would be) for an object to have *F*. Let me give an example of an opaque concept. Suppose Kev and I are playing a fun game called, ‘Guess which property Kev is thinking about’. As it happens, the property Kev is thinking about is the feeling of pain. I am clueless, but I try to be clever, ‘Of course I know which property you are thinking about. You are thinking of “the property Kev is thinking about right now” (where I use this description as a rigid designator).’ Hilarity ensues.

Of course I have managed here to denote the property Kev is thinking about, the property of how pain feels. I have managed to refer to the feeling of pain in thought and language. But there is a clear sense in which I do not

³I am assuming that when two human beings think about the (determinable property of the) feeling of pain in terms of how it feels, they employ the same kind of concept. Hence I can talk about ‘the phenomenal concept of pain’.

know what it is for an object to have that property. I might say, ‘Of course I know what it is for an object to have that property . . . it’s for it to be such that it has “the property Kev is thinking about right now” (where I use this description as a rigid designator).’ But this is more trickery. Again, there is a clear sense in which, although I am able to refer to the property (which happens to be how pain feels), I have no idea what its nature is. The concept of *the property which Kev is thinking about right now* (rigidly designated) is opaque.

When I reflect carefully on the phenomenal concept of pain, it seems to me that it is not opaque in this sense: at least something of what it is for something to feel pain is knowable *a priori*. It is because reflection on pain in terms of how it feels reveals something of what it is to feel pain that I am moved to help those in pain. There is nothing mysterious about this. I am not having some mysterious *a priori* insight into the nature of reality. This is merely an intuition about a concept. I have the same intuition about the concept of *being a friend*, or the concept of *being spherical in Euclidean geometry*. To take the latter example, I know *a priori* at least something of what it is for an object to be spherical in Euclidean geometry, i.e. I know it’s for that object to have all points on its surface equidistant from its centre. This does not yet get me to reality, because I do not know *a priori* whether this concept is satisfied, that is, whether there are any objects which have the property of being spherical. In the same way, knowing something *a priori* about what it is for something to feel pain tells me nothing, in and of itself, about the nature of reality, because it is not *a priori* whether the phenomenal concept of pain is satisfied. Unless the ontological argument is sound, all conceptual analysis can tell me is *what it would take for a concept to be satisfied*, not whether or not such conditions are satisfied.

I argue below that *a posteriori* physicalism implies that the phenomenal concept of pain is opaque. Given this, if (i) I have good reason to trust my carefully considered armchair judgments about my concepts, (ii) the phenomenal concept of pain is one of my concepts, (iii) my carefully considered armchair judgments about the phenomenal concept of pain strongly suggest to me that it is not opaque, it follows that I have good reason to doubt the truth of *a posteriori* physicalism. Any philosopher whose carefully considered judgments about the phenomenal concept of pain agree with mine will likewise have reason to doubt the truth of *a posteriori* physicalism.

3. Avoiding Contentious Meta-Semantic Premises

The aim of this paper is to argue that *a posteriori* physicalism is committed to phenomenal concepts’ being opaque, and that this is a counter-intuitive conception of phenomenal concepts. David Chalmers’s well-known arguments against physicalism [1996, 2002, 2009] implicitly involve a commitment to phenomenal concepts’ not being opaque. But rather than offering arguments to the effect that *phenomenal concepts in particular* are not

opaque, Chalmers gets to this position by making very general claims about the nature of *all concepts*.

Consider the following four categories of concept:

Transparent: A transparent concept reveals the nature of its referent (where the referent is a property, the concept reveals what it is for an object to have that property).

Translucent: A translucent concept reveals part (but not all) of the nature of its referent (where the referent is a property, the concept reveals part of what it is for an object to have that property).

Mildly opaque: A mildly opaque concept does not reveal essential properties of the referent,⁴ but does reveal accidental features of the referent which uniquely identify it in the actual world.⁵

Radically opaque: A radically opaque concept reveals neither essential nor accidental properties of its referent.⁶

Chalmers's two-dimensional semantic framework rules out the existence of radically opaque concepts by involving the following two principles:

1. Every concept has a primary intension, which is a function from possible worlds (considered as actual) to referents, i.e. the primary intension of the concept *water* is a function which returns the referent H_2O to the world (considered as actual) where H_2O is the colourless, odourless stuff in oceans and lakes, returns the referent *XYZ* to the world (considered as actual) where *XYZ* is the colourless, odourless stuff in oceans and lakes, etc.

2. The primary intension can be evaluated *a priori*.

These two principles are inconsistent with the existence of concepts which reveal no properties of the referent. If we can evaluate the primary intension *a priori*, then we can know *a priori* the property of the referent which uniquely identifies it in the actual world; in the above example we can know *a priori* that water is the colourless, odourless stuff in oceans and lakes in the actual world. It is reasonably uncontroversial, for the reasons Kripke gives in *Naming and Necessity* [1980: 144–55], that phenomenal concepts are not mildly opaque. If phenomenal concepts are not mildly opaque, and there are

⁴Beyond essential properties which are trivial, e.g. being self-identical, or concern the very general ontological kind of the referent, e.g. being a property, being a countable object.

⁵We might alternatively define mildly external concepts as concepts possession of which involves implicit knowledge of how reference is fixed. This definition would avoid many of the tricky issues described in note 9, but departs from Chalmers's way of setting things up.

⁶Same qualification as with mildly opaque concepts, see note 4.

no radically opaque concepts, it follows that phenomenal concepts are not opaque.⁷

The banning of radically opaque concepts is essential for Chalmers's arguments against physicalism. A crucial premise in the argument, the premise that links mere conceivability to genuine possibility, is that every coherent proposition, even if necessarily false, is 'verified' at some genuinely possible world. A coherent proposition P is verified by world W iff under the hypothesis that W actually obtains, one ought to conclude that P is true. For example the proposition $\langle \text{water is XYZ} \rangle$ is necessarily false, but it is verified at the genuinely possible world where the colourless, odourless stuff in oceans and lakes is XYZ. On the assumption that this possibility is actual, one ought to conclude that $\langle \text{water is XYZ} \rangle$ is true.⁸

The reason $\langle \text{water is XYZ} \rangle$ is verified at the XYZ world is that the concept *water* reveals *a priori* an accidental property which uniquely identifies its referent in the actual world: the property of being the colourless, odourless stuff in oceans and lakes. The fact that this proposition is verified at the XYZ world is explained in terms of the mild opacity of the concept *water*. Suppose there is a radically opaque concept, the concept of *that stuff*, which refers to H₂O. The proposition $\langle \text{that stuff is XYZ} \rangle$ would be coherent, in the sense that it cannot be ruled out *a priori*, and yet would not be verified at any possible world. Whatever possible world I am considering, given that the concept of *that stuff* reveals no properties of its referent, I would have no way of picking out *that stuff* in the world that I am considering as actual. If there are radically opaque concepts, then Chalmers's crucial premise linking conceivability to possibility is false.⁹

⁷Jackson [1998] has a very similar two-dimensional framework which also excludes radically opaque concepts, and which he uses to oppose *a posteriori* physicalism.

⁸We need to add 'centres' to the worlds considered as actual to evaluate the primary intension. Centres would also have to be built into a satisfactory definition of a mildly opaque concept, e.g. we would need centring to distinguish the watery stuff on Earth from the watery stuff on Twin Earth. I ignore this complication.

⁹There are extra complexities here. Could we not bring facts about reference into the worlds being considered for verification of $\langle \text{that stuff is XYZ} \rangle$? For example, we could suggest that this proposition is true at the world considered as actual where a recognitional capacity picking out XYZ was activated in me at the time (we can pick out the relevant individual and time with reference to the centre of the world) I thought this thought (and no other referential capacity was activated, etc.). If this is fair game, then the physicalist can reply that the proposition Q $\langle \text{all the actual physical facts obtain but no badger has ever felt pain} \rangle$ is true at the world considered as actual where (i) all the actual physical facts obtain (ii) a Chalmers-style phenomenal concept referring to a non-physical property F is activated in me when I think about pain, (iii) badgers don't have F , but is false at the world considered as actual where (i) all the actual physical facts obtain, (ii) a Papineau-style phenomenal concept referring to c-fibres' firing is activated in me when I think about pain. This would result in pain having distinct primary and secondary intensions (the secondary intension being dependent on the referent of the actual concept activated in me when I think about pain), which would render false another crucial premise in Chalmers's argument: that phenomenal concepts have identical primary and secondary intensions (which allows the move from there being a world considered as actual where there are zombie badgers to there being a world considered as counterfactual where there are zombie badgers).

But there is another problem for the physicalist. One might think that the physicalist ought to hold that complete knowledge of the actual physical facts would reveal that 'pain' refers to c-fibres' firing, in which case we can rule out *a priori* that all the actual physical facts obtain in the absence of pain (because the actual physical facts entail that 'pain' refers to c-fibres' firing). I think there is room for the physicalist to deny this. She could say that, if there is actually a non-physical property closely associated with c-fibres' firing, then 'pain' refers to that non-physical property, but that if there is actually no such non-physical property closely associated with c-fibres' firing, then 'pain' refers to c-fibres' firing (this strategy would be similar to the 'conditional analysis' defence of physicalism: Hawthorne [2002], Stalnaker [2002], Braddon-Mitchell [2003]). We could tell a similar story about the conceivability of all the actual physical facts obtaining in the absence of *that stuff*, so there is no reason to think that phenomenal concepts are unique in this regard.

Chalmers also suggests that the mere fact that there is a minimal physical duplicate of our world, call it W , which verifies the proposition Q , is enough to refute physicalism, as the actual world differs from W in that Q

Chalmers's arguments, then, rely on a highly contentious meta-semantic assumption: there are no radically opaque concepts.¹⁰ The argument I present in this paper does not. I am happy to allow that there are radically opaque concepts; perhaps we have certain radically opaque concepts grounded in sub-personal recognitional capacities.¹¹ My argument relies specifically on a claim about phenomenal concepts: phenomenal concepts are not opaque.

4. *A Posteriori* Physicalism Implies That Phenomenal Concepts Are Opaque

A posteriori physicalists, as physicalists, hold that to feel pain just is to instantiate some functional or physiological state, which we can call, as is the tradition, "c-fibres' firing". If this were the case, and if the concept of pain were *transparent*, that is, if it revealed what it is for someone to feel pain, it would follow that the phenomenal concept of pain reveals that to feel pain is to have your c-fibres fire. But this is precisely what the *a posteriori* physicalist denies. *A posteriori* physicalism is inconsistent with the claim that the phenomenal concept of pain is transparent.

Objection 1: Fair enough, the *a posteriori* physicalist cannot claim that the phenomenal concept of the feeling of pain reveals *everything* about what it is for something to feel pain (after all, it doesn't reveal that to be in pain is to have your c-fibres fire), but she can claim that the phenomenal concept of pain reveals *something* of what it is for something to feel pain. Perhaps our armchair intuitions about the phenomenal concept of pain are strongly against its being opaque, but they are not obviously against its being translucent. The

is actually false (actual badgers feel pain). I think the physicalist can plausibly deny that there is any minimal physical duplicate of the actual world which verifies *Q*—on the grounds that if there are in actuality only physical properties 'pain' refers to a physical property—whilst making sense of the conceivability of *Q* in the following way: under the hypothesis that dualism actually obtains, there is a world considered as counterfactual at which *Q* is true.

Alternatively, I think the physicalist could concede that strictly speaking zombies (including zombie badgers) are ruled out *a priori*, but *only* because the physical facts entail facts about the reference of phenomenal concepts, not because phenomenal concepts can be analysed into functional concepts. This would still avoid the difficulties with the standard *a priori* physicalist position: it would allow that we have a special way of thinking about experiences which is not conceptually connected to functional role, which is the new concept Mary learns when leaving her black and white room, etc. The physicalist may still be obliged to explain why we find it difficult to accept that 'pain' refers to c-fibres' firing, in a way that we don't find it difficult to accept that *that stuff* refers to H₂O, but physicalists have gone to some lengths to do this, e.g. Papineau [1993a].

¹⁰See Dretske [1994] and Millikan [1984] for theories that involve a commitment to radically opaque concepts. In the theories of phenomenal concepts offered by *a posteriori* physicalists, phenomenal concepts are generally conceived of as some kind of radically opaque concept, e.g. demonstratives [Papineau 1993b; Perry 2001], or grounded in recognitional capacities [Loar 1990] or facts about teleology [Papineau 2002, 2007] rather than *a priori* content. Chalmers and Jackson [e.g. Jackson 1998: 37–42] often claim that their two-dimensional framework is consistent with the causal theory of reference, and indeed it is in the sense that the primary intension might express certain causal/historical connections between the referent and employments of the concept. However, their framework is inconsistent with meta-semantic theories according to which those causal/historical relations, or whatever facts determine reference, are not *a priori* associated with the concept. This is still a highly contentious meta-semantic assumption, which takes away from the dialectical strength of Chalmers's anti-physicalist arguments.

¹¹Suppose I meet Bob at a party and form a recognitional capacity which allows me to refer to Bob. Two months later I wake up and think, 'I wonder if I'll met *that guy* again', where I employ my recognitional capacity in order to think about Bob. I can't remember Bob's name, or what he looks like. The only thing I remember about Bob is that he is a guy I met at a party, but this does not uniquely identify him, as I have met lots of guys at parties. This seems to me a plausible example of a radically opaque concept.

a posteriori physicalist, then, can hold that the phenomenal concept of pain is translucent, and in this way appease our armchair intuitions.

One way to interpret this objection, which one hears a lot in conversation if not in print, is as follows:

*Objection 1**: Pain has number of aspects. The phenomenal concept of pain reveals the phenomenal aspect of pain, i.e. how it feels, but it may also have a physical aspect, e.g. c-fibres' firing, which is not revealed by the phenomenal concept of pain.

To agree with *Objection 1** is to concede property dualism: there is one state—call it 'pain' if you like¹²—which has on the one hand a physical aspect, and on the other hand a phenomenal aspect. The consistent physicalist cannot settle with this. The physicalist is obliged to claim that *the feeling of pain itself* is a physical property. If the feeling of pain is not a physical property, then there is a property which is not physical (assuming the feeling of pain exists) and physicalism is false. This quite straightforward point is generally conceded by physicalists, but it is good to be clear about it.¹³

There is another way to interpret *Objection 1*:

*Objection 1***: It is not that there are two properties here, one phenomenal and one physical. There is just a single property, how pain feels, but the phenomenal concept reveals some, but not all, of what it is for an object to have that property.

Even interpreted this way, I don't think it can help the *a posteriori* physicalist. For the *a posteriori* physicalist, how pain feels is a wholly physical property. If the phenomenal concept of pain is translucent, then it reveals an aspect of that property. But, crucially, *an aspect of a wholly physical state is itself a physical state*. Therefore, if the phenomenal concept of pain were translucent, it would reveal that how pain feels involves a physical state (that physical state being an aspect of the physical state which is the feeling of pain). But this is precisely what the *a posteriori* physicalist denies. *A posteriori* physicalism is inconsistent with the claim that phenomenal concepts are translucent.

Let's take an example to make this clearer. Suppose that for a thing to feel pain is for its c-fibres and d-fibres to fire. If the phenomenal concept of pain revealed that the feeling of pain involved c-fibres' firing, but did not reveal that the feeling of pain involved d-fibres' firing, then the phenomenal concept of pain would be translucent: it would reveal an aspect of what it is to feel pain (to have c-fibres' firing), but not the

¹²My linguistic intuitions tell me that 'to be in pain' just is, and is nothing more than, to feel pain (to think this is not to beg the question against materialism; many materialists, e.g. David Lewis [1980: §VIII], and David Papineau [2002: 82] agree). But this terminological debate over whether 'being in pain' is semantically equivalent to 'feeling pain' is utterly unimportant, for the simple reason that we can choose to explicitly focus, as I have mostly done here, on the feeling of pain itself, and ask whether that is a physical property.

¹³See, for example, Lewis [1980: §VIII] and Papineau [2002: 82].

whole of what it is to feel pain (to have c-fibres firing and d-fibres firing). But this would clearly be inconsistent with the *a posteriori* physicalist's contention that it is not *a priori* that the phenomenal concept of pain denotes a physical property.¹⁴

Objection 2: The phenomenal concept of pain does reveal what it is for something to feel pain. The physical concept of c-fibres' firing also reveals what it is for something to feel pain. The physical and the phenomenal concept are two conceptually distinct ways of knowing what it is for something to feel pain.

It is not at all obvious that the claim expressed in *Objection 2* is intelligible. Of course we can *refer* to a property in lots of different conceptually distinct ways: in terms of its accidental properties, in virtue of a causal or historical connection, etc. But it does not follow from this that we can *know what it is for a specific property to be instantiated* in numerous conceptually distinct ways. Someone understands what it is for an object to be spherical (in Euclidean geometry) just in case they know that a spherical object is a three-dimensional object which has all points on its surface equidistant from its centre. Someone understands what it is for something to be a bachelor just in case they know that a bachelor is an unmarried man, and they know what the institution of marriage is. It is difficult to make sense of the thought that the notion of a bachelor, or the notion of sphericity, could be understood in two conceptually distinct ways.

Of course I can know the proposition, <for x to be a bachelor is for x to have the property Kev is thinking about right now> (where Kev is thinking about the property of being a bachelor), which is conceptually distinct from, but at the level of reference identical to, the proposition <for x to be a bachelor is for x to be an unmarried man>. But I take it that there is a clear sense in which to know the former proposition is not to know what it is for something to be a bachelor.

Call the following 'The Thesis of Dubious Intelligibility' (TDI):

For some property F , there are two conceptually distinct ways of knowing what it is for F to be instantiated.

It has become apparent to me in conversation that many people take the *a posteriori* physicalist Brian Loar's view to involve a commitment to TDI. Superficial attention to Loar's way of putting things gives that impression. But closer attention reveals that at no point does Loar commit to TDI. His way of talking superficially suggests that he takes phenomenal concepts to be transparent, talking as he does of phenomenal concepts affording us a 'grasp of the essence' of phenomenal properties. But closer examination reveals that all he means by a concept affording a 'grasp of essence' is that the concept is not mildly opaque:

¹⁴Or suppose that to feel pain is to be in a certain a functional state. If the concept of *the feeling of pain* is translucent, then the concept will partly reveal the functional nature of the feeling of pain, which is again inconsistent with the claims of the *a posteriori* physicalist.

Phenomenal concepts, as we have seen, do not conceive of their reference via contingent modes of presentation. And so they can be counted as conceiving phenomenal qualities directly. Calling this a grasp of essence seems to me all right, for phenomenal concepts *do not conceive their references by way of their accidental properties* . . .

[Loar 1990: 305, my emphasis]

The most natural reading of the phrase, ‘Phenomenal concepts give us a grasp of essence’, is that phenomenal concepts reveal what it is for a conscious state to be instantiated, i.e. that phenomenal concepts are transparent. But this is not what Brian Loar uses this phrase to mean. Rather, he uses it to mean that phenomenal concepts do not conceive of their referents by way of their accidental properties, i.e. phenomenal concepts are not mildly opaque. But this leaves two very different options open: either phenomenal concepts might be *transparent/translucent* or they might be *radically opaque*.

In other words, Loar dodges the issue. By lumping together transparent/translucent concepts and radically opaque concepts, under the category of ‘direct concepts’, he declines to tell us whether on his view phenomenal concepts are transparent/translucent or radically opaque, and so declines to tell us whether or not his view is committed to TDI (which it would be if he held both that phenomenal concepts were transparent/translucent and that physical concepts were transparent).¹⁵

Because Loar does not explicitly defend TDI, he does not put the work in to making this thesis intelligible. At the very least, if the *a posteriori* physicalist wants explicitly to defend the idea that we could know what it is for a single property to be instantiated under two conceptually distinct modes of presentation, she ought to spell out the details of this thesis, to put some flesh on the bone, and in doing so help us to make sense of this initially obscure notion. Could the same be true of other properties, e.g. could we know what it is for something to be a bachelor under two conceptually distinct modes of presentation? Or is it only in the case of pain that we can grasp its nature in two conceptually distinct ways? Why is pain special in this way? No *a posteriori* physicalist has even begun to tackle these issues, or indeed even explicitly endorsed TDI.¹⁶

¹⁵Why would Loar be committed to TDI if he thought that phenomenal concepts were translucent and that physical concepts were transparent? Answer: if the phenomenal concept of pain is translucent, then it reveals an aspect, call it *a*, of what it is to feel pain. If the physical concept that denotes the feeling of pain reveals the complete nature of the feeling of pain, then it will reveal the nature of *a* (as an aspect of that complete nature). If the physical concept conceives of *a* in physical terms, and the phenomenal concept conceives of *a* in phenomenal terms, then the physical and the phenomenal concept reveal the nature of *a* under conceptually distinct modes of presentation (according to *a posteriori* physicalism).

¹⁶Both Loar [1990] and Papineau [2002] have put effort into explaining what is special about phenomenal concepts, but not specifically *what is special about phenomenal concepts such that TDI is or might be true*. Loar avoids the issue by not distinguishing between transparent/translucent concepts on the one hand and radically opaque concepts on the other. Papineau, as I explain below, has recently explicitly committed to phenomenal concepts’ being opaque, and so does not need to accept TDI. As I also explain below, the primary aim of Papineau and Loar is to give an account of phenomenal concepts which enables them to make sense of psycho-physical identities’ being exceptions to what Papineau calls ‘the transparency thesis’ and Loar calls ‘the semantic premise’, a principle which is not equivalent to TDI.

Loar avoids explicitly endorsing the (highly counterintuitive) claim that phenomenal concepts are opaque. But in the more recent writings of David Papineau we find an *a posteriori* physicalist willing to face up to this implication of his view. Indeed, Papineau not only claims that phenomenal concepts are opaque, but goes as far as to claim that *all* atomic concepts are opaque, which we can see in the following attack on the notion of a transparent atomic concept:

... the notion of direct acquaintance on which it rests strikes me as highly suspicious. It assumes some mode of thought where objects become completely transparent to the mind and all their essential properties are thereby laid bare. It is hard not to see this as inspired by some misplaced visual model, in which we are able to peer in at some immaculately illuminated scene. I myself doubt that there is any such mode of thought. No doubt there are ways of thinking of things that make certain essential properties *a priori* knowable. But I take such *a priori* knowledge to derive from (possibly implicit) compositionality in the relevant modes of thinking, and so not to be associated with the most basic ways in which thought makes contact with reality ... When it comes to these basic points of contact, I find it hard to take seriously any alternative to the assumption that our atomic concepts are related to reality by facts external to our *a priori* grasp, such as causal or historical facts.

[Papineau 2006: 102]

... straightforward physicalists do not accept *any* transparent modes of thought ... I don't recognize any way in which the mind 'captures' something, apart from simply referring to it.

[ibid: 106]

We can see, then, that Chalmers and Papineau make polar opposite, but arguably equally contentious, meta-semantic assumptions. Both deny the existence of at least one of the four categories of concept I have formulated: Chalmers denies the existence of radically opaque concepts whilst Papineau denies the existence of (atomic) transparent concepts and translucent concepts.¹⁷

I will not here go into the details of Papineau's account of phenomenal concepts, but it is clear that he takes these concepts to be basic (indeed it is difficult to see how an *a posteriori* physicalist can say otherwise given that she resists giving functionalist analyses of phenomenal concepts). Hence, we can see that Papineau holds that a phenomenal conception of pain reveals *none* of the essential features of the feeling of pain, that there is no sense in which this way of thinking of pain 'captures' the feeling other than simply referring to it. Papineau believes that thinking of the feeling of pain in terms of how it feels reveals nothing of what it is to feel pain. Stop and reflect for a moment on what Papineau believes (often a clear exposition of a view can

¹⁷I take it that it follows from atomic concepts' being 'related to reality by facts external to our *a priori* grasp' that atomic concepts are neither transparent nor translucent.

constitute an argument against it): *to think of the feeling of pain in terms of how it feels reveals nothing of what it is to feel pain.*

There are two ways in which Papineau caricatures his opponent in the above passage. First, he does not distinguish between a commitment to transparency and a commitment to translucency (at least with regard to atomic concepts). Compare:

The phenomenal concept of pain reveals *something* of what it is to feel pain (reveals some essential features of the feeling of pain, captures the feeling of pain in some way which goes beyond merely referring to it).

The phenomenal concept of pain reveals *everything* of what it is to feel pain (reveals all essential features of pain).

It is clear that the latter claim is much stronger than the former, but Papineau lumps the two together, and in so doing makes his opponent seem less moderate than she need be. As I have argued, *a posteriori* physicalism is at odds not only with the claim that phenomenal concepts are transparent, but also with the much more moderate claim that phenomenal concepts are translucent.

Secondly, and perhaps more importantly, the visual metaphor of ‘peering at some immaculately illuminated scene’, implies that a commitment to the transparency or translucency of an atomic concept involves some kind of mysterious *a priori* insight into the nature of reality. But to claim that a concept is transparent or translucent is just to claim something about the concept, not about extra-conceptual reality. To claim that a concept is translucent is to claim that that concept reveals something of what it demands of reality. To claim that a concept is transparent is to claim that that concept reveals everything of what it demands of reality. Of course, if we have independent reason to think those demands are satisfied then this will have implications for how the world is. But a mere commitment to the translucency, or indeed the transparency, of a concept reveals nothing about the world outside the concept.¹⁸

We can see *a posteriori* physicalists as caught on the horns of a dilemma. Either they hold that phenomenal concepts are opaque, and put themselves at odds with deep intuitions we have about our phenomenal concepts, or they accept that phenomenal concepts are transparent/translucent and accept TDI. Until *a posteriori* physicalists explicitly claim that phenomenal concepts are transparent/translucent, rather than merely not mildly opaque, I will read them as embracing the first horn of this dilemma (as Papineau does explicitly).

¹⁸As I will discuss below, it is plausible to think that, in having a phenomenal concept of pain whilst introspecting one’s pain, one can know for certain that that concept is perfectly satisfied. If this is the case, and if the phenomenal concept of pain is transparent, then mere introspection can give us insight into the nature of reality. Merely by introspecting I can know the nature of a property, and know that that property is instantiated. But even if this is so, it is not the mere *transparency* of the phenomenal concept which allows us insight into the nature of reality; it is transparency in conjunction with the certainty that the concept is perfectly satisfied. There is nothing mysterious about the notion of transparency itself.

Slight Interlude: Breaking the Impasse in the Consciousness Debate

The primary aim of the theories of phenomenal concepts offered by Papineau and Loar is to show how a principle which they take to be employed in Kripke-inspired anti-physicalist arguments—a principle which Papineau calls ‘the transparency thesis’, and which Loar calls ‘the semantic premise’—can be shown to be false. This is the principle that all *a posteriori* identities involve at least one mildly opaque concept [Loar 1990; Papineau 2002: ch. 3]. We can roughly interpret Chalmers as arguing that: (i) The transparency thesis is true (on the grounds that the standard examples of *a posteriori* identities, e.g. $\langle \text{water is H}_2\text{O} \rangle$ involve at least one mildly opaque concept), (ii) Phenomenal concepts are not mildly opaque (for the reasons Kripke gives in *Naming and Necessity*), (iii) Therefore, $\langle \text{pain is c-fibres' firing} \rangle$ is false (because neither the concept of *pain* nor the concept of *c-fibres' firing* is mildly opaque¹⁹).

I am perfectly happy to agree with Papineau and Loar that the transparency thesis is false: the proposition $\langle \text{that stuff is XYZ} \rangle$, in which *that stuff* is a radically opaque concept which refers to H₂O, is an exception to the transparency thesis. But I find it a great deal less plausible that TDI is false. The consciousness debate, as it has been held in the last fifteen or so years between *a posteriori* physicalists like Papineau and Loar and dualists/*a priori* physicalists like Chalmers and Jackson, has placed too much emphasis on the transparency thesis, or something like it, as a result of failing to clearly distinguish between transparent/translucent concepts on the one hand and radically opaque concepts on the other (both of these kinds of concept are usually just lumped together under the category of ‘direct concepts’). It is this misplaced emphasis which I hope to correct with this paper. It would be better at this point to move the argument on to the questions of: (i) whether or not phenomenal concepts are opaque, (ii) whether or not TDI is true.

Objection 3: You have been all along assuming that *physical concepts*, such as the concept of c-fibres’ firing, are transparent. But why should they not be opaque? If the phenomenal concept of pain is transparent, and the physical concept of c-fibres’ firing is opaque, then both concepts can denote the same property without the need to accept TDI.

There is a non-standard form of physicalism, a view I like to call ‘funny physicalism’, which makes sense of the identity of physical and phenomenal properties in this way.²⁰ Bertrand Russell, one of the first funny physicalists, claimed that there is an important sense in which physics, for all its success,

¹⁹More precisely: the physicalist cannot rely on the concept of c-fibres’ firing being mildly opaque unless she is a funny physicalist (see response to *Objection 3* below).

²⁰‘Funny peculiar’ (in a non-pejorative sense) rather than ‘funny ha ha’. For some examples of funny physicalism, see Russell [1927], Eddington [1928], Feigl [1967], Maxwell [1979], Lockwood [1989], Strawson [1994, 2003, 2006], Chalmers [1996], Griffen [1998], and Stoljar [2001].

tells us very little about the nature of the physical world. The language of physics is mathematical, and, because of this, physics is suited to tell us only about the *structure* of the physical world, not about the qualitative nature of whatever it is that has that structure. The only window we have onto the qualitative nature of physical reality is via our immediate acquaintance with our conscious experience. Physics tells us about the structure of our brain, but in introspection we peer in on the reality of the thing that has that structure.²¹

This view is very different from ‘physicalism’ as it is normally understood (the view that Papineau [2006] calls ‘straightforward physicalism’, and Galen Strawson [2003] calls ‘physicSalism’), and is arguably only terminologically distinct from property dualism. For the funny physicalist, conscious experience is an aspect of the nature of the brain which goes beyond what physics can ever tell us about the brain. If we define ‘physical property’, as I have been implicitly doing throughout this paper, such that a property is physical iff it is (or is *a priori* by) the kind of property which physics transparently reveals to us, then funny ‘physicalism’ is committed to reality’s containing non-physical properties, and so is not really a form of physicalism. It is only if we define ‘physical property’ in some way that is not so rigidly tied to physics (for example, a property *F* is physical iff it is had by physical objects, i.e. the kind of *objects* physics tells us about, regardless of whether or not *F* is one of the properties transparently revealed by physics) that the view counts as a form of physicalism.²²

Regardless of whether or not we are happy to call this view ‘physicalism’, it is metaphysically much closer to dualism than it is to views which are generally called ‘physicalist’ in the literature; standard forms of physicalism hold that phenomenal properties are identical to, or *a priori* entailed by, the kind of properties neuroscience reveals to us the nature of. The attraction of standard physicalist views is that they are very economical, committed only to the kind of properties physics transparently reveals to us (which according to Russell are structural properties²³).²⁴ Funny physicalism believes in these properties *and* in properties of physical objects which go beyond these properties. David Chalmers puts the matter well when he suggests that funny physicalism (which he calls ‘type-F monism’), ‘fits the letter of materialism’ but ‘shares the spirit of antimaterialism’ [2002: 265].

It is standard physicalism, rather than funny physicalism, which is the target of this paper. I accept that *Objection 3* avoids denying the transparency of the phenomenal concept of pain, and does so in a way which avoids embracing TDI, but the view that results is not a form of

²¹Again, if we read this as merely a commitment to transparency, it need not imply *a priori* insight into extra conceptual reality. It is only the conjunction of (i) phenomenal concepts are transparent (ii) phenomenal concepts are perfectly satisfied, which implies that our phenomenal concepts can tell us something about reality.

²²See Stoljar [2001] for a detailed account of these two ways of defining what a ‘physical property’ is.

²³Most funny physicalists take views according to which there are only structural properties to be incoherent (Strawson [2003], Stoljar [2001]), but other philosophers defend the coherence of such a metaphysical view (Ladyman and Ross [2007], van Fraassen [2006]).

²⁴More precisely: standard physicalists are only committed to the kind of properties which physics reveals to us the nature of, and to properties which are *a priori* entailed by the kind of properties which physics reveals to us the nature of.

standard physicalism. Hence, I will set this objection on one side, and focus only on forms of physicalism which do not rely on a denial of the transparency of physical concepts in order to make sense of the *a posteriori* identity between conscious states and physical states.^{25,26}

Objection 4: You are assuming throughout that our phenomenal concepts are *perfectly satisfied*. Consider the following analogy. Suppose our folk concept of time is transparent, and that it reveals time to be essentially something that flows. Now suppose that, in terms of the metaphysics of the matter, the B-theory is true. Does this mean that our concept of time is not satisfied? Does this mean that our concept of time is not transparent? We need say neither of these things. We can say the folk concept of time transparently makes a series of demands of reality (one of these demands being that time flows), and that the static B-series satisfies some, but not all, of those demands (it does not satisfy the demand that time flows). We can say that our concept of time is *imperfectly realized* by the static B-series. Could we not say the same of our phenomenal concept of the feeling of pain? The phenomenal concept of pain may make a number of transparent demands of reality, some, but not all, of which are satisfied by C-fibres' firing. C-fibres' firing may be an imperfect realizer of the phenomenal concept of pain.

It is easy to see how this story goes with reference to time, because there are a number of aspects to our folk concept of time. At the very least, there are the two aspects of time folkishly conceived to which McTaggart drew our attention: first, that time consists of events distinguished by the relations of *earlier* and *later*, and, second, that time consists of events distinguished by the categories of *past*, *present* and *future*. The static B-series matches up to the former but not the latter aspect of our folk conception of time.

But it is not easy to see how we could tell such a story about *the feeling of pain*, about *how pain feels*. Our phenomenal concept of pain demands merely that we *feel pained*. How could this demand be partially satisfied (remember that for the *a posteriori* physicalist the phenomenal concept of pain does not transparently demand that any functional properties be instantiated)? There may be a number of aspects to my pain—a throbbing feeling, a nauseous tugging—but they are all *phenomenal aspects*. If our concept of any one of these aspects is perfectly satisfied, then our concept of *feeling pain* is perfectly satisfied, since these specific aspects of pain are just determinate ways of instantiating the determinable property of feeling pain.

Moreover, it is plausible to think, in contrast to the case of the folk concept of time, that I do know for certain that the demand my phenomenal concept of pain makes upon reality is perfectly satisfied. When I cut my hand, and feel pain, and focus my attention on the pain, and think of it in terms of its feeling pained, there is no epistemic possibility that I am not feeling pained. I said earlier that, assuming that the ontological argument

²⁵See Goff [forthcoming] for a decisive refutation of funny physicalism, and also of *a priori* physicalism. That paper also contains an argument against *a posteriori* physicalism, but one of the premises (that ghosts are possible if conceivable) relies on the arguments in this paper. Goff [forthcoming] and the paper you are now reading can be read in conjunction as an attack on all forms of physicalism.

²⁶See note 5 of Papineau [2002: 85] for an explicit rejection of this strategy for making sense of the *a posteriori* identity between physical states and conscious states.

fails, conceptual analysis can only tell us what concepts demand of reality, it can never tell us whether those demands are satisfied. In the case of our concepts of our conscious states, we (arguably) have a unique case whereby we know for certain that those concepts are perfectly satisfied.²⁷ This certainty, in conjunction with conceptual analysis of phenomenal concepts, can get us all the way to reality.²⁸

At the very least, if the *a posteriori* physicalist wants to claim that our phenomenal concept of pain is imperfectly realized, she has a story to tell. She has so far shown no signs of telling that story.²⁹

5. Comparisons to Nida-Rümelin's Argument for Property Dualism

There are similarities between the argument in this paper against *a posteriori* physicalism, and Martine Nida-Rümelin's [2007] argument for property dualism. Nida-Rümelin focuses her argument around the claim that phenomenal concepts are transparent (in her terminology phenomenal concepts allow us to 'grasp' phenomenal properties), arguing that this, together with some plausible assumptions (one of which, her 'Principle of Cognitive Transparency', is similar to my denial of TDI), entails property dualism. Nida-Rümelin claims, as I do with reference to my own argument, that her argument relies on less contentious assumptions than the arguments of David Chalmers; specifically she claims to remain neutral with regards to Chalmers's assumptions (A) that all facts apart from the phenomenal facts are *a priori* entailed by the physical facts, (B) that every coherent proposition is verified at some metaphysically possible world.³⁰

However, it can be seen that Nida-Rümelin makes the same contentious meta-semantic assumption as Chalmers: her framework—like Chalmers's—has no place for radically opaque concepts. For Nida-Rümelin, possession of a concept involves having implicit knowledge of how the counterfactual extension of the concept is dependent on the way the actual world is:

Understanding the concept of water thus involves knowing that a liquid in counterfactual circumstances falls into the extension of the concept just in case it

²⁷As I said in notes 18 and 21, this does not contradict what I have claimed earlier about the transparency of a concept not implying *a priori* insight into the nature of reality. It is only the transparency of a phenomenal concept in conjunction with the certainty that that concept is satisfied which has implications for how the world is. An eliminativist about consciousness could consistently claim that our phenomenal concepts are transparent but not satisfied, and hence that those concepts afford us no insight into the nature of reality.

²⁸I am not here making the more controversial claim that I grasp all aspects of my pain. We can take it that the certainty I am talking about is with reference to some determinable concept of pain, rather than some highly specific determinate of it. I can be certain, in the moment of feeling pain, that the highly determinable concept of pain I am employing is perfectly satisfied.

²⁹Other kinds of physicalists have appeared to suggest that folk concepts of conscious states are imperfectly realized, e.g. Lewis [1995] and arguably Dennett [1991]. However, *a posteriori* physicalists tend not to want to be revisionary about our phenomenal concepts. Perhaps there is room for a hybrid between the physicalism of Papineau/Loar and the physicalism of Lewis/Dennett: a form of *a posteriori* physicalism which takes our phenomenal concepts to be imperfectly satisfied. Until such a view is spelt out, I shall ignore it.

³⁰Strictly speaking, Nida-Rümelin's argument is against all forms of physicalism, rather than merely *a posteriori* physicalism. However, the premise which excludes *a priori* physicalism, the principle of 'Cognitive Independence of Physical and Phenomenal Properties', is only very briefly defended. If this premise were added to my argument, my argument would also be (at least as) effective against all forms of standard physicalism.

is composed of H₂O, if the real world is a world with liquids composed of H₂O in the rivers lakes and oceans on Earth ... If you are given the relevant information about an arbitrarily chosen possible world *w* (the information about the chemical composition of the liquids in the rivers, lakes and oceans in *w*), then you know what would be the secondary intension of the concept of water if *w* were actual. [2007: 318]³¹

Having implicit knowledge of how the counterfactual extension is dependent on how the actual world turns out surely entails having implicit knowledge of what features identify the referent in worlds considered as actual. If I wasn't able to identify the referent in a given world considered as actual, then I wouldn't know what facts in that world are relevant to determining the counterfactual extension. To take the specific case of water, if I did not know that water is the actual liquid in the rivers, oceans and lakes, then I wouldn't know that the facts about the actual world relevant to determining the counterfactual extension are the facts concerning the chemical composition of the liquid in oceans and lakes. It seems, then, that for Nida-Rümelin possessing a concept involves implicit knowledge of the identifying properties of the referent in a world considered as actual. If this is what is involved in possessing a concept, then no concept is radically opaque.

In the absence of radically opaque concepts, there seems much less reason to doubt Chalmers's assumptions (A) and (B). When it comes to concepts which, as in the case of the concept of water, reveal the properties of the referent that identify them in worlds considered as actual, it is quite plausible to think that (i) the fact that such a concept is satisfied is *a priori* entailed by a complete description of the fundamental features of reality (because we will be able to work out from such a description whether the identifying properties are instantiated), (ii) that a coherent but false identity statement involving such a concept will be verified at some possible world (the world where the identifying properties are instantiated by some thing other than the referent). Whilst there may be no direct implication from Nida-Rümelin's theory of concept possession to Chalmers's assumptions (A) and (B), her theory does exclude her from the kind of meta-semantic view in which the denial of (A) and (B) is most plausible.

There is another respect in which my argument relies on less contentious premises than Nida-Rümelin's. Her argument relies on the premise that phenomenal concepts are *transparent*. As I argue in my response to *Objection 1*** above, my argument requires only that phenomenal concepts are *translucent*.

6. Conclusion: Three Theories of Mental Concepts

A posteriori physicalists would like to present the options to choose from when adopting a theory of consciousness as follows:

³¹It is clear I think from the context that Nida-Rümelin intends something like this story to generalize to all concepts.

Two theories of mental concepts

Analytic functionalism: All mental concepts conceive of conscious states in terms of their causal role.

Conceptual dualism: There are certain mental concepts, i.e. phenomenal concepts, which do not conceive of conscious states in terms of their causal role.

Two theories of conscious states

Physicalism: Conscious states are physical or functional states.

Dualism: Conscious states are not physical or functional states.

A posteriori physicalists (*conceptual dualism + physicalism*) claim that their view combines the best theory of mental concepts, i.e. the one which fits best with our armchair intuitions about our mental concepts, with the best theory of the metaphysics of conscious states, i.e. the most economical or the best suited for making sense of mental causation.

Actually, there are at least three theories of our mental concepts:

A: Conceptual dualism + transparency/translucency: (i) There are certain mental concepts, i.e. phenomenal concepts, which do not conceive of conscious states in terms of their causal role, (ii) phenomenal concepts are transparent/translucent.

B: Conceptual monism + transparency/translucency: (i) All mental concepts conceive of conscious states in terms of their causal role, (ii) all mental concepts are transparent/translucent.³²

C: Conceptual dualism + opacity: (i) There are certain mental concepts, i.e. phenomenal concepts, which do not conceive of conscious states in terms of their causal role, (ii) phenomenal concepts are opaque.

A posteriori physicalists reject *B*, the position accepted by *a priori* physicalists, on the grounds that it is contrary to our armchair intuitions about phenomenal concepts. But their own view *C* is inconsistent with other significant armchair intuitions about our phenomenal concepts, i.e. that our concept of the feeling of pain reveals something of what it is for something to feel pain. The theory of mental concepts put forward by the *a posteriori* physicalist is arguably no less contrary to our carefully considered judgments about our mental concepts than that put forward by the *a priori* physicalist. For the philosopher who shares my intuitions about our mental concepts (*conceptual dualism + transparency/translucency*), in so far as she has reason to trust her carefully considered judgments about her concepts, has reason to reject both *a priori* and *a posteriori* physicalism.

³²I interpret David Lewis as claiming that we have a transparent flaccid concept of *being a pain*, and a mildly opaque rigid concept of *being the actual state which is human pain*, both of which conceive of pain in terms of its causal role [Lewis 1980].

I do not take what I have offered here to constitute a knock-down argument against *a posteriori* physicalism. Our carefully considered intuitions about phenomenal concepts may turn out to be wrong. However, the main attraction of *a posteriori* physicalism is its reputation for agreeing with the dualist about concepts, while disagreeing with the dualist about ontology. I have tried to show in this paper that *a posteriori* physicalism departs from the dualist's intuitive picture of our mental concepts just as radically as does *a priori* physicalism.³³

University of Hertfordshire

Received: August 2009

Revised: January 2010

References

- Armstrong, D. 1968. *A Materialist Theory of Mind*, London: Routledge and Kegan Paul.
- Block, N. and R. Stalnaker, 1999. Conceptual Analysis, Dualism and the Explanatory Gap, *Philosophical Review* 108/1: 1–46.
- Braddon-Mitchell, D. 2003. Qualia and Analytical Conditionals, *Journal of Philosophy* 100/3: 111–35.
- Chalmers, D. J. 1996. *The Conscious Mind: Towards a Fundamental Theory*, Oxford: Oxford University Press.
- Chalmers, D. J. 2002. Consciousness and its Place in Nature, in *Philosophy of Mind: Classical and Contemporary Readings*, ed. D. J. Chalmers, Oxford and New York: Oxford University Press: 247–72.
- Chalmers, D. J. 2009. The Two-Dimensional Argument Against Materialism, in *Oxford Handbook of the Philosophy of Mind*, ed. B. McLaughlin, Oxford: Oxford University Press: 313–39.
- Dennett, D. 1991. *Consciousness Explained*, Boston: Little, Brown.
- Dretske, F. 1994. If You Can't Make One, You Don't Know How it Works, in *Midwest Studies in Philosophy* 19/1: 468–82.
- Dretske, F. 1995. *Naturalizing the Mind*, Cambridge, MA: MIT Press.
- Eddington, A. 1928. *The Nature of the Physical World*, New York: Macmillan.
- Feigl, H. 1967. *The "Mental" and the "Physical": the Essay and a Postscript*, Minneapolis: University of Minnesota Press.
- Fraassen, B. C. van 2006. Structure: Its Shadow and Substance, *The British Journal for the Philosophy of Science* 57/2: 275–307.
- Goff, P. forthcoming. Ghosts and Sparse Properties: Why Physicalists Have More to Fear from Ghosts than Zombies, *Philosophy and Phenomenological Research*.
- Griffen, D. R. 1998. *Unsnarling the World-Knot: Consciousness, Freedom, and the Mind-Body Problem*, Minneapolis: University of Minnesota Press.
- Harman, G. 1990. The Intrinsic Quality of Experience, *Philosophical Perspectives* 4: 31–52.
- Hawthorne, J. 2002. Advice for Physicalists, *Philosophical Studies* 109/1: 17–52.
- Hill, C. 1997. Imaginability, Conceivability, Possibility and the Mind-Body Problem, *Philosophical Studies* 87/1: 61–85.
- Hill, C. and B. McLaughlin 1999. There are Fewer Things in Reality than are Dreamt of in Chalmers's Philosophy, *Philosophy and Phenomenological Research* 59/2: 445–54.
- Jackson, F. 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*, Oxford: Clarendon Press.
- Kripke, S. 1980. *Naming and Necessity*, Oxford: Blackwell.
- Ladyman, J. and D. Ross (with D. Spurrett and J. Collier) 2007. *Every Thing Must Go: Metaphysics Naturalized*, Oxford: Oxford University Press.
- Levine, J. 1983. Materialism and Qualia: the Explanatory Gap, *Pacific Philosophical Quarterly* 64/9: 354–61.
- Lewis, D. 1966. An Argument for the Identity Theory, *Journal of Philosophy* 63/1: 17–25.
- Lewis, D. 1970. How to Define Theoretical Terms, *Journal of Philosophy* 67/13: 427–46.
- Lewis, D. 1980. Mad Pain and Martian Pain, in *Readings in the Philosophy of Psychology, Vol. I*, ed. N. Block, Harvard: Harvard University Press: 216–22.
- Lewis, D. 1994. Reduction of Mind, in *Companion to the Philosophy of Mind*, ed. S. Guttenplan, Oxford: Blackwell: 412–31.
- Lewis, D. 1995. Should a Materialist Believe in Qualia?, *Australasian Journal of Philosophy* 73/1: 412–31.
- Loar, B. 1990. Phenomenal States, in *Philosophical Perspectives* 4: 81–108.
- Loar, B. 2003. Qualia, Properties, Modality, *Philosophical Issues* 13/1: 113–29.
- Lockwood, M. 1989. *Mind, Brain and the Quantum*, Oxford: Oxford University Press.
- Lycan, W. G. 1996. *Consciousness and Experience*, Cambridge, MA: MIT Press.

³³I owe a huge debt of gratitude to David Papineau for invaluable assistance over the last five years in developing my arguments against him, and to Galen Strawson for giving me the confidence to pursue a wacky view just because it happens to be true. I am also grateful for the comments of the journal's anonymous reviewers, which resulted in a much improved paper. I revised this paper while I was a Post-Doctoral Research Fellow in the AHRC project *Phenomenal Qualities*, at the University of Hertfordshire.

- Maxwell, G. 1979. Rigid Designators and Mind-Brain Identity, in *Perception and Cognition*, ed. C. W. Savage, Minneapolis: University of Minnesota Press: 365–404.
- Millikan, R. G. 1984. *Language, Thought, and Other Biological Categories*, Cambridge, MA: Bradford Books/MIT Press.
- Nida-Rümelin, M. 2007. Grasping Phenomenal Properties, in *Phenomenal Concepts and Phenomenal Knowledge*, ed. T. Alter & S. Walter, New York: Oxford University Press: 307–36.
- Papineau, D. 1993a. Physicalism, Consciousness, and the Antipathetic Fallacy, *Australasian Journal of Philosophy* 71/2: 169–83.
- Papineau, D. 1993b. *Philosophical Naturalism*, Oxford: Blackwell.
- Papineau, D. 2002. *Thinking about Consciousness*, Clarendon Press: Oxford.
- Papineau, D. 2006. Comments on Galen Strawson, *Journal of Consciousness Studies* 13/6: 100–9.
- Papineau, D. 2007. Phenomenal Concepts and Perceptual Concepts, in *Phenomenal Concepts and Phenomenal Knowledge*, ed. T. Alter & S. Walter, New York: Oxford University Press: 307–36.
- Perry, J. 2001. *Knowledge, Possibility and Consciousness*, Cambridge, MA: MIT Press.
- Rey, G. 1995. Towards a Projectivist Account of Conscious Experience, in *Conscious Experience*, ed. T. Metzinger, Paderborn: Ferdinand Schöningh: 123–42.
- Russell, B. 1927. *The Analysis of Matter*, London: Kegan Paul.
- Stalnaker, R. 2002. What is it Like to be a Zombie?, in *Conceivability and Possibility*, ed. T. S. Gendler & J. Hawthorne, Oxford: Oxford University Press: 385–400.
- Stoljar, D. 2001. Two Conceptions of the Physical, *Philosophy and Phenomenological Research* 62/2: 253–81.
- Strawson, G. 1994. *Mental Reality*, Cambridge, MA and London: MIT Press.
- Strawson, G. 2003. Real Materialism, in *Chomsky and his Critics*, ed. L. Antony and N. Hornstein, Oxford: Blackwell: 48–88.
- Strawson, G. 2006. Realistic Monism: Why Physicalism Entails Panpsychism, in *Consciousness and its Place in Nature*, ed. Anthony Freeman, Exeter: Imprint Academic: 3–31.
- Tye, M. 1995. *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*, Cambridge, MA: MIT Press.