

William S. Robinson

Developing Dualism and Approaching the Hard Problem

Abstract: *Arguments for property dualism offer a strong challenge to materialist views, but even if they are regarded as successful, a large task remains, namely, to develop a positive account of the place of non-physical properties in the world — one that holds some promise of eventual satisfaction regarding the hard problem. After noting some difficulties in current approaches to this task, this paper outlines one possible line of development for a dualistic view. Like all other suggestions for routes to progress in this area, this one is speculative. However, the empirical findings that would support this line of development lie within current epistemic possibility. Moreover, the conceptual changes that would be required are intelligible from our present vantage point, and have parallels in views that are accepted in science and by non-dualist philosophers.*

The dualism to be discussed in this paper claims that there are phenomenal qualities, and that these are different from, and not compossible from, the properties and relations found in our natural sciences. Phenomenal qualities include the colours, the tastes, the odours, and other sensory qualities; painfulness, itchiness, and other bodily sensation qualities; and qualities that occur in emotions such as anger, jealousy, and elation.¹ Dualists hold that many words have a

Correspondence:

William S. Robinson, Iowa State University. Email: wsrob@iastate.edu

[1] Some may wish to add to this list, likely candidates being feelings in a non-sensory ‘fringe’ such as familiarity or confidence (Mangan, 2001), and specifically cognitive phenomenology (see Bayne and Montague, 2011, for a variety of positions). Issues connected with these additions must be left for another day: the focus of this paper is on sensory qualities.

double use: they can be used for phenomenal qualities and they can also be used for physical properties, typically ones whose instances are causally related to instances of phenomenal qualities. The characteristic claim of the dualism of this paper is not that there are no physical colours — for example, not that there is no set of reflectance profiles that is appropriately classified under a term such as ‘red’. The distinctive claim is that our experiences (a) are non-physical, qualitative events; that is, they are, or essentially involve, instances of phenomenal qualities, i.e. instances of properties that are not instances of physical properties; and (b) these property instances have no further, or hidden, physical nature.²

There are many arguments for property dualism (see, for example, Chalmers, 1996; White, 2007; Goff, 2011; Nida-Rümelin, 2007; Robinson, 2004). This paper does not add to these arguments and, although its author is evidently sympathetic to them, the paper does not, strictly speaking, assume that any of these arguments are successful against materialist views. The stance of this paper is, instead, that arguments for property dualism are sufficiently strong to keep dualism in play, as a view that deserves to be further developed. Churchland (1984, p. 19) has taunted dualists with this thought: ‘Compared to the rich resources and explanatory successes of current materialism, dualism is less a theory of mind than it is an empty space waiting for a genuine theory of mind to be put in it.’ We might call this the charge of Poverty of dualism.³ One aim of this paper is to describe a series of developments that are at present epistemically possible, and that, taken together, would answer the charge of Poverty.

There is currently much interest in panpsychism (see, for example, Strawson, 2009; Seager and Allen-Hermanson, 2012) and in Russellian monism (See, for example, Chalmers, 1996; Stoljar, 2006). Difficulties for both views are well-known. Consciousness for

[2] Most qualitative events are instantiations of several phenomenal qualities. This fact raises some significant issues, but since they do not affect the arguments of this paper, they will be set aside here. They have been addressed in Robinson (2004).

It is common for contemporary philosophers to insist that colour terms (and terms for other sensory qualities) have their primary use, or their only correct use, as denoters of properties of physical objects (or, perhaps, physical surfaces). But it is also common to recognize the intelligibility of a dissenting view, although it is often insisted that a dualist’s experiential properties must properly be denoted by some special term such as ‘red-dish’ or ‘phenomenal red’. The double use of words affirmed in the text is thus not substantively different from familiar understandings of dualism. It differs only in its denial of the appropriateness of special terms for quality words in their phenomenal use.

[3] *Cp.* Owen Flanagan (1992, p. 35): ‘The view that subjectivity is easy to understand if we think of it as part of an immaterial world has proved to be an illusion. The view actually explains nothing...’

inanimate objects, whether fundamental particles or tables, is counter-intuitive, and the combination problem for panpsychism remains formidable. Russellian monism is equally counter-intuitive if it attributes phenomenal qualities to fundamental particles, but if it avoids this problem by appeal to protophenomenal qualities, it faces its own combination problem — how can protophenomenal qualities combine to produce a red sensation, or a pain? (See, for example, Nagasawa, 2008.)

This paper neither elaborates on nor adds to these difficulties. They are mentioned here only as motivations for an alternative development. At present, panpsychism and Russellian monism are proposals whose development is speculative and incompletely worked out. Perhaps they will eventually be given a more satisfactory articulation. In the meantime, it is reasonable to ask whether there can be an articulation of dualism that is alternative to these views. This paper aims to develop one such alternative.

There is one further motivation that will be important for our discussion. A mainspring of writings on consciousness in the last 18 years has been Chalmers' (1995) hard problem. Developing dualism in the present climate requires offering some sort of response to this problem. Some authors, however, seek to dismiss this problem in a way that has been given a succinct expression by O'Hara and Scutt (1996): 'For a problem to be a genuine problem, some sort of idea of a solution must be available...' But the hardness of the hard problem is that it is not easy, where 'easy' means that 'we have a clear idea of how we might go about explaining them' (Chalmers, 1995, p. 203). Putting these two thoughts together suggests the result that the very thing that is supposed to make the hard problem hard actually makes it not a genuine problem.⁴ Let us call this the Hard Problem Problem (HPP) — the problem of maintaining that the hard problem (HP) is a genuine problem.

Neither materialists nor non-materialists need accept the premise of O'Hara and Scutt that leads to the HPP. However, it would count as a virtue of a dualistic theory if it could lay the HPP to rest. We shall see that the way of developing dualism that is outlined in this paper is

[4] Michael Tye (2000, p. 34) raises a similar problem about the explanatory gap: 'Since an explanatory *gap* exists only if there is something unexplained that needs explaining, and something needs explaining only if it can be explained (whether or not it lies within the power of *human beings* to explain it), there is again, no gap.' In *Consciousness Online 4*, a paper by Glenn Carruthers and Elizabeth Schier (2012) has the title 'Dissolving the Hard Problem of Consciousness'. Its first sentence is: 'In this paper we attempt to dissolve worries around the hard problem of consciousness by showing that there is no good argument for the existence of such a problem.'

entitled to claim this virtue. This is not to say that it offers a solution to the HP itself. That solution depends on empirical matters that are not yet known, and on possible conceptual changes that are envisioned in this paper, but not (yet) actually in force. It will, however, be useful to begin by considering a constraint on the solutions of both problems.

1. A Constraint on Solutions of the HP and the HPP

It will be helpful to formulate the HP as consisting of the following two questions. Why, given that there are our neural events and their causal relations to each other and to the world at large, should there be any phenomenally qualified events at all? Why, given that there is a particular kind of neural event (standing in its particular web of causal relations) should there be this particular kind of qualitative event that it yields? ('Yields' is a studiously neutral term covering identity, correlation, and causation. Identity theorists will bridle at 'correlation', representationalists will reject 'identity', and dualists will reject both. But all agree that where there are neural events of certain kinds there are phenomenal qualities, and 'yields' is introduced to have a brief way of capturing this commonality.)

These questions express an intellectual dissatisfaction. For whatever reason, we have an 'intuition of distinctness' (Papineau, 2002) that seems to tell us that phenomenal qualities just cannot be the same as any physical properties, or any combination of them or relation among items with physical properties.⁵ This leaves us with two sets of properties and no obvious reason why one should yield the other. A solution to the HP requires providing a reason why neural events yield qualitative events.

Correlatively, a solution to the HPP requires providing a sketch of such a reason, or at least a proof that such a reason could be worked out.⁶ It requires providing a set of ideas for which it is plausible that further developments could remove our intellectual dissatisfaction. What is needed is a point of view, about which we can imagine our descendants feeling that *of course*, given a defensible view of how things might very well be, our neural events would yield qualitative

[5] The phenomenal concept strategy seeks to explain away this intuition. It is clear that dualists are not convinced by this strategy but, as in other cases, it is beyond the scope of this paper to explain why.

[6] Thanks to an anonymous referee for suggesting the last clause of this sentence. It is, however, difficult to imagine how there could be an existence proof of a reason that would provide intellectual satisfaction even though no sketch of the reason itself could be derived from it.

events, and *of course* particular neural events would yield just the kind of qualitative events that they do.⁷

It may be that such a feeling of understanding of the neural event–qualitative event relation is not possible. Price (1996), following work by Rosch (1994), has argued that causal relations are by nature explanatorily gappy, and that we have several means of hiding this gappiness from ourselves. On this view, what is special about the mind/body case is not that there is an explanatory gap, but that in this case we are uniquely unable to apply our customary means of hiding the gap from ourselves.

Perhaps we will eventually have to concur in some such account. But we should not accept it without first making a serious attempt to achieve intellectual satisfaction about the relation between neural events and qualitative events. The first step of such a project is to set out a possible route by which such satisfaction might be obtainable. This is the project to be pursued in this paper.

2. A Short but Unsuccessful Route

Some philosophers already hold the view that causal relations are essential to properties (e.g. Shoemaker, 1980; 1998; 2007; Swoyer, 1982; Bird, 2005). So, it is not difficult to imagine that our descendants become persuaded to accept this way of thinking. It should be noted that causal essentialism can be adopted by dualists. ‘In all metaphysically possible worlds in which A is instantiated, A causes B (in circumstances C)’ does not imply that A and B are the same property. Indeed, in general, effect properties are distinct from the properties of their causes, so causal essentialism can quite naturally be taken as the assertion of necessary connections between distinct properties.⁸

There is a second belief that we can imagine our descendants to hold, namely, the view that neural events cause qualitative events. It is true that references to the ‘NCC problem’ — the problem of the neural causes of consciousness — are sometimes avoided in favour of references to the ‘NCCC problem’ — i.e. the problem of the neural causes or correlates of consciousness. But there are many philosophers even now who would be willing to embrace the causal arm of this

[7] Note that this formulation does not imply that qualitative events are reducible to the physical. Thus, although such a reduction would likely provide the desired intellectual satisfaction, proposing reduction as a requirement would be to impose a constraint on the form of solutions that dualists (and others) are free to reject.

[8] Of course, this view offends Humean sensibilities. For its defence, see, for example, Wilson (2010).

disjunction, and one can surely imagine the possibility of this view becoming more popular than it is at present.

Someone who holds both of these views might well conclude that the causal relations between certain of our neural events and our qualitative events are metaphysically necessary. We would, of course, find out which neural event properties yield our qualitative events only by empirical investigation. That fact might account for our tendency to think that the NE–QE relations are contingent, and to regard the relation as mysterious. But, on the combination of views being considered, the sense of mystery is founded on an epistemic illusion of contingency. If our descendants become thoroughly persuaded of causal essentialism, they might very well feel that events instantiating our phenomenal qualities *have to occur* given the neural events of the kinds that have arisen in the brains of many animals in the course of evolution (or, *have to be caused by* those neural events).

We shall see that something like this suggestion will play an important role in our sketch of a solution to the HP. As it stands, however, this simple account will not provide what is required of it. The reason is that the most persuasive argument for causal essentialism turns on the idea that transworld identification of properties depends on using causal profiles; but this reason does not evidently apply in the case of phenomenal qualities.⁹

To explain, consider the proposal that in W1 — a possible world different from the actual world — certain items are electrons, but the laws of W1 are different from ours, and these items do not repel each other. A very plausible reaction to this proposal is that there could be no principled reason that would justify the proponent of such a view in calling the indicated items ‘electrons’. They might just as well be counted as neutrinos with a mass different from actual-world neutrinos, or as particles of a kind that do not exist in our world. The point is generalizable to all particles, and to ordinary objects such as bricks and windows. There is no principled way of defending a claim such as ‘In W1 throwing bricks at glass windows does not break them’. It is just as plausible to deny that W1 really contains bricks, or windows that are made of glass, and not some other substance that presents the same superficial appearance.

This kind of point, however, does not seem to generalize to phenomenal qualities. It at least appears to us that ‘blue’ or ‘sweet’, in their uses to describe the qualities in our experience, are not identified

[9] Shoemaker (2007, p. 142): ‘Basically, my case for CTP [causal theory of properties] in my earlier work comes down to the claim that there is no plausible truthmaker for the identification of properties in different worlds having different causal profiles.’

by us through any causal relations. If we want to *communicate* about our experiences, we need to use common predicates, and, very plausibly, that will involve reliance on causal relations, as occurs, for example, in ‘the colour I normally see when I look at a cloudless sky’. But we do not appear to have any difficulty in imagining possible worlds in which looking at cloudless skies causes a different kind of experience. It at least appears to us that our words for phenomenal qualities are rigid designators whose reference is fixed by us through acquaintance with those qualities, without any involvement with assumptions about causal relations.

It may be that this way of looking at our words for phenomenal qualities harbours some kind of error. If we come to see such an error clearly, and learn how to convincingly dispel it from our thinking, the short route might become available. But as long as it seems to us that phenomenal qualities do not have essential causal relations, and as long as we have no reason to think otherwise, we will not have intellectual satisfaction about the proposition that neural events of particular kinds cause particular kinds of qualitative events. We will not be able to feel that particular kinds of neural events *have to yield* the particular kinds of experiences that they do — or, indeed, any kinds of experiences at all.

Progress may, however, be made by attending to some facts about complexity. The following section introduces some key considerations that will be developed in the remainder of the paper.

3. Some Reflections on Complexity

Our qualitative events and their properties are complex in various ways. Colours have degrees of saturation as well as hue and brightness, and hues themselves may be unique or mixed. Sounds have timbre and loudness as well as pitch. Tastes and smells produced by good cooks or good vintners are often described as ‘complex’.

Where we find such complexity, we also find a foothold for plausible suggestions for plurality in its causes. Cooks use several herbs and spices in the same dish, there are many chemicals in wines, musicians distinguish between fundamentals and overtones, artists mix their paints, and so on.

The complexity of our phenomenal qualities is, however, orders of magnitude less than the complexity of neural properties that we must suppose to be instantiated in neural events that yield qualitative events, and orders of magnitude less than the field properties proposed by electromagnetic field theories (Pockett, 2002; McFadden,

2013). There are thousands of distinguishable phenomenal qualities, and therefore thousands of relevantly different kinds of neural event types, or field types, that yield them. So far as we have any reason to believe, the properties of neurons that are relevant to their yielding of qualitative events are their degrees of activation. To obtain sufficient diversity in the yielding neural events, we must therefore suppose that they consist of the activation levels of many neurons.¹⁰ Field theories will require fields that extend over regions that include many neurons.

A consequence of the disparity between the degree of complexity found in phenomenal qualities and in brain events that yield our qualitative events is that we are unlikely to find anything plausible about a connection between a particular kind of qualitative event and the particular kind of brain event that yields it. For dualists, this means that no amount of focusing on the nature of a phenomenal quality is likely to provide any intuitive sense that the cause of its instances should be a brain event of type N, where 'N' is given as a set of activation levels of a group of neurons, or fields that depend on neural activity. It may very well feel natural that mixing pure yellow and pure green should result in a colour that strikes us as more complex than either element in the mixture, but the *relative* simplicity of phenomenal qualities, as compared with the likely brain event causes of their instances, will fail to suggest a natural connection between them, so long as we focus simply on one phenomenal quality and one brain event type.

This line of thinking may be resisted in various ways:

(a) We each have ~100 billion neurons, so perhaps there are single neurons, each of which causes its characteristic phenomenal quality when activated. However, aside from its ill fit with contemporary vision science, this suggestion does not seem capable of providing intellectual satisfaction. 'Activating *this* neuron makes you have a salty taste, activating *that* one makes you see red' seems about as unexplanatory as a claim could be.

(b) Taxonomizing neuron types is an ongoing project, but we already know there are hundreds of distinguishable types. Perhaps a fully developed taxonomy would reveal enough different types so that

[10] We can think of such sets of activation levels as vectors, and different yielders as different points in a high-dimensional vector space. For elaboration of this way of thinking see P.M. Churchland, e.g. his (2012).

Some discussions of sensory systems emphasize distinct regions, e.g. primary visual cortex vs. primary auditory cortex. But nothing in our sciences gives mere difference of location a causal role. The different brain regions must contribute different *kinds* of events, if they are to yield different qualitative events. And these different kinds, so far as we have any reason to believe, must be differences in patterns of activations of sets of neurons, or, perhaps, differences in the values of field strengths over some region.

we could say ‘Activation of type 1 neurons produces blue experiences, activation of type 2 neurons produces an experience as of oregano...’ This is an improvement on (a), in that it relates types to types, rather than individuals to types. But it is intellectually unsatisfying: what does the shape, or number of dendrites, or branching pattern have to do with colours or odours?

(c) A more philosophically motivated resistance may come from representationalists, for whom phenomenal qualities are in external objects, while neural events distinctively track them and mediate discriminative responses.¹¹ This proposal, however, does not evidently lay questions about complexity to rest. Colours in objects will have to be reflectance profiles (or, perhaps, dispositions to produce characteristic reflectance profiles in normal lighting conditions, or molecular structures that ground such dispositions), tastes will be molecular structures, sounds will be complex waves. These properties have far greater complexity than what is apparent in our phenomenal qualities. If we say that phenomenal properties *just are* these complex physical properties, we will have to explain how a property can appear in some way, while failing to appear to have the complexity it actually has. And, of course, we will have to do so without appealing to the idea that external properties are merely causes of experiences that either have or represent different, less complex properties. This fact makes it problematic to say that we represent external properties, but do not represent them *as* they actually are. But if we cannot say that, the only way to preserve representationalism will be to hold that, *in addition* to the complex external properties that are usually taken to cause events in our sensory cortices, there are properties of external objects that are simpler — and therefore different — from reflectance profiles, molecular structures, and so on. This amounts to a form of dualism, since these simpler properties do not figure in any of our natural sciences. Such a view would not likely lead to intellectual satisfaction, because it is unparsimonious, and because it would naturally invite questions of the form ‘Why does complex external property P1 give rise to representations of the particular simpler external property Q1 that it does?’

4. Complexity and Simplicity

Our reflections on complexity can, however, be turned to significant advantage, once we notice that the complexity of neural causes and

[11] Tye (2009) gives an especially clear version of this kind of view.

the relative simplicity of qualitative event types locate neural events and qualitative events with respect to a common property. We have no name for this property, but there is no bar to introducing one. Following Spinoza's use of 'motion and rest' (a pairing of apparent contraries) as a term of art designed to bring out the fact that being at rest is a property on the same scale as being in motion, let us use 'complexity and simplicity' — abbreviated to CAS — to designate the common property that is manifested in a certain kind of complexity in neural events and a certain kind of simplicity in qualitative events. This section and the next explain both the terms in this definition, and the justification for regarding complexity and simplicity as two manifestations of a common property.

'A *certain kind of complexity*' has to be taken seriously. CAS is not manifested in just any kind of complexity, but only in complexity of the kind that occurs in the brain events that we may discover to be the causes of qualitative events. It is, of course, an empirical assumption that there is a property that fits this description.¹² Perhaps, at some point in the future, we will come to see the search for such a property as hopeless. But at present, it seems to many to be a worthwhile research programme: it's the one commonly referred to as the NCC (or NCCC) problem. There are complex neural events that yield qualitative events (either directly or through production of fields), but there are also neural events that must be quite complex in their way, but do not yield qualitative events — or do not do so, at least as far as we have any reason to believe.¹³ Examples are (a) neural events in the visual system that give rise to colour constancy. The apparent colour of a patch in the visual field depends on the character of illumination, but that can be detected only from a visual field that is larger than the patch. What we see thus depends on a considerable amount of neural processing; but the earlier stages of this processing do not, as far as we have any reason to believe, yield qualitative events of their own directly — they do so only by causally contributing to later events that are directly correlated with visual experiences; (b) our brains control secretions to our digestive systems, contractions of digestive muscles, release of substances from glands, and so on. These processes are

[12] If 'disjunctive properties' were allowed, there would be no empirical assumption here — a 'property' would be guaranteed so long as there were some brain event property or other for each event that yielded a qualitative event. This understanding is to be excluded; here and throughout, 'properties' are to be understood as natural, non-disjunctive, projectible commonalities.

[13] But see Schwitzgebel (2013) for dissent on this assumption.

complex, in their way, but have no correlates in consciousness, as far as we have any reason to believe.

It is, of course, logically possible that there is no feature of brain events that is common to all of those that yield qualitative events and absent from all of those that do not. But it is not an unreasonable expectation that such a property will be found. And it is not beyond imagining that it will turn out to be some particular kind of complexity — one that will have as its species complex neural or field properties whose instances yield qualitative events of different kinds. It is *that* kind of complexity that will be the manifestation of CAS in neural events. It is also possible that the properties that distinguish occurrences that yield qualitative events will strike us as having a higher degree of complexity than events such as those in (a) and (b) in the previous paragraph.

Work on the olfactory bulb in rabbits by Skarda and Freeman (1987; see Freeman, 2008) led them to the view that a sniffed odorant would push the olfactory system into the basin of a chaotic attractor distinctive of that odorant. If that view is right, a property that is common to all and only events that yield qualitative events would be a property that picked out a certain class of attractors. A perspicuous description of such a property would have to entail that its instances are very complex events, and it would have to differentiate those events from others that, in their own way, are also very complex.

This paper is, of course, not committed to the view that the approach inspired by Freeman and colleagues is correct. Even so, their work can help us understand the possibility of a structure, involving both temporal duration and a large set of neurons, that has a distinctive form of exceptionally high intricacy.

‘A certain kind of simplicity’ likewise has to be taken seriously. The most fundamental manifestation of this kind of simplicity is the persistence of qualitative sameness through a temporal interval. This simplicity is on display by itself in olfactory experience. Smells can, of course, change fairly rapidly. One can get a whiff of something for a very brief time. But even a whiff has a noticeable duration; it is not a strictly instantaneous occurrence. And if one puts one’s nose up to a flower, one can have a sweet smell that lasts through one’s sniff. That is a kind of constancy, or stasis that persists through time. One’s neurons are engaged in complex firing patterns during the sniff, but the qualitative event is a continuum in time of the fragrance. The fragrance may have distinguishable components, and so not be absolutely simple. But (a) it is quite simple relative to the complexity of the

neural or field events that yield it, and (b) each of the components shares the property of persisting as it is through a duration.

All other qualitative events persist through durations, but some offer other species of the kind of simplicity through which CAS manifests itself in qualitative events. Visual events, for example, exhibit a spatial, as well as a temporal, continuity. Colour experiences are not collections of ‘colour points’ — they are collections of colour *expanses*. The auditory experience produced by tightening a string while bowing it is a glissando that exhibits a continuity of change of pitch. (Of course, that also involves temporal persistence, but the change itself is a simple continuity.)¹⁴

It may be suggested here that there are properties that have the required kind of relative simplicity, but are not instantiated in qualitative events.¹⁵ Now, if considerations of complexity and simplicity are to prove useful to a solution to the HP, this possibility will eventually have to be ruled out. It is, however, *prima facie* plausible that the kind of positive, continuous persistence that is found in qualitative events can be found uniquely in them. Our physical sciences typically progress by finding analyses that result in more, not less, complexity in the properties we attribute to the world. It is thus certainly imaginable that our descendants will continue to think that there is a distinctive kind of simplicity that pertains to the properties in qualitative events, and it will be assumed in what follows that they do so.

CAS, then, is the property that manifests itself in brain events as the kind of complexity indicated, and manifests itself in qualitative events in the ways indicated. Because our science of the brain is not fully developed, we have only an imperfect understanding of the complexity in CAS, and we cannot even be sure that the empirical presuppositions of its coherence will prove to be borne out. This is one of the main reasons why we cannot now have a solution to the HP. But we do understand how our descendants may be able to form a better conception of the relevant kind of complexity of brain events. Let us imagine

[14] Consider a moving dot that changes colour part way through its journey (or, an illusion of this — the colour phi phenomenon). This case involves several species of continuity: simple temporal persistence, spatial spread of the colours, and continuity of motion. This severalness is one kind of complexity, but it must be distinguished from the special kind of complexity that yields qualitative events. Presumably, each aspect of continuity will be underlain by its appropriate species of the kind of neural complexity that yields qualitative events. Their combination in a single event accounts for the co-instantiation of several kinds of continuity, but it is being underlain by the appropriate species of neural complexity that yields each qualitative aspect.

[15] There is a little more on this matter in the next note. Thanks to Bill Lycan for pointing out the need to consider this possibility.

that they have done so, and ask how they might proceed to build on that understanding.

5. Explanatoriness and Generality

Complexity and simplicity, even when qualified in the ways just indicated, seem intuitively to be 'opposites', which suggests that they may 'offset' each other, i.e. more complexity of a certain kind in neural events may require more of the kind of simplicity that is exemplified only in qualitative events. We can express this relation in the following principle.

(ConCAS) CAS is conserved.

If this principle were accepted, then there would be an answer to the general formulation of the HP: there are qualitative events because events with their kinds of simplicity are required by conservation of CAS, given that we have neural or field events that have their particular kind of complexity.¹⁶

Since we have at present only a tenuous grip on CAS, we cannot hope to argue directly that (ConCAS) is explanatory. But we can argue indirectly that conservation principles can be explanatory, and thus (ConCAS) might be so.

The parallel of interest here is the principle of conservation of energy.¹⁷ To begin the argument, let us note that appeals to this principle do seem to provide explanations. It is, to take one example, a well known generalization that water does not flow uphill. But why should that be so? Well, if water did flow uphill (without a pump of some kind) that would violate the conservation of energy. So, *of course* it does not flow uphill. Perhaps we ought not to find such appeals to provide intellectual satisfaction, but it seems evident that we do.

[16] Could CAS be conserved in virtue of the simplicity of simple properties in our physical sciences, such as charge or mass which, after all, persist unchanged through time? No, because it is not remotely plausible that they are systematically related to complexity: they are properties of our neural (and, of course, our other) parts whether those neurons are firing in some complex way or not.

[17] General relativity raises questions about conservation of energy, and it may turn out that the best thing to say is that the principle is false, except where certain conditions are met. However, unless physics is doomed, there will be other principles of similar generality. The argument in the text concerns conditions for intellectual satisfaction, and the principle of conservation of energy has the advantage of being familiar, and often felt to be explanatory. The conclusions drawn here from reflection on conservation of energy can plausibly be derived equally from any accepted physical principles of a comparable level of generality.

Now, what is it about the conservation of energy that can enable it to provide the kind of satisfaction that it does? There are three candidates. It is very general. It is very basic, in the sense that to give it up would require the wholesale revision of our sciences. It is (perhaps) necessary that energy is conserved.

The parallel for the first point would be that (ConCAS) is very general. This seems to be the case: it is a relation among properties. It applies to all instances of the right kind of complexity and the right kind of simplicity. It applies across the whole territory of phenomenality, and the whole territory of physical events with the right sort of complexity.

The parallel for the second point would be that to give up (ConCAS) would require wholesale revision of our account of the neural (or field) event—qualitative event relation. Now, evidently such a remark could be legitimately made only from the point of view of a future in which (ConCAS) has been found to be of explanatory value. But if the point about generality is accepted, it seems possible that our descendants could come to regard (ConCAS) as being as equally basic to our understanding of the world as conservation of energy is to our understanding of physics.¹⁸

In evaluating this suggestion, it is important *not* to fall into the error of supposing that (ConCAS) is basic to *our physics*. Evidently, it is not. CAS is not a property that is presently found in our natural sciences. Neither is it a phenomenal quality. It is a property that is genuinely and irreducibly a common nature of properties that are instantiated in some physical events and also of properties that are instantiated in qualitative events. The point of view we have to imagine is one in which CAS is taken to be a fundamental property of our world, where ‘our world’ includes both physical and qualitative events.

Turning to the third point, it may be held that the principle of conservation of energy is (i) not necessary, or that (ii) it is necessary. If (i), then the explanatoriness of the principle cannot come from its necessity. In that case, we can imagine our descendants finding (ConCAS) sufficiently general and sufficiently basic that they regard it as explanatory.

If conservation of energy is necessary (or, necessary under certain conditions), then if a quantity is not conserved in some world, it cannot be *energy*. It might share some aspects with energy, but it would only be *shenergy*, not the real thing. Now, if one thinks this is right,

[18] Or, at least, to the physics of uncurved spaces. See previous note.

one should also be able to imagine that our descendants think the same way about CAS. If they are inclined to think that *basic* properties have their conservation (or lack thereof) essentially, then they might think that the same holds of genuine CAS.

They could, of course, ask the question ‘Why are we so lucky as to live in a world in which CAS is instantiated (rather than, say, SHMAS, which is not conserved)?’ This question, however, is parallel to ‘Why do we live in a world in which there is energy (rather than, say, shmenergy)?’ We do not have a good answer and, perhaps, we cannot have good answers to such questions, when we get down to the most basic properties of our universe. But lack of such an answer does not seem to prevent appeals to conservation of energy from giving us intellectual satisfaction about such matters as why water does not flow uphill (when unaided by pumps). So, we do not have a reason that shows us that our descendants could not derive intellectual satisfaction from (ConCAS), even though they had no good answer to why they live in a world in which CAS is instantiated.

In considering this matter, it is important not to confuse questions about CAS with questions about brain event complexity of a particular kind. If we focus on that property alone, it will always seem contingent, and unexplained, why it should yield phenomenality. In contrast, the conservation principle supports the conception of CAS as a *common* nature that is manifested both in certain kinds of brain events and in qualitative events, and this commonality is crucial to its explanatory force.

An analogy may help here. We can imagine people having formed a concept that they call ‘energy’, but that we would think is more appropriately called ‘kinetic energy’. They do not find that their ‘energy’ is conserved. To get a better physics, they have to form another concept — a concept of a property that is possessed by boulders rolling downhill, but is also possessed by boulders that are still resting near the top of a mountain. One can imagine a period in which these concepts are confused, but if they do improve their physics, they will have to clearly distinguish the new concept (energy properly so called) from their old, restricted one.

Let us end this section by noting a source of intellectual satisfaction that might arise if our sciences develop in a way that now seems epistemically possible. We may begin by recalling that we have noted more than one kind of persistence. All qualitative events are able to persist with little or no qualitative change through time, but some others, e.g. visual qualitative events, also ‘persist through space’ — i.e. colours have extensity, or are spread out over apparent surfaces

(which may be flat or curved). We may say the same thing by using ‘continuity’: colour expanses have spatial continuity, and all qualitative events have temporal continuity. Now, it is plausible that these continuities are related. Since we do not at present know the brain event properties whose instances yield qualitative events, we are in no position to know how these continuities might be reflected in our conception of CAS. But it does not seem beyond possibility that at some point our descendants may be able to find a structure within CAS, according to which it has several ‘subdimensions’ that can correspond to different species of continuity. If they were able to see such a structure, that would likely increase their sense of intellectual satisfaction.

6. Smart’s Obstacle, and the Particular Form of the HP

In a famous paper, J.J.C. Smart wrote that

I cannot believe that ultimate laws of nature could relate simple constituents to configurations consisting of perhaps billions of neurons (and goodness knows how many billions of billions of ultimate particles) all put together for all the world as though their main purpose in life was to be a negative feedback mechanism of a complicated sort. (Smart, 1959, p. 143)

Smart goes on to say that such laws have a queer ‘smell’ to them; and, when the matter is put as Smart puts it, it is hard not to be sympathetic to his characterization. The bad odour, however, plausibly arises from Smart’s rhetoric. On the CAS theory, a fundamental law of nature is (ConCAS), and this is not a relation among billions of anything: it is a relation between a certain kind of complexity and a certain kind of simplicity. What’s true is that the instantiations of the right kind of complexity require perhaps billions of neurons. But no one finds the law of gravitation malodorous on the ground that it relates bodies that may consist of billions of atoms.

One may reply to this point by noting that gravitation also holds between single particles, and that gravitational interactions of large bodies can, in principle, be constructed from the resultants of forces between pairs of particles. But no such procedure can be applied to CAS. This correct observation, however, need not remove our descendants’ satisfaction with CAS theory, because they will note that neural (or field) complexity of the right kind *necessarily* involves a large number of neurons in various states of activation (or regions with patterns of many different field strengths). To deny the relevance of this observation is to commit oneself to a principle that no law of nature

can hold with respect to a property whose instantiation requires a plurality of particles. When stated in this way, such a commitment seems far from evident. Our descendants may be able to reject it as incompatible with the best theory of the whole of what they take to be true, including the regular relation between brain event complexity of a certain kind and qualitative events that have a certain sort of relative simplicity.

These remarks bring us to the end of our discussion of the HP in its general form, i.e. when it is phrased as ‘Why, given that there are our neural events and their causal relations to each other and to the world at large, should there be any phenomenally qualified events at all?’ The answer, in brief, is that CAS is conserved, and so, when brain events having the right kind of complexity occur, there have to be events with the corresponding kind of relative simplicity, and these are the qualitative events.

But Smart writes of ‘laws’ in the plural, which suggests that he may also have been thinking of a further problem that we have not yet considered. This is the HP in its particular form, namely, ‘Why, given that there is a particular kind of neural event (standing in its particular web of causal relations) should there be this particular kind of qualitative event that it yields?’ Failure to have any idea how to answer this question might very well dim any sense of intellectual satisfaction that our descendants might derive from filling out the sketch suggested so far. So, let us turn to a possible way of responding to the HP in its particular form.

7. Parallel Structure in Quality Spaces

Austen Clark (2000) has offered a suggestion for overcoming the explanatory gap which, unsurprisingly, can be put to use in addressing the HP. The early parts of what follows are devoted to explaining the core of Clark’s idea. As we proceed, we will see how far this approach can take us.

A key claim is that there is structure in our quality spaces. This structure is built on similarity relations, and since all sensory modalities give rise to judgments of similarity, they all have structured quality spaces. We shall, however, follow Clark in focusing on vision, which provides the easiest examples.

Consider the claim that:

- (O) Orange is similar to red and also to yellow, and it is far more similar to each of these than it is to turquoise.

(O) is one of countless claims of the same form, the truth of which is evident to those with normal human visual systems. Taken together, they give a set of relations among colour qualities, and this set of relations is the structure of colour quality space.

Clark explains in detail why the relata in quality spaces cannot be stimuli or discriminative dispositions.¹⁹ The relata in our quality spaces are held to be the qualities themselves.²⁰ Moreover, and crucially, these relations among qualities seem to be necessary. To deny this would be to say, for example, that in some possible world, orange might be more similar to blue than it is to yellow; and this seems impossible.²¹ We may say (though Clark does not so formulate the point) that *similarity essentialism* holds for qualities: their similarity relations are the same in all possible worlds.

A second key claim is that the structure of causes of a set of events cannot be simpler than the structure present in the set of effects. It could be more complex, although parsimony counsels against supposing unneeded complexity. But it cannot be simpler; it cannot be that the same cause produces sometimes one effect and sometimes another.

The third claim we need is empirical, and well established. Colour space is not regular and symmetrical. For example, fully saturated yellow is more similar to white than fully saturated red, blue, etc. In some regions small differences in wavelength are easily discriminated, in others not.²²

If we put these three claims together, we can easily imagine that future science may identify a set of neural events that stand in a set of similarity relations that is isomorphic to the structure of similarity relations in colour quality space (and analogously for other sets of neural events and other quality spaces). Once we had such a set, it might prove irresistible to suppose that *of course* it is the item that has

[19] Briefly, (i) metamers are distinct stimuli that give rise to the same qualities, and (ii) similarity judgments differ across individuals, even though the structure of their quality spaces is the same.

[20] Clark makes a distinction between phenomenal properties and qualitative properties, but questions can be raised about the clarity of this distinction, or whether it can be made in quite the way he suggests. Fortunately, the plausibility of the views discussed in this section does not depend on this distinction, or on Clark's particular way of understanding 'qualities'. Still, readers should be aware that the appropriation of Clark's ideas used here abstracts from some of Clark's own commitments.

[21] Of course, in some possible worlds our *words* might be permuted so that (O) would express a falsehood in that world. But where (O) means in W what it means here, (O) is true.

[22] Clark (2000, p. 25) notes five kinds of asymmetries among colours.

a certain place in the neural structure that causes the item that has the corresponding place in quality space.

Clark puts the point in terms of identity.

To close the explanatory gap, the world must be such that there is one best way that the spacing of qualities could be implemented in the nervous systems with which we are endowed. But here the asymmetries in the human colour quality space... come to our rescue. Once we have discerned the lumpy, asymmetric, anisotropic structure of that space, and we have detailed the home-grown capacities of our sensory nervous systems, we will see that there is just one best way to fit the two together.

...So I claim that we can conceive of a conceptual scheme and of empirical details under which it would be inconceivable that sensing redly be anything other than brain process squiggle-squiggle. (Clark, 2000, pp. 37–8)

Structural isomorphism in general does not imply identity — sets of pits in a CD, for example, are not the same thing as the compression waves that emanate from stereo speakers. So it is evident that the identity claim does not follow from the imagined isomorphism between neural event types and sensory qualities. But the background provided by Clark does plausibly show us how we might very well come to feel that it *has to be* that yellow in a qualitative event is brought about by one kind of brain event, red by a certain other kind of brain event, and so on.

The question now is whether this feeling would be sufficient to provide our descendants with intellectual satisfaction concerning the explanatory gap and the HP. One kind of doubt that it would be sufficient is that it does not seem to be metaphysically necessary that structural isomorphism holds. If it is not necessary, then our descendants would be exposed to the question of why structural isomorphism holds in our world.

A possible suggestion in response is that our descendants might be persuaded of causal essentialism. Perhaps causal essentialism *plus* similarity essentialism can succeed, where causal essentialism alone could not. This suggestion, however, seems inadequate, because it does not address the problem that was raised for causal essentialism. To remind, that was that there does not seem to be anything in particular phenomenal qualities that would lead us to say that we cannot have identified instances of them in other possible worlds unless their causal relations were the same as they are in this world. This point holds even if we imagine our descendants as having accepted (ConCAS). They would think that neural events require the

occurrence of qualitative events of relative simplicity, but that would not by itself carry them to the view that each particular kind of instance of the right kind of brain event complexity must be the cause of the same qualitative event in all worlds as it causes in this world. It seems, then, that we must look for something else to add to (ConCAS) to deal with the HP in its particular form.

A second source of doubt whether the background provided by Clark can help with the particular form of the HP is the possibility that there is more than one set of qualities whose structure is isomorphic to the structure of the set of brain events that have the right kind of complexity. If there are such multiple sets, our descendants would be faced with the question of why our brain events cause elements of one of these sets rather than elements of the other.

8. The Minimum Arbitrariness Principle

In this section, we will see that we can add something to the background provided by Clark that will give us a more robust sense that our brain events have to yield what they do.

Let us begin with a response to the second source of doubt. To put the idea in a way that is natural for a dualist, perhaps the reason why our brain events do not cause other properties with the same similarity structure is that there simply *are no* other property sets that have that structure. In that case, it would be true that *if* the events in our set of brain events (of the right kind) cause elements of a set of qualitative events with a parallel structure, then it is just *these* kinds of qualitative events that they must cause (where *these* kinds are the kinds that we actually have, and that we empirically discover to be caused by brain events of the kinds that our future science will discover).

This thought suggests a converse. Perhaps there are no sets of properties, other than those found in our brain events, that have a similarity structure that parallels that of our sensory quality spaces. In that case, the occurrence of our qualitative events would require the occurrence of brain events of the kinds we have.

Turning to the first source of doubt, we can note that the structural isomorphism of causes and effects plausibly follows from what we may call the Minimum Arbitrariness Principle.

- (MAP) (a) So far as is possible, proximate effects of closely similar causes are closely similar; and (b) So far as is possible, proximate causes of closely similar effects are closely similar.

The ‘so far as is possible’ clauses in (MAP) require some explanation. To this end, imagine that we have two sets of properties (e.g. brain event properties and phenomenal qualities) with several kinds of similarity relations. To illustrate what this means, consider that colours have at least three dimensions of possible similarity — namely in hue, saturation, and brightness. Likewise, neural events may resemble and differ in phase, in rates of firings of some of the neurons involved in the event, in patterns of intervals between bursts of activation and relative quiescence, and perhaps in other ways. Now, we cannot rule out that in order to obtain the maximal degree of preservation of close similarities in effects given similarities in their proximate causes, there will have to be *some* cases in which causes that are closely similar in one dimension produce effects that are considerably different in the corresponding effect dimension (and/or conversely). (MAP) does not rule out the possibility of such cases. Instead it implies that there are as few such cases as possible.

A further source of unavailability of possibilities is that it may be that there are no two sets of properties that are perfectly isomorphic in structure. In that case, (MAP) implies that the actual causal relations will hold between those sets of properties for which there is a best fit.

Both this latter point and the response suggested for the second source of doubt implicitly commit us to a strong claim, which we should explicitly recognize. To wit,

(uMAP) There is a unique solution to the constraints that (MAP) imposes.

To further explain the view that is based on (MAP) and (uMAP) let us begin by explaining the reason for using ‘minimum arbitrariness’ in naming these principles. The idea is that if proximate effects of closely similar causes were *not* as closely similar as possible, then there would be many possible candidates for what they might be. Without (MAP), it would seem arbitrary which one, among many dissimilar possible effects, would be the effect of any particular cause. It is the minimization of arbitrariness that constrains possible causal relations.

To introduce a possibly helpful analogy, we may imagine a Martian with no understanding of music who ponders the fact that small discs with pits are often found in proximity to vinyl discs with grooves. We may imagine that after lengthy investigation, our Martian is able to find a classification of properties of variations in grooves and a classification of patterns of pits that enables discovery of a close correspondence, in which sequences of similar groove properties can be

matched to sequences of similar pit patterns. We may further imagine that the samples with which our Martian works include records that have undergone various degrees of wear. In that case, there may be no perfect correspondence of groove variation properties and pit patterns, but it is still likely that there would be a best fit.

We shall shortly turn to some caveats and difficulties in accepting (MAP), and explain why (MAP) and (uMAP) give only a sketch of a solution to the HP and not an actual solution. But let us continue for a moment with the explanation of the view that results if one does accept these principles. It is important to note that these principles concern sets of properties. Now, it is independently plausible that the set of properties is the same at all possible worlds. Let us express this as the Necessity of Properties.

- (NP) The set of properties at any possible world is the same as the set of properties at any other possible world.

Of course, not all properties are instantiated in all possible worlds. For a property to be in the set of properties at a world, understood as intended here, is for it to be true at that world that in some world or other that property is instantiated. So, what (NP) says is that if it is true in some world that a property can be instantiated (i.e. that it is instantiated in some possible world), it is true in all worlds that that property can be instantiated (i.e. it is true in all worlds that there is some world in which that property is instantiated). It follows from that, that if a property is instantiated in any world, then it is true for every possible world, that that property is in the set of properties at that world.

In view of (NP), and the essentialism of similarity relations, applying (MAP) will give the same result no matter which world we suppose we are in. So, the constraint on causal relations that (MAP) provides is the same in all worlds in which (MAP) holds, and if (uMAP) is true, then the causal relations are the same in all worlds in which (MAP) holds. This result provides a kind of generality that our descendants may find satisfying.

The upshot of this discussion is that (MAP) provides something like causal essentialism, and therefore something like the sense of necessity provided by the latter view. The grounds, however, are different: they rest on minimization of arbitrariness and similarity essentialism. And similarity essentialism applies to phenomenal qualities, so the objection that causal essentialism seems least plausible for phenomenal qualities does not apply here.

9. Limitations

Part of the explanation of (MAP) relied on the idea of maximal overall preservation of similarity in effects given similarity of proximate causes. There is no guarantee that such similarity preservation can have a unique maximum, even if we accept that degrees of similarity in a particular respect can always be well defined. For it could happen that departure from similarity of effects of causes that are similar in respect A would allow close similarity preservation in respects B and C, while departure from similarity of effects of causes that are similar in respect B would allow close similarity preservation in respects A and C.

However, the possibility of non-arbitrary solutions to such ‘trade off’ problems remains open, and increasing knowledge of the relevant kinds of neural complexity might suggest principles for their resolution. (uMAP) has been identified as a very substantive assumption; part of what it commits us to is that there are non-arbitrary solutions to problems of the kind just described.

(MAP) itself is not a necessary truth: there could be (as far as we have any reason to believe) worlds in which arbitrariness reigns. So, even if our descendants think that (MAP) holds in our world and that (uMAP) is true, they will not regard the causal relations in our world as necessities. They will not be able to say without qualification that our neural events *have to* cause our qualitative events. The generality to which they will be entitled is this one: in any world in which (MAP) holds, our brain events would cause the same qualitative events that they do in our world.

We have no *a priori* proof that we live in a world in which (MAP) holds. However, if our sciences turn out to support massive parallelism in the structures of brain event causes and qualitative events, that would be some reason to believe that we inhabit a (MAP) world.

Our descendants are unlikely ever to have an *a priori* proof that we live in a (MAP) world, because it is doubtful that we can ever be quite sure that there are *not* certain properties. To explain what this means, consider that there do not seem to be any ‘tastolours’ (or ‘colourastes’) i.e. properties such that their instances would be as much like an instance of some colour as they are like an instance of some taste.

It is tempting to say that there is no such property. But we cannot be certain. It might be that there are possible brain event properties that would give rise to tastolours and that we cannot have such events because our brains are simply not constructed so that those kinds of events can occur in them. We cannot imagine passing by small

differences from, say, red to oreganoish, as we can imagine passing by small differences from red to green (e.g. through oranges, yellows, and chartreuses). But again, it might be that imagining such a series of qualities requires having neural equipment that is arranged in a way that the construction of our brains does not permit.

The possibility just described gains some support from the likelihood that there are animals on earth that have phenomenal qualities that we cannot have. Some birds are tetrachromats, and we are likely permanently unable to imagine what their ‘extra’ colours are like. Our neural systems cannot get into the states that those of electric eels can get into. In these cases, we are in no position to know about similarity or dissimilarity between our phenomenal qualities and these other animals’ phenomenal qualities — except, of course, by the *a posteriori* route which *assumes* (MAP), and reasons that to the extent that there may be neural similarities or differences, these are reflected in similarities and differences among phenomenal qualities.

If we assume that our descendants will never be able to overcome this kind of uncertainty, we should conclude that they will never be in a position to survey all the phenomenal qualities, all the possible brain event properties, and calculate the consequences of (MAP) so as to give an *a priori* derivation of phenomenal qualities from our brain event properties.

In consequence, the epistemology of causes of qualitative events will always appear to us just as Hume described it — a thoroughly *a posteriori* affair dealing with events of kinds that apparently could be differently related. And, in view of the non-necessity of (MAP), that appearance is correct. But our descendants could accept all this, and still find intellectual satisfaction in the view that *if* we live in a (MAP) world, as we apparently do, then, if (uMAP) is true, then there is a reason why just *these* qualitative events ensue upon just *these* brain events.

10. Conclusion

We cannot predict the future of science, and so we cannot know that the empirical discoveries envisioned here will actually be made, or that the conceptual changes outlined will come to pass. It is, however, an epistemic possibility for us that these things will happen. If they do, then it is conceivable that our descendants may also find that (ConCAS) — together with greater specification of the ‘certain kind’ of brain event complexity — gives intellectual satisfaction with regard to the general form of the HP. And it is conceivable that they

will also find that (MAP) and (uMAP) give intellectual satisfaction with regard to the HP in its particular form.

These possibilities cannot actually provide the intellectual satisfactions that they envisage, so they are not a solution to the HP. But they do outline a conceivable path to such satisfaction; and that is to say that they do solve the HPP. We need not be discouraged by those who dismiss the HP on the ground that a problem cannot be genuine unless we have some idea of how it might be solved — because we do have such an idea.

One may worry that possession of the foregoing sketch of a solution to the HP shows that O'Hara and Scutt (1996) were half right: there is no HP — not because there cannot be a problem where we cannot envisage a solution, but because ability to sketch a possible solution shows that the alleged hard problem is not actually *hard*. In response, it should first be noted that we need not accept that hardness requires eternal or in-principle unsolvability. An alternative characterization would be unsolvability within our present conceptual scheme.

Would the developments outlined in this paper amount to a change in our conceptual scheme? An affirmative answer seems supported by the following observations. The concept of CAS, and *a fortiori* appeals to (ConCAS), are not familiar in the literature on consciousness. To the extent that intellectual satisfaction in the face of the HP may depend on use of this concept, it should thus be no surprise that the HP strikes us as hard. Further, this concept and its associated principle cannot be fully understood until we experimentally establish a detailed understanding of what kinds of events yield qualitative events. Here, we may compare particle physics, which contains many concepts that arose only from the need to explain experimental results, and would likely never have been thought of without them. Further, while the essential idea behind (MAP) is present in Clark (2000), it is bound up there with identity theory, and is in any case not very frequently mentioned in discussions of consciousness.

These considerations do not conclusively establish that the development of dualism outlined here would amount to a significant change of conceptual scheme; for the criteria for what should count as such a change are not sufficiently clear. They do, however, make an intelligible case for the following view: if our descendants developed dualism along the lines described here, they would find it quite intelligible that their predecessors who lacked the concept of CAS and a developed form of (MAP) would have written of a hard problem of consciousness.

This paper has not claimed that the series of developments it describes is the only way that a dualist ascendancy could play out. It should, however, reduce fear of Poverty, because it exhibits a possible route to a future, articulated, intellectually satisfying account.

Acknowledgments

I would like to thank Stephen Biggs, Fred Kroon, Bill Lycan, Denis Robinson, and Jonathan Tsou for their very helpful comments on drafts of this paper.

References

- Bayne, T. & Montague, M. (eds.) (2011) *Cognitive Phenomenology*, Oxford: Oxford University Press.
- Bird, A. (2005) Laws and essences, *Ratio*, **18** (4), pp. 437–461.
- Carruthers, G. & Schier, E. (2012) Dissolving the hard problem of consciousness, *Consciousness Online 4*, [Online], <http://consciousnessonline.files.wordpress.com/2012/01/dissolvinghardproblem.pdf>.
- Chalmers, D.J. (1995) Facing up to the problem of consciousness, *Journal of Consciousness Studies*, **2** (3), pp. 200–219.
- Chalmers, D.J. (1996) *The Conscious Mind: In Search of a Fundamental Theory*, Oxford: Oxford University Press.
- Churchland, P.M. (1984) *Matter and Consciousness*, Cambridge, MA: MIT Press/Bradford Books.
- Churchland, P.M. (2012) *Plato's Camera: How the Physical Brain Captures a Landscape of Abstract Universals*, Cambridge, MA: MIT Press.
- Clark, A. (2000) *A Theory of Sentience*, Oxford: Oxford University Press.
- Flanagan, O. (1992) *Consciousness Reconsidered*, Cambridge, MA: MIT Press/Bradford Books.
- Freeman, W.J. (2008) A pseudo-equilibrium thermodynamic model of information processing in nonlinear brain dynamics, *Neural Networks*, **21** (2–3), pp. 257–265.
- Goff, P. (2011) A posteriori physicalists get our phenomenal concepts wrong, *Australasian Journal of Philosophy*, **89** (2), pp. 191–209.
- Mangan, B. (2001) Sensation's ghost: The non-sensory 'fringe' of consciousness, *Psyche*, **7** (18), [Online], www.theassc.org/files/assc/2509.pdf.
- McFadden, J. (2013) The CEMI field theory: Closing the loop, *Journal of Consciousness Studies*, **20** (1–2), pp. 153–168.
- Nagasawa, Y. (2008) *God and Phenomenal Consciousness*, Cambridge: Cambridge University Press.
- Nida-Rümelin, M. (2007) Grasping phenomenal properties, in Alter, T. & Walter, S. (eds.) *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*, New York: Oxford University Press.
- O'Hara, K. & Scutt, T. (1996) There is no hard problem of consciousness, *Journal of Consciousness Studies*, **3** (4), pp. 290–302.
- Papineau, D. (2002) *Thinking about Consciousness*, Oxford: Oxford University Press.
- Pockett, S. (2002) Difficulties with the electromagnetic field theory of consciousness, *Journal of Consciousness Studies*, **9** (4), pp. 51–56.
- Price, M.C. (1996) Should we expect to feel as if we understand consciousness?, *Journal of Consciousness Studies*, **3** (4), pp. 303–312.

- Robinson, W.S. (2004) *Understanding Phenomenal Consciousness*, Cambridge: Cambridge University Press.
- Rosch, E. (1994) Is causality circular? Event structure in folk psychology, cognitive science and Buddhist logic, *Journal of Consciousness Studies*, **1** (1), pp. 50–65.
- Schwitzgebel, E. (2013) If materialism is true, the United States is probably conscious, *Consciousness Online Conference 5*, [Online], consciousnessonline.com/2013/02/15/if-materialism-is-true-the-United-States-is-probably-conscious/.
- Seager, W. & Allen-Hermanson, S. (2012) Panpsychism, in Zalta, E.N. (ed.) *Stanford Encyclopedia of Philosophy*, [Online], plato.stanford.edu/.
- Shoemaker, S. (1980) Causality and properties, in van Inwagen, P. (ed.) *Time and Cause*, pp. 109–135, Dordrecht: D. Reidel.
- Shoemaker, S. (1998) Causation and metaphysical necessity, *Pacific Philosophical Quarterly*, **79**, pp. 59–77.
- Shoemaker, S. (2007) *Physical Realization*, Oxford: Oxford University Press.
- Skarda, C.A. & Freeman, W.J. (1987) How brains make chaos in order to make sense of the world, *Behavioral and Brain Sciences*, **10**, pp. 161–195.
- Smart, J.J.C. (1959) Sensations and brain processes, *Philosophical Review*, **68**, pp. 141–156.
- Stoljar, D. (2006) *Ignorance and Imagination*, Oxford: Oxford University Press.
- Strawson, G. (2009) Realistic monism: Why physicalism entails panpsychism, in Skrbina, D. (ed.) *Mind That Abides: Panpsychism in the New Millennium*, pp. 33–57, Amsterdam: John Benjamins.
- Swoyer, C. (1982) The nature of natural laws, *Australasian Journal of Philosophy*, **60** (3), pp. 203–223.
- Tye, M. (2000) *Consciousness, Color, and Content*, Cambridge, MA: MIT Press/Bradford Books.
- Tye, M. (2009) *Consciousness Revisited*, Cambridge, MA: MIT Press/Bradford Books.
- White, S. (2007) Property dualism, phenomenal concepts, and the semantic premise, in Alter, T. & Walter, S. (eds.) *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*, pp. 210–248, New York: Oxford University Press.
- Wilson, J. (2010) What is Hume's dictum, and why believe it?, *Philosophy and Phenomenological Research*, **80** (3), pp. 595–637.

Paper received March 2013; revised June 2013.