

Mark Bradley

The Causal Efficacy of Qualia

Abstract: *Qualia are the elements of phenomenal consciousness — the raw feels which constitute what it is like to be in a conscious mental state. Some claim that qualia are epiphenomenal properties — mere by-products of brain function which are causally inert. Though this is an implausible theory, it is difficult to show that it is false. Here I present an ad hominem argument — the argument from coincidence — which shows that epiphenomenalism about qualia is explanatorily deficient because it leaves unexplained a highly improbable state of affairs, and so we have good reason to suppose that qualitative properties must be causally efficacious, and must be so in virtue of their qualitative nature.*

Introduction

Qualia are the elements of phenomenal consciousness — the raw feels which constitute what it is like to be in a conscious mental state, which constitute experience itself. Imagine the taste of lemonade, the flash of a camera, the agony from a burn, and the nervous feeling brought about by some forthcoming event. All involve experiences which feel a certain way, all involve phenomenal properties or qualities, all involve *qualia*. The lemonade tastes fizzy, tangy, or sharp; the camera flash is bright and blinding, instantaneous and all-encompassing; the burn on my hand is persistent and demanding, excruciating and sickening; and my nervousness is fluttery and flighty, distracting and distressing.

Some claim that qualia are not physical or functional properties, that they cannot be reduced to or explained in terms of physical entities. A contemporary version of this view is naturalistic dualism: the

Correspondence:
Email: m.c.bradley@leeds.ac.uk

idea that qualia are non-physical properties which are nonetheless fundamental natural properties of the universe. There is nothing supernatural or spooky about them, but they cannot and should not be counted as among the physical entities that make up the world.

But if this is true it raises the question of whether qualia are causally efficacious properties, whether they have any causal (or explanatory) role to play at all in the world. If the physical properties of our brains or central nervous systems are fully sufficient for the causation and explanation of our behaviour, aren't qualia merely redundant by-products of these physical systems, real yet causally impotent features of the universe? Aren't they, in short, epiphenomena? Some lean toward such a view. In the same paper in which his knowledge argument appeared, Frank Jackson claimed that qualia are causally impotent properties, and David Chalmers claims that his conceivability argument shows that some form of property dualism is true, and writes that 'I am not endorsing epiphenomenalism, but I regard it as one of the three serious options that remain... [t]he other two are interactionism and a Russellian "panprotopsychism"' (Chalmers, 2003, p. 255, n. 17).¹

Of course epiphenomenalism is a very unpopular view, since most of us think that mental properties have an essential causal role to play in the production of behaviour; if they didn't play such a role then there would be some very unsavoury consequences for our notions of human agency, free will, and responsibility. Such considerations, however, do not show that epiphenomenalism is false — running counter to common sense and coming into conflict with many of the views that we have of ourselves and our place in the world does not really add up to a cogent argument. Indeed, it is a very difficult theory to argue against, partly because it can be claimed that for the explanation of any particular kind of behaviour there is always a wholly physical explanation available.

Given this, what I shall attempt to do below is present an argument which shows that epiphenomenalism about qualia leaves unexplained what would be a highly improbable state of affairs and so is an *explanatorily deficient* theory.² We therefore have good reason to suppose that qualitative properties must be causally efficacious properties,

[1] See also Chalmers (1999, p. 493). Jackson's defence of epiphenomenalism is in his (1982). He has now changed his view (2004). Other epiphenomenalists about qualia include Campbell (1984) and Robinson (1988; 2006; 2007). See also the papers in the special edition of *Journal of Consciousness Studies* (2006) for a good discussion of the issues.

[2] In what follows, the kind of epiphenomenalism at issue involves non-physical properties. It is not a thesis about the content of otherwise physically explainable mental states.

and must be so *in virtue of* their qualitative nature. I shall not be proposing a form of interactive naturalistic dualism, but merely presenting an *ad hominem* argument to those who present an epiphenomenal form of this new kind of dualism.

Let's first get clear on exactly what we mean by epiphenomenalism about qualia and its motivation. Epiphenomenalism in the sense being used here is a form of (naturalistic) property dualism about qualitative properties, and so qualia are taken to be non-physical properties. The motivation for such a dualism invariably has its basis in the irreducibility of qualitative properties. Qualia are or possess certain features which make it impossible to reduce them to physical properties — amongst them might be the intrinsic nature of such properties, their subjectivity, or perhaps their immediate apprehensibility in experience. But whatever the features, there must be at least one which makes a physical reduction impossible, otherwise there would be no motivation for epiphenomenalism in the first place. Let's say then that epiphenomenalism is committed to the following principle:

Irreducibility of Qualia [IQ]: Qualia are or have special features which make them irreducible non-physical properties.

And let's characterize qualia epiphenomenalism itself along the following lines:

Epiphenomenalism about Qualia [EQ]: For any instance of behaviour [b] of any individual x at time t₃, there is a physical property [p] instantiated at time t₂, and a physical stimulus [s] which occurs at time t₁ such that [s] causally contributes to the instantiation of [p] and [p] causally contributes to the occurrence of [b]; but [p] also causes the instantiation of qualitative property [q] some time after t₂, though [q] does not (and cannot) cause anything in turn.

The Argument from Coincidence

I shall first present this argument and then defend each premise in turn. Throughout I'll stick with the example of pain as it serves as the most vivid example of a qualitative property to which the argument applies. Whether the argument will apply to other kinds of qualitative property will be briefly discussed after the argument has been defended:


1. There are a series of correlations between a certain set of physical properties [P] and a certain set of qualitative properties [Q];

2. If there are a series of correlations between a certain set of physical properties [P] and a certain set of qualitative properties [Q], then there must be an explanation for these correlations;
3. Epiphenomenalism cannot provide an explanation for these correlations;
4. If a theory cannot provide an explanation for these correlations, then it is **explanatorily deficient**;
5. So epiphenomenalism is explanatorily deficient.

I call this the argument from coincidence since it trades upon the fact that there are a large number of correlations between properties of type [P] and properties of type [Q] which add up to a large scale coincidence. It will be argued that epiphenomenalism cannot provide an explanation for this coincidence, while those who hold that qualia are causally efficacious can. It follows that a theory which attributes a causal role to qualia is to be preferred over a theory which does not, other things being equal.³ Let's look at each premise in turn.

Premise 1: *There are a series of correlations between a certain set of physical properties [P] and a certain set of qualitative properties [Q].*

The diagram below captures this:



It illustrates the idea that for a particular individual x in a particular situation in which x receives a physical detrimental stimulus (x 's hand being burned, for example), and there is actual or potential tissue damage [s] to x , there is a physical property [p] which is partly caused by [s]. Physical property [p] might be the firing of a whole group of neurons, it really doesn't matter for the purposes of the argument. Now [p] also causally contributes to a further physical (behavioural) property [b] which consists of certain kinds of appropriate damage

[3] This is a reworked version of the argument which appears in chapter 7 of my book (2010). So far as I can tell, the only other place where this kind of argument appears in the literature is in James (1879, pp. 17–8; and 1890, pp. 143–4), where he devotes just a few lines to it. I have recently found something similar in Pauen (2000, pp. 57–8), who also says very little. Robinson (2004; 2007) is the only philosopher that I know of who has put forward a defence (see below). Latham (2000) discusses the example of a pleasure-pain inverted world, but uses it to argue for the logical impossibility of inverted qualia.

limitation behaviour, but [p] also causes the instantiation of a qualitative property [q] — an unpleasant experience or pain. According to epiphenomenalism, [q] has no causal effects.

Instances such as [s] belong to the set of instances of tissue damage [S], and instances of behaviour [b] belong to the set of damage limitation behaviours [B]. It would be difficult, if not impossible, to provide criteria for the conditions that must be satisfied in order for a physical property to be included in either [S] or [B]. I think that, intuitively, for something to be a member of [S] it must present a current or potential threat of injury to an individual, or compromise that individual's survival or well-being in some way. Examples are burns, blows, stings, lacerations, and fractures. For a physical property to be a member of [B] it must be, or must be a disposition to engage in, certain kinds of typical appropriate behaviour, such as the indication of instances of [S] (squealing or crying), the avoidance of instances of [S] (retraction of one's hand from a flame), and the nursing of instances of [S] (pouring cold water over a scald). The exact kind of behaviour which results from instances of [S] will be a function of a number of different factors, such as the location and the duration of the stimulus, and cognitive factors such as anticipation and attention.

Instances such as [p] belong to the set [P]. All and only its members are partly caused by instances of [S] and in turn causally contribute to instances of [B]. Instances such as [q] belong to the set [Q]. All and only its members are qualitative properties caused by instances of [P]. We classify members of [Q] as unpleasant experiences or pains.⁴

So instances of [Q] are correlated with instances of [S], [P], and [B]. The correlation of instances of [Q] with instances of [S] can be explained by appeal to instances of [P]. The properties which cause instances of [Q] are instances of [P] which also happen to be the kind of properties that are partly caused by instances of [S]. So whenever instances of [S] occur, instances of [P] occur, and whenever instances of [P] occur, instances of [Q] occur. The correlation of instances of [Q] with instances of [B] can also be explained by appeal to instances of [P], since instances of [P] are also the kinds of property which partly cause instances of [B]. So whenever instances of [P] occur, instances of [B] occur, and because whenever instances of [P] occur, instances of [Q] occur, whenever instances of [Q] occur, instances of [B] occur, though, according to epiphenomenalism, instances of [Q]

[4] There might be a few exceptions here, such as those instances involved in masochistic behaviour. But there are a sufficient number of cases that are the norm for the argument to go through.

play no causal part in bringing about instances of [B].⁵ We only need to focus on the correlation of instances of [P] with instances of [Q] from now on.

Premise 2: If there are a series of correlations between a certain set of physical properties [P] and a certain set of qualitative properties [Q], then there must be an explanation for these correlations.

We now have a condition that must be satisfied in order for a physical property to be included in the class [P]. All and only instances of [P] are partly caused by instances of [S] and in turn causally contribute to instances of [B]. This ensures that membership of [P] is not arbitrary and that all members of the set have something significant in common. We have also seen that instances of [P] have a certain consequence: they cause instances of [Q], members of which are properly describable as unpleasant experiences or pains. This condition and this consequence are correlated thus: physical properties of the set [P] always bring about qualitative properties of the set [Q]. Given that we have this correlation, we now require an explanation for why it exists: why it is that there is (almost) always a correlation between the condition and the consequence. In fact the explanation would need to do two things: first, explain why it is that only instances of [P] bring about instances of [Q] — why it is, for example, that other kinds of physical property such as instances of the set [P1] or the set [P2] which might be neural properties which are not instantiated between instances of [S] and instances of [B], do not bring about instances of [Q]; and second, why it is, for example, that instances of [P] do not bring about instances of other kinds of qualitative property, such as instances of the set [Q1] of pleasurable experiences, or the set [Q2] of fearful experiences, or the set [Q3] of experiences of disgust, and so on. Consider our example of a painful experience of being burned. This experience will be correlated with an instance of [P]. Another kind of experience, say, an experience of orgasm, will be correlated with an instance of another set which has no overlapping members with [P], say, [P4], whose condition of membership is that all and only its instances must play an intermediate causal role between contact with certain kinds of beneficial stimuli and certain kinds of behaviour that is appropriate for such stimuli. Instances of [P4] cause, other things being equal, pleasurable qualitative properties of the set [Q4]. But why is it that no

[5] It should also be noted that instances of [P] and [B] might be brought about by physical states other than instances of [S], such as, for example, a case of artificial stimulation of the brain. Again, these cases will be the exception rather than the rule.

(or at least very few) instances of [P] cause instances of [Q4], and why is it that no (or at least very few) instances of [P4] cause instances of [Q]? The fact that these states of affairs do not obtain requires an explanation.

It is important to note here two points about the explanation required. First, because the existence of the set [P] has simply been stipulated, it might be thought that the demand for an explanation is in some way empty. But this is not so, since membership of the set [P] has the necessary and sufficient conditions that each of its members must be partly caused by the right kind of stimulus and result in the right kind of behaviour. Membership of [P] is not arbitrary. Second, the question is not how instances of [P] cause instances of [Q] — that is the question which perhaps might be answered by appeal to some kind of fundamental psychophysical law. Rather, the question is *why* this particular group of physical properties always seem to cause qualitative properties that we experience as unpleasant (pains), rather than properties of any other kind (and different groups of physical properties do not), something which really wouldn't matter if those qualitative properties couldn't have any causal contact with us. It wouldn't and couldn't matter if those physical properties caused any other kind of qualitative property, since the latter wouldn't and couldn't have any effects on our behaviour. In other words, it really wouldn't make any non-mental difference if individuals who were burned underwent qualitative experiences not of being burned, but only of jealousy, or of fear, or of ecstasy, etc.⁶

Premise 3: *Epiphenomenalism cannot provide an explanation for these correlations.*

One might use one of the assumptions of the argument to explain why the correlation holds. Remember that all properties of [P] have something in common — they are all partly caused by instances of [S] and, in turn, causally contribute to instances of [B]. So it might be argued that the kinds of physical property which come about from instances of [S] and which in turn causally contribute to instances of [B] are just the right kinds of property to bring about certain kinds of qualitative experience, and those experiences are classified as pains. In order for a painful experience to come about you need a certain kind of physical property — just those properties which are likely to be the effects of

[6] And just aside, would we really expect to see sexually promiscuous individuals seeking out casual sex if it hurt so much, and would we really expect to see individuals refraining from causing themselves to have at least minor burns if it were so pleasurable?

certain stimuli and are likely to have certain physical effects on behaviour. Such properties may have been selected for because of the advantages that they confer on the individual who has them. We find them playing a causal role between instances of [S] and instances of [B]. Any instance of a physical property which is embedded in the causal network in this way will bring about the instantiation of qualitative properties of the kind [Q] (and not any other kind of qualitative property).⁷

The short response to this claim is that it does not add up to a satisfactory explanation, since no reason is provided for why all and only those physical properties which cause painful qualitative experiences are the ones which play the intermediate causal role between instances of [S] and instances of [B]. We still have no explanation for the fact that such properties only cause painful experiences when such experiences have no causal effects, and the claim that these physical properties have been selected for provides no reason why these physical properties should also carry along with them qualitative properties which are always of the unpleasant kind.

A second response might involve the claim that we classify our painful experiences in a negative way due to the very fact that we associate them with negative effects on our body — with tissue damage caused by detrimental stimuli. The reason that there is a correlation between members of [P] and painful experiences is simply due to the fact that all members of [P] have been brought about in this negative way. Members of [Q] are those kinds of property that we classify in a negative way — perhaps in terms of negative judgments or evaluations — because of the very fact that they are caused by members of [S] and thus [P]. However, this response would violate one of the assumptions that the epiphenomenalist is committed to above. Remember that for the epiphenomenalist (or indeed any other kind of property dualist) qualitative properties are irreducible, so painful experiences are *irreducibly* painful — they cannot be exhaustively analysed in terms of anything else. To analyse them in terms of further properties, such as negative judgments or evaluations which we might make about them, would violate the property dualist thesis [IQ] and would arguably confer some kind of causal role upon qualitative properties — an obvious violation of the very thesis [EQ] which epiphenomenalism

[7] Robinson (2004; 2006; 2007) seems to use this kind of defence.

is characterized in terms of. So the epiphenomenalist cannot provide an explanation by taking this route unless she is willing to give up [IQ].⁸

A similar kind of problem will arise for a response which makes a distinction between the sensory and motivational features of painful qualitative properties. The sensory core of a pain might be the thumping, pricking, or stabbing experience of the sensation, for example. The motivational feature can be thought of as the ‘hurtfulness’ of that sensation — its ‘awfulness’ or ‘unpleasantness’. It might be claimed that qualia have a sensory core, but that any motivational features which are usually associated with them are in fact distinct and can be given some kind of reductive analysis in terms of dispositions or attitudes to the sensory core. Since the motivational aspect of pain can be divorced from its sensory core, it can be claimed that pains aren’t irreducibly *awful* or *unpleasant* after all, and so there isn’t anything that all the members of set [Q] really have in common. The argument from coincidence cannot get off the ground since there are no significant correlations to explain in the first place.

In one sense, this seems to violate [IQ] and so would be contrary to the epiphenomenalist view. There is something that it is like for a pain to be unpleasant or awful which I think the epiphenomenalist would agree is not explicable in reductive terms. And even if the motivational aspect of pain could be divorced from its sensory core, and there might be at least one sense in which qualitative states are physically reducible (perhaps in terms of attitudes or dispositions), why stop at the motivational aspect? If the ‘awfulness’ or ‘unpleasantness’ of a pain could be physically reduced, why not say that the sensory core might be too? Again, taking this approach seems to threaten the motivation for property dualism in the first place.

Premise 4: If a theory cannot provide an explanation for these correlations, then it is explanatorily deficient.

We have a number of correlations of the form:

an instance of [P] causes an instance of [Q]

Such correlations require an explanation because without one we are faced with the existence of a highly improbable state of affairs — a massive coincidence. Any theory which cannot provide such an explanation will be explanatorily deficient. Epiphenomenalism cannot provide such an explanation, so epiphenomenalism is explanatorily

[8] Also, it is surely counter-intuitive to claim that we classify pains only in virtue of their association with tissue damage. We classify pains as pains because of their painfulness.

deficient. The conclusion at 5 follows. Of course this argument can be taken as a challenge to provide such an explanation, but then the burden of proof lies with the epiphenomenalist to do just that.

Moreover, an explanation for the existence of these correlations *can* be provided by a non-epiphenomenal version of naturalistic dualism — that is, by a theory which involves the claim that qualia *are* causally efficacious. It is due to the very fact that qualia are causally efficacious that such correlations can be made. If painful experiences have some causal role to play in the indication of instances of [S] and the explanation of behaviour such as those included in [B], then the reason that instances of [P] are always correlated with instances of [Q] is that instances of [P] have been selected for because they serve to indicate to a subject *x* the presence of instances of [S] by causing the instantiation of an instance of [Q] in *x*. Instances of [P] bring about instances of [Q] because only by doing so can instances of [S] bring about instances of [B]. The nature of the qualitative property that is caused is *itself* causally effective in bringing about an instance of behaviour of type [B] which it is appropriate to take in such situations. The nature of the properties that are members of [Q] are ‘unpleasant’ or ‘awful’ so that they cause the individuals who have them to take the behaviour required in order that such experiences either cease or are at least assuaged in some way, and this has the secondary function of minimizing the level of tissue damage or harm which the individual is subjected to. The qualitative properties which are involved we call pains. So qualitative properties such as pains are causally efficacious *qua* their qualitative nature. They are causally efficacious *in virtue of* the qualitative nature that they have.⁹

Not only does this show that for certain kinds of behaviour to come about qualia must be *causally efficacious* properties, but also that in order to explain certain kinds of behaviour we need to invoke the nature of qualitative properties, and so qualia are also *explanatorily*

[9] Popper put forward the argument that consciousness must be causally efficacious because it evolved in organisms, and that every trait of an organism must have some survival value and so be causally efficacious (Popper and Eccles, 1977, pp. 73–4). But Broad anticipated this kind of argument way back when he wrote that ‘natural selection is a purely negative process, it simply tends to eliminate individuals and species which have variations unfavourable to survival... the possession of mind is not unfavourable, it simply makes no difference’ (Broad, 1925, p. 118). Jackson (1982) also points to the evolution of traits which have no survival value. Indeed some traits, such as the heavy coat of the polar bear, actually seem to be detrimental to survival. Traits such as a heavy coat may be necessary by-products of traits that *are* conducive to survival, such as a thick (and therefore warm) coat. See also Gould and Lewontin (1979). The argument from coincidence, on the other hand, does not *assume* that qualia must have some evolutionary benefit, it shows that they *must* have some.

relevant too. Both of these theses are rejected by the epiphenomenalist. So if an epiphenomenal theory cannot provide an explanation for why certain correlations between physical properties and qualitative properties hold, but a non-epiphenomenal theory can, we have good reason to prefer the non-epiphenomenal theory, and the idea that qualia *are* causally efficacious properties.

This conclusion can only be established if the claim that pain qualia are causally efficacious can be shown to be an instance of the more general thesis that *all* qualia are causally efficacious. One might have reason to doubt this. What exactly is the *distinctive* causal role of an experience of reddishness, for example? Considerations of spectrum inversion (red for green, and blue for yellow) might lead one to ask whether these kinds of experience are causally efficacious at all. In response, there are three things we might say: (i) we might claim that the distinctive causal role of a particular qualitative property depends on the rest of the qualitative and cognitive architecture of which it is part, so that a reddish quale will cause you to stop at a red traffic light if you understand what the red light represents, and if you wish to drive safely, etc.; (ii) we might, with Lowe (2000, pp. 58–9), claim that qualia are causally efficacious though they do not necessarily have *distinctive* causal roles. This fact does not, however, prevent us from attributing the *same* causal roles to two different kinds of qualia, and so maintaining that they are causally efficacious nonetheless; (iii) it seems natural to extend the non-epiphenomenal explanation to other cases of qualitative experience even if we cannot distinguish a particular causal role for each type of quale. In fact it seems that we would need a further explanation if we did not take this option — an explanation for why some kinds of qualia are causally efficacious while others are not.

In sum then, the argument from coincidence shows that either qualia must be causally efficacious or that there must be some explanation for what would be a highly improbable state of affairs. Without such an explanation, it seems that the non-epiphenomenal view has a theoretical advantage over the epiphenomenal one.¹⁰

[10] Of course, there are lots of problems for what would still be a form of naturalistic dualism, such as the causal closure of the physical domain, the problem of overdetermination, and the question of the causal mechanism involved, etc. But my aim here has only been to show that epiphenomenalism about qualia faces another serious problem.

The Function of Qualia?

One implication of the argument from coincidence is that it suggests what the general role of qualia might be. It was claimed that qualia are causally efficacious *qua* their qualitative nature — that it is due to the very nature of qualia that they bring about the effects that they do. The role seems to be that qualia represent certain parts of the world and motivate us to perform certain kinds of action that are typically conducive toward our survival. There may be occasions when they fail to do this — the existence of conflicting attitudes, for example, or the presence of mental illness, the use of drugs, or any number of other reasons. But in the typical case, the role or function of qualia seems to be twofold: there is the primary function of representing the environment and stimulating the subject to behave in appropriate ways, and there is the secondary function of conferring some kind of survival advantage on the individuals who have them. The secondary function is a result of the primary one. These considerations seem to be borne out by a wide variety of stimuli and behavioural reactions, and across a number of different species. One reason why we have sex is because it is pleasurable, but it is also conducive to our survival (or at least to the survival of a sufficient number of our genes); one reason we avoid flames is because they cause pain, but also because they are detrimental to our general well-being; one reason why we drink orange juice is because it tastes good, but also because of the nutritional benefit that it provides; one reason we recoil at rotting meat is because it smells dreadful, but also because it poses a risk of contamination and disease, etc. On a speculative note, then, qualia have a dual role: their function is to represent and motivate, and thus confer some survival advantage on those who have them. The important point is that qualia do this in virtue of the nature that they have and, typically, we negotiate our way around the world more successfully when we use them as our guides.

Acknowledgments

I would like to thank two anonymous reviewers for their helpful comments on an earlier draft.

References

- Bradley, M. (2010) *The Elements of Phenomenal Consciousness*, Saarbrücken: Lambert Academic Publishing.
- Broad, C.D. (1925) *The Mind and Its Place in Nature*, Cambridge: Routledge & Kegan Paul.
- Campbell, K. (1984) *Body and Mind*, revised ed., Notre Dame, IN: University of Notre Dame Press.

- Chalmers, D. (1999) Materialism and the metaphysics of modality, *Philosophy and Phenomenological Research*, **59**, pp. 473–496.
- Chalmers, D. (2003) The content and epistemology of phenomenal belief, in Smith, Q. & Jovic, A. (eds.) *Consciousness: New Philosophical Perspectives*, pp. 220–272, New York: Oxford University Press.
- Gould, S. & Lewontin, R. (1979) The spandrels of San Marco and the Panglossian Paradigm, *Proceedings of the Royal Society of London B*, **205**, pp. 581–598.
- Jackson, F. (1982) Epiphenomenal qualia, in Ludlow, P., Nagasawa, Y. & Stoljar, D. (eds.), *There's Something About Mary: Essays on Phenomenal Consciousness and Frack Jackson's Knowledge Argument*, pp. 39–50, Cambridge, MA: MIT Press.
- Jackson, F. (2004) Mind and illusion, in Ludlow, P., Nagasawa, Y. & Stoljar, D. (eds.), *There's Something About Mary: Essays on Phenomenal Consciousness and Frack Jackson's Knowledge Argument*, pp. 421–442, Cambridge, MA: MIT Press.
- James, W. (1879) Are we automata?, *Mind*, **4**, pp. 1–22.
- James, W. (1890) *The Principles of Psychology*, vols. 1 & 2, New York: Dover Publications.
- Journal of Consciousness Studies* (2006) Special edition on epiphenomenalism, **3** (1–2), pp. 1–175.
- Latham, N. (2000) Chalmers on the addition of consciousness to the physical world, *Philosophical Studies*, **98**, pp. 71–97.
- Lowe, J. (2000) *An Introduction to the Philosophy of Mind*, Cambridge: Cambridge University Press.
- Ludlow, P., Nagasawa, Y. & Stoljar, D. (eds.) (2004) *There's Something About Mary: Essays on Phenomenal Consciousness and Frank Jackson's Knowledge Argument*, Cambridge, MA: MIT Press.
- Pauen, M. (2000) Painless pain: Property dualism and the causal role of phenomenal consciousness, *American Philosophical Quarterly*, **37**, pp. 51–63.
- Popper, K. & Eccles, J. (1977) *The Self and Its Brain*, Cambridge: Routledge & Kegan Paul.
- Robinson, W. (1988) *Brains and People*, Philadelphia, PA: Temple University Press.
- Robinson, W. (2004) *Understanding Phenomenal Consciousness*, Cambridge: Cambridge University Press.
- Robinson, W. (2006) Knowing epiphenomena, *Journal of Consciousness Studies*, **13** (1–2), pp. 85–100.
- Robinson, W. (2007) Evolution and epiphenomenalism, *Journal of Consciousness Studies*, **14** (11), pp. 27–42.

Paper received April 2010; revised September 2010.