Qualia and Introspection

Abstract

The claim that behaviourally undetectable inverted spectra are possible has been endorsed by many physicalists. I explain why this starting point rules out standard forms of scientific explanation for qualia. The modern 'phenomenal concept strategy' is an updated way of defending problematic intuitions like these, but I show that it cannot help to recover standard scientific explanation. I argue that Chalmers is right: we should accept the falsity of physicalism if we accept this problematic starting point. I further argue that accepting this starting point amounts to at least implicitly endorsing certain theoretical claims about the nature of introspection. I therefore suggest that we allow ourselves to be guided, in our quest to understand qualia, by whatever independently plausible theories of introspection we have. I propose that we adopt a more moderate definition of qualia, as those introspectible properties which cannot be fully specified simply by specifying the non-controversially introspectible 'propositional attitude' mental states (including seeing x, experiencing x, and so on, where x is a specification of a potentially public state of affairs). Qualia thus defined may well fit plausible, naturalisable accounts of introspection. If so, such accounts have the potential to explain, rather than explain away, the problematic intuitions discussed earlier; an approach that should allow integration of our understanding of qualia with the rest of science.

1. Introduction

We are concerned with the definition of consciousness; that is, we are discussing the nature of the *target* of explanation in our scientific or philosophical study of consciousness. As Vimal (this issue) clarifies, there are many views about how to pin the target down. Ought we to be trying to explain consciousness conceived of as a cognitive property? As a phenomenal property? As somehow related to awareness and attention?

The present paper is concerned with the phenomenal aspect of consciousness: with *qualia*; with the 'something it is like' to have an experience[1]. This is not to completely ignore the many other aspects present within the broader concept of 'consciousness', as covered by Vimal. Indeed, it is my hope that many or most of these aspects will prove to be intimately related to each other, within the right theoretical framework. Nevertheless, there is a certain mystery to the phenomenal aspect of consciousness in particular. It seems especially hard to find a place for that aspect within our growing understanding of the natural world (Levine, 1983; Chalmers, 1995).

The aim here will be to critique a particular approach to phenomenal consciousness which 'defines in', from the start, certain problematic features of qualia. Specifically, I will critique that class of approaches which entail that our knowledge of phenomenal facts is *a posteriori* with respect to our knowledge the physical facts.

There is quite a lot to be unpacked here, about what philosophers mean when they talk like this. To get the discussion started, I need to introduce two assumptions which I share with the position I am critiquing. The first is this: when I introspect and come

[1] Qualia are the characteristic properties of phenomenal consciousness: something is a state of phenomenal consciousness if and only if it has such properties.

to think that it is 'like this' for me to see red (for example), then my thought refers to some fact: a fact about 'what it is like' (or, equivalently, about what the phenomenal feel is). We can call such facts *phenomenal facts*, and knowledge of such facts *phenomenal knowledge*. The second shared starting point is this: it is possible to discover the existence of regular co-occurrence between physical facts and introspectible phenomenal facts. If so, we would be able to discover that when certain physical facts about a creature ('neural correlates of consciousness') or, perhaps better, about a creature in its world (physical correlates of extended mind), are true, then certain phenomenal facts are always true.

Given these shared starting points, the *a posteriori* approach which I am critiquing goes on to claim that the existence of this regular co-occurrence between public physical facts and introspectible phenomenal facts could not have been worked out in advance, purely by conceptual analysis, however well we understand what we mean, when we say that we 'know what it is like' and however well we understand the public physical facts which co-occur with the phenomenal facts.

David Chalmers has called this kind of approach *phenomenal realism* (Chalmers, 2003). As Chalmers rightly states (2003), and as I will show below, certain very common presuppositions about phenomenal facts (specifically, either or both of the inverted spectrum[2] or zombie[3] claims about qualia) directly entail that there is this kind of *a posteriori* relation between physical and phenomenal facts. Chalmers also states that it is not possible to "take consciousness seriously" (Chalmers, 1996, p.xii), without adopting starting points which lead directly to such a view. For the purposes of the present paper, I will use the label *strong phenomenal realism* for such views, since my main aim will be to claim that there are other ways to take qualia seriously.

The biggest problem with such *a posteriori* approaches is that they rule out (on the basis of presuppositions built into their definition of qualia) a certain extremely standard form of scientific explanation. In section 2, I will outline the model of explanation in question. Then, in section 3, I will present one historically popular (and still influential) approach to naturalising qualia which I will use as an example, to make clear why these starting points rule out this type of explanation. In section 4, I will outline the modern phenomenal concept strategy, which claims that physicalism can be preserved, even if we adopt such *a posteriori* claims about qualia. I will present reasons to agree with Chalmers that this cannot work.

The final parts of this paper question whether theorists really are entitled to such problematic starting assumptions. In section 5, I will argue that such starting points amount to implicit *theoretical* claims about the nature of introspection: claims which, if true, are themselves justified by introspection. I will point out that there is widespread disagreement about the nature of introspection, and I will suggest that there is a widespread tendency to build presuppositions about it into our theories of sensory experience. As such, I will argue that the theorists I am critiquing are not justified in endorsing such problematic starting points.

Finally, in section 6, I argue that it is possible to preserve a moderate form of phenomenal realism (there really are qualia, we really do know them in introspection), without these problematic starting points. To do this, I propose a more moderate definition of qualia, which allows our theorising about them to be guided by whatever

[2] The claim that there can be creatures which are physically (or functionally) identical to each other, but which have different phenomenal mental lives.

[3] The claim that there can be creatures which are physically (or functionally) just like us, but with no phenomenal mental lives at all.

independently plausible theory of introspection we have. I will argue that this moderate definition still looks to have the ability to explain, rather than completely explain away, many intuitions about qualia, including some of the problematic starting points above.

2. Normal Scientific Explanation

In this section I will briefly present an account of a very standard form of scientific explanation. My claim is that this form of explanation is so ubiquitous, that for any property which science recognises, the existence of that property is either a) believed to be explicable in terms of more fundamental properties in this way, or b) is treated as a fundamental fact about our universe.

A paradigm example is the explanation of the properties of water (the way it freezes and boils, its transparency, its viscosity, and so on) in terms of the properties of, and interactions between, H_2O molecules (the shape of the molecule, the forming and breaking of hydrogen bonds between molecules, and so on).

Philosophers often like to emphasize the fact that the relation between water and H_2O molecules can only be known *a posteriori*: that the existence of such a relation could not have been worked out in advance of the relevant empirical discovery, even with the most careful reasoning. But this is a misdescription, or at least an over-simplification. As Loar (1997, p.608) and Chalmers (2006), amongst others, have noted, there is an *a priori* entailment between the low level properties of H_2O and the high-level properties of water.

It is possible to be too prescriptive about exactly what such an *a priori* entailment involves (see note 4), so I will try to put it as neutrally as possible: having mastered the concepts involved in describing the low and high levels, it would not be rational to believe that certain high level facts do *not* apply (e.g. that there is stuff which behaves like water round here) when certain low level facts apply (that there is a large number of H_2O molecules with a certain energy distribution, etc., around here). This is an *a priori* conceptual entailment, in that the existence of the rational link in question follows purely from an understanding the concepts involved, with no further empirical research necessary[4].

Note, also, that it is a *one way* conceptual entailment: the facts[5] about H_2O molecules entail that a mass of them behaves the way water behaves, but the facts about the way water behaves do not entail that it is made of a mass of H_2O molecules. I would agree that it is not rational for someone informed by modern science to claim that water is *not* (mainly) made of H_2O. But the logic in this direction is fundamentally *a posteriori*, based on induction from the *discovery* that what has been

[4] In fact, this is not an *a priori* entailment in the strict philosophical sense: a step requiring *no* empirical knowledge whatsoever. This is because the kind of practical mastery of the concepts required to see the connection between the high and low levels *does* require empirical knowledge and experience. The account I'm giving therefore claims that we use common sense, at the point where the more traditional 'deductive nomological' account of scientific explanation/reduction would claim that we use 'bridge laws'; but I don't think anything in the main line of argument hinges on this difference from the perhaps more familiar account. For these and various other reasons, the account I am giving is not quite that of Chalmers and Jackson (2001).

[5] A note on how I individuate facts in this paper: I treat the fact that 'H_2O molecules are present' as a different fact from the fact that 'water is present' (even when they refer to one and the same state of affairs), because of the (one-way) conceptual independence between the levels of description involved; conversely, I would treat the fact that 'a bachelor is present' and the fact that 'an unmarried male is present' (when they refer to the same state of affairs) as the same fact, because there is no conceptual independence between the two descriptions involved.

found to explain wateriness round here always has been H_20. This relationship between concepts, which is *a priori* in one direction but *a posteriori* in the other, can be contrasted with *two way* cases such as the relationship between 'bachelor' and 'unmarried male' (*a priori* in both directions), or that between 'son of Barack Obama, Senior' and '44th president of the United States of America' (*a posteriori* in both directions). I am claiming that the special, one way kind of relationship is essential for scientific explanation[6]. It is important to be clear that the high level properties do not somehow disappear once we have such an explanation: it is only in talking at the high level that we can express what needed to be explained in the first place. In fact, the concepts of the high level need not even be applicable at the low level.

This pattern is not specific to water and H_2O; it is widely repeated, in scientific explanation. The same pattern holds between the micro-facts of modern genetic theory (transmission of DNA, gene-expression during embryonic development, etc.) and the macro-facts of inheritance with variation required for Darwinian evolution[7], or between the micro-facts of statistical mechanics and the macro-facts of thermodynamics, and so on and so on.

Unfortunately, many views which take qualia seriously, including many which see themselves as varieties of physicalism, build elements into their definition of qualia which rule out any chance of providing explanations of this type.

3. The Nature of Functionalism

There is a historically popular brand of functionalism which tries to argue that inverted spectra are perfectly possible, and are compatible with normal science. The view was advocated (with subtle differences, on which see more below) by Lewis (1980), the Churchlands (1982) and Shoemaker (1975), amongst others. Lewis says:

> "As philosophers, we would like to characterize pain a priori. … As materialists, we want to characterize pain as a physical phenomenon." (Lewis, 1980, p.123)

An *a priori* characterisation of pain would be one which makes clear that certain facts (e.g. wincing, groaning, withdrawing from noxious stimuli[8], etc.) are two way conceptually identical to facts about pain. Such an *a priori* characterisation of pain would presumably be just a small part of an *a priori* characterisation of the entire mental level (including belief, desire, perception and so on) applicable to any agent with a mental life.

It is a general characteristic of functionalism (not just of the particular variant being discussed here) that it supposes that there exists some level of characterisation of a creature which is 'the mental level', and that there are other facts about that creature which can vary, independently of the mental level. This seems to me to be the right kind of approach (with caveats about exactly how this approach should be understood, which I will explain below). In the case of the type of functionalism I am discussing here, however, this strategy is *not* followed through to what might seem its logical conclusion. For the *a priori* characterisation of the mental level is supposed, by these authors, *not* to capture everything mental which there is to say about the subject. Specifically, it does not capture what it is like to be the subject; it is supposed that

[6] See section 3.1 for a brief discussion of an opposing view.

[7] As in many such cases, we have enough of the detail so that the relation between the levels no longer seems 'in principle' mysterious – even though many details remain to be discovered, and our understanding of both levels may no doubt be refined in the process.

[8] Or, at least, a tendency towards such behaviours, which may be masked by other factors but which could be revealed by suitable experimentation.

there could be two subjects who are the same, in terms of this publicly observable mental level of behaviour, but where it nevertheless feels one way to be one subject, and another way to be the other.

In making this point, the Churchlands mention the classic inverted spectrum case, in which we are asked:

> "to imagine someone … [who has] a sensation of red in all and only those circumstances where you have a sensation of green, and so forth." (Churchland and Churchland, 1982, p.122)

The Churchlands explicitly claim that:

> "These cases are indeed imaginable, and the connection between quale and functional syndrome is indeed a contingent one." (Churchland and Churchland, 1982, p.122)

In a similar vein, Lewis asks us to:

> "Suppose that the state that plays the role of pain for us plays instead the role of thirst for a small subpopulation of mankind, and vice versa." (Lewis, 1980, p.128)

Lewis argues that in such a case:

> "there is no determinate fact of the matter about whether the victim of the interchange undergoes pain or thirst." (Lewis, 1980, p.128)

This claim would be false if the phenomenal feel were fully determined by the functional role: if so, a groaning, writhing[9] agent would be unequivocally in pain, whatever was the case about the physical states constituting the agent. But the authors quoted here think that there are two meanings to pain, the *a priori* meaning, where pain simply refers to that state where a creature displays (or tends to display) pain behaviour, and the *a posteriori* meaning, which refers to whatever physical state science has determined to fill this functional role (in a population) [10].

I will describe such views as *hybrid functionalism*, since they combine elements of the earlier identity theory ('the physical stuff determines the feel', c.f. Lewis, 1980, p.124) with what would otherwise be 'pure' functionalism (the claim that the mental facts are fully captured at the in principle publicly observable mental level).

Why, though, believe that a difference in "physical realization" has any "bearing on" the introspectible facts about "how that state feels"? (The quotes are from Lewis, 1980, p.130.) The Churchlands flesh out this part of the view in more detail:

> "the spiking frequency of the impulses in a certain neural pathway need not prompt the non-inferential belief, "My pain has a searing quality." But withal, the property you opaquely distinguish as "searingness" may be precisely the property of having 60 Hz as a spiking frequency." (Churchland and Churchland, 1982, p.128)

The claim is that the physical state of 60 Hz neural firing (or whatever physical state it really turns out to be) *is* what we introspect, when we introspect a searing pain. Equally, in some other agent, the same functional role might be filled by a different physical state, such as inflation in hydraulic cavities in the feet (Lewis' semi-humorous suggestion as to the state which might play the role of pain in Martians). A difference like this is supposed to be the right kind of difference to account for a

[9] On any plausible *a priori* account of the mental, it must be supposed that the groaning and writhing is suitably integrated with other aspects of the agent's behaviour, quite possibly including their rationality. This is a point which both the Churchlands (1982, p.128) and Shoemaker (1990, p.71) make.

[10] There are issues here, to do with whether, and in what sense, sub-system states could possibly be role fillers for mental level states such as pain (see, e.g. Shoemaker, 1990, p.67). I won't say much about this, though I will point out later that such sub-system states *can't* be the states we introspect, if (or to the extent that) the states we introspect are whole-agent mental states (footnote 27).

difference in introspectible feel, of the kind involved in the inverted spectrum (see also Shoemaker, 1975, e.g. p.310).

There is a problem with such views, though, if we want to look for an explanation of qualitative feel of the kind outlined in section 2. It is not that there are no low level differences with which to explain the alleged difference in feel; as we have just seen, there are. The problem is that there would seem to be no high level difference at all, in the central case of behaviourally undetectable inverted spectra. For two such creatures will do and say exactly the same things. Each will say "it feels like this". If you ask them *how* it feels, they will say all the same things as each other (e.g., "it feels searing"). And so on, and so on. The above model of scientific explanation can only work if we have differences at the low *and* the high levels (e.g. certain stable hydrogen bonds are formed; water freezes). With no difference at the publicly observable mental level, we are left looking for a reason to suppose that there is any mental level difference at all. It is at this point that the various authors mentioned differ.

3.1 Explanation and Reduction

We only have a problem, as regards giving an explanation of the type outlined in section 2, if there are indeed two different levels to relate: a level of mental facts (which do not entail any lower level, non-mental facts), and some non-mental facts (whose existence is not entailed by the mental facts, but which might – if a standard explanation can be given – entail those facts). As we have seen, this is no more nor less than is the case with water versus H_2O, or with heat and temperature versus statistical distribution of energy across microstates. However, in the case of the mental, the existence of such a conceptually separate higher level can be denied.

To see what would be involved in this denial, we need to notice that there are two different ways of understanding the proposal that we should look for an *a priori* analysis of the mental, only one of which I would endorse. I endorse the claim that there is an *a priori* relation between the public notion of pain, and a tendency towards certain behaviours such as wincing, groaning, withdrawing from painful stimuli, etc. But I am endorsing this as a relation amongst facts *at the same level*. Thus pain, wincing, groaning, etc. are all (in the first instance) *mental* level facts[11], just as the properties of macroscopic water (boiling, melting, etc., etc.) are all 'water level' facts.

There is an entirely different reading of the same claim which I would *not* endorse. On this reading, wincing, groaning, withdrawing, etc. are to be read as entirely non-mental facts, and the claim being made, in that case, is that the mental level is not conceptually independent of such entirely non-mental facts. If it is right that the mental is identical (on careful reflection) to some non-mental level of description, then it could be coherently claimed that introspecting the feel of pain is conceptually the same thing as subpersonal detection of a subpersonal state such as 60 Hz neural firing (when this occurs within the right, surrounding subpersonal context).

If this fully 'operationalized' *a priori* analysis of the mental can be carried out, then we don't need to look for an explanatory relation between two levels of description (as outlined in section 2), because there is really only one level of understanding in play.

Endorsement of this latter kind of *a priori* analysis is a very strong form of reductionism about the mental (which is sometimes not clearly enough distinguished

[11] At least, wincing and groaning are mental facts, to the extent that they occur with the right connections to the rest of the mental – see note 9.

from the process of *explanation* outlined in section 2). In many ways, this strongly reductive approach looks like a denial of the reality of the mental level[12], especially when it is made clear that no such conceptual reduction is involved in the explanation of many much less contentious properties[13]. As such, in the rest of the paper, I will discuss what follows if we assume that there *is* a conceptually separate mental level, and that what we are looking for is an *explanatory* relationship between non-mental facts and mental facts (or, at least, an understanding of why we cannot have such a explanatory relationship). On this, at least, I agree with Chalmers, with the authors working on the phenomenal concept strategy (section 4) and with a least one of the authors who historically argued for hybrid functionalism.

3.2 Phenomenal Knowledge

The above strongly reductive analysis would indeed give us a reason to believe in a mental difference between some functionally identical agents: if the analysis is correct, a physical difference of the right type *is* a mental difference. However, if we don't accept the reductive analysis, then we have no third-person reason to believe that there is a mental difference between the supposed inverts. That this is so follows in two steps. Firstly, there is no reason at the public mental level to suppose that there is a difference, for such agents are *the same* at that level. Secondly, the publicly accessible difference which does exist between alleged inverts (on the hybrid-functionalist view) is a physical difference: it lies at a level of description which is not (without further argument) mental. Without the reductive claim, and considering *purely* the third-person facts, there is no reason to believe that that public difference is (or causes, or amounts to) a mental difference.

But not all of the authors who have argued for the compatibility of functionalism and the inverted spectrum intuition endorse a strongly reductionist analysis. Shoemaker, for instance, was not and is not a reductionist about the mental, but he took and takes the inverted spectrum intuition seriously as a starting point for theorising about qualia (Shoemaker, 1975, 1994a, 1994b). It should be emphasized, then, that it follows logically that, if one endorses the strong phenomenal realist view, but rejects reductionism, one *must* take oneself to have a *first-person* reason to believe that the inverted spectrum is possible[14]. It is worth emphasising clearly what this means. Without reductionism, there can be *no reason* to believe in inverted spectra *at all*, unless it is a reason which fundamentally involves first-person knowledge. If such views are right, we *must* be able to come to know by introspection[15] that 'what it feels

[12] One might call such an approach *eliminative reduction*, but it is *not* the same thing as the outright *eliminativism* which the Churchlands argued for elsewhere, concerning the belief-desire framework of folk psychology (see, for instance, the sections on eliminativism in Churchland and Churchland, 1998); one cannot hope to show that 'introspecting phenomenal feel' is conceptually identical to some reasonably well-defined set of subpersonal processes, if one also wishes to show that 'introspecting phenomenal feel' is part of a bad conceptual scheme which does not refer very well to anything at all.

[13] In fairness to the Churchlands' position, I should make clear that they did not accept the analysis of scientific explanation which I have given. Instead they asserted that the pattern of conceptual analysis of role, coupled with *a posteriori* discovery about role filler, is normal elsewhere in science (Churchland and Churchland, 1988, e.g. p.78).

[14] I am ignoring the complications which might follow if, for instance, someone claimed that the inverted spectrum intuition was grounded in fundamentally second-person knowledge.

[15] I will treat 'introspection' as identical to 'the ability to gain knowledge in a fundamentally first-person way'; even if the relevant knowledge is not gained *entirely* through introspection (in this sense), it must be gained in a way which *essentially* involves introspection.

like' is the kind of thing which could differ, even as between two agents who act in all the same ways[16].

Now we can see the connection between the strong phenomenal realist starting points (specifically, the zombie or inverted spectrum claims; though we are mainly considering the inverted spectrum claim, since this is the one popular with many physicalists) and *a posteriori* knowledge. For the knowledge which one is supposed to have, on such accounts, is knowledge which cannot be entailed by (just) the third-person facts, since none of those facts (taken apart from introspective knowledge) give us any reason to believe that there is a mental difference, as we have seen. Equally, if we are not being reductionist about the mental level, then there is no reason to suppose that the mental facts on their own (including any facts known by introspection) entail any lower level, non-mental facts. So here, we have a pure (i.e. two way) *a posteriori* discovery – there are certain phenomenal facts which I know, when I 'look' inwards (i.e. introspect) which I could not have known by looking outwards[17].

It turns out, then, that the same starting points which entailed that there was no publicly accessible high level to explain (in certain key cases) must also entail that phenomenal knowledge is entirely *a posteriori* with respect to (neither entailing nor entailed by) our knowledge of publicly observable facts[18]. Note that this kind of knowledge is strange in that (if it really exists) its existence is *a posteriori* with respect to (i.e. it could not have been deduced from) all knowledge of the third-person facts, however clear thinking and detailed.

Even with the need for this unusual kind of knowledge, perhaps it might still be argued that these views are not so implausible after all. For whilst this is a very special kind of knowledge (c.f. Chalmers, 1996, p.193), it is also knowledge of a special kind of state. Perhaps we should *expect* ourselves to have non-standard and intimate knowledge of those states which partly constitute us? Indeed, in some sense of this suggestion, I would agree with it. But perhaps it is the case that such intimate knowledge *ought* to have these strange *a posteriori* features? Considerably more would need to be said here, to defend this suggestion. As far as I am aware, the hybrid

[16] Informal conversation indicates to me that a large number of (though not all) thoughtful non-philosophers do indeed take themselves to know exactly this; they take themselves to know, presumably on the basis of introspection, that the inverted-spectrum scenario is 'obviously' possible. So this starting intuition, if wrong, is widely (though not universally) shared, at least in this culture.

[17] The disconnect between this alleged knowledge and knowledge of publicly accessible facts is much *stronger* than the 'disconnect' between public knowledge and indexical knowledge (first-person knowledge such as "I am in Sussex", "It is Sunday", etc.). This is because the fact that I can only have indexical knowledge when I am *in* a certain state *follows* from the publicly observable facts, plus an understanding of the concept of indexical knowledge (see Chalmers and Jackson, 2001; related points are made in Beaton, 2005). Whereas the phenomenal knowledge which (allegedly) grounds our belief in the possibility of the inverted spectrum has to be of a quite different type: it might well be possible to *learn* (*a posteriori*) that when I am in a certain physical state, I will be in a certain phenomenal state, but there can be no communicable understanding of the nature of this phenomenal state which could let someone *work out* (*a priori*) that when an agent is in the physical state, the agent *must* be in the related phenomenal state.

[18] Actually, these starting points only strictly rule out an entailment from physical facts to mental facts. There could still be (just) the reverse entailment. This would make (at least some) mental facts more fundamental than any physical facts. This is a form of idealism, and certainly not a rebuttal of the claim that strong phenomenal realism rules out physicalism, which is what I am trying to establish.

functionalists whose position was outlined above never said it[19], but more recent work in the philosophy of mind has stepped in to fill the gap.

4. The Phenomenal Concept Strategy

Loar (1997), and the other proponents of 'the phenomenal concept strategy', embrace the point which I have just made, that what we know about qualia from the first-person is two-way conceptually independent of any facts which science might access. Thus Loar says:

> "Phenomenal concepts are conceptually irreducible in this sense: they neither a priori imply, nor are implied by, physical-functional concepts. Although that is denied by analytical functionalists[20] … , many other physicalists, including me, find it intuitively appealing." (Loar, 1997, p.597)

But Loar also argues that this need not be a problem for physicalism:

> "It is my view that we can have it both ways. We may take the phenomenological intuition at face value, accepting introspective concepts and their conceptual irreducibility, and at the same time take phenomenal qualities to be identical with physical-functional properties of the sort envisaged by contemporary brain science." (Loar, 1997, p.598)

How could such a view work? The general strategy (shared by Loar and others who've published variants of this view) is to concentrate on the special way which we have of introspectively thinking about our own phenomenal states. The claim is that the phenomenal concepts[21] involved in such thoughts ('*this* feeling'; 'like *this*') are special, in that they are "conceptually isolated" (Carruthers and Veillet, 2007) from the third-person concepts which we use when we think about publicly accessible facts. The claim that phenomenal concepts are conceptually isolated does not mean that they cannot occur in the same thoughts as publicly applicable concepts. But it does mean that no amount of *reasoning* can lead from facts expressed using phenomenal concepts (e.g. 'my experience is like *this* now') to facts expressed using publicly applicable concepts (e.g. 'my physical-functional state is this, now'), or *vice versa*.

Apart from this general point about conceptual isolation, the views vary as regards the specific nature of phenomenal concepts which is supposed to explain the isolation. Loar (1997) and others have equated phenomenal concepts with some form of recognitional concept; Perry (2001) has equated phenomenal concepts with some form of indexical concept; Papineau (2002) has suggested that phenomenal concepts are 'quotational' ("my red is like this: ____", where the blank is filled in by the experience itself). As such, all these views are trying to give a more detailed account of the first-person *acquaintance* which we have with our own qualia[22] – i.e. an

[19] Of course, for the reasons outlined, the Churchlands needed no such account. For suggestions from Shoemaker along these lines in more recent work, see Shoemaker (1994b, Section IV).

[20] Loar is referring to the thoroughgoing variety of functionalism which takes *everything* mental to be analysable in terms of its (at least counterfactual) relation to publicly accessible behaviour (i.e. he is not referring to the hybrid variety of functionalism I have just been discussing).

[21] Concepts, in the sense used here, do not require language: rather, they are the recombinable components of rational thought. In the same vein, rationality itself, as used here, should be understood in a sense whereby a rational agent is one which can *make* rational decisions, not necessarily one which can make rational decisions by thinking them through, step by step, in the manner of the most complex human thought.

[22] More accurately (c.f. Chalmers, 2003), an account of the knowledge which such acquaintance can grant us. In the sense in which Chalmers uses the term, the *acquaintance* itself comes in simply having the quale; but this acquaintance is the fundamental ground for later first-person, conceptual knowledge of the quale. It should be noted that a moderate phenomenal realist, type-A (Chalmers, 1996) materialist (i.e. the position which I am trying to defend, or at least open a space for, in the present

account of exactly what seemed to be missing, in the variant of functionalism outlined above.

Can such a view successfully preserve physicalism? A lot has been written about the phenomenal concept strategy, and I don't wish to dismiss it out of hand. Nevertheless, there is a very general argument against the possibility of phenomenal concepts preserving physicalism[23], if physicalism is understood as requiring an explanation of the presence of consciousness in the manner outlined in section 2.

First of all, it is worth noting that the *a posteriori* claim about the nature of phenomenal knowledge (which is so central to the phenomenal concept strategist's position) is not merely entailed by the inverted spectrum starting point (as I have already shown, in section 3.2), but also entails it. To see why this is so, note that the denial of the inverted spectrum starting point amounts to the claim that there *is* always a behaviourally detectable difference, for every difference in qualia. The notion that there exists special *a posteriori* knowledge of the phenomenal is not compatible with this denial of the inverted spectrum. That is, the phenomenal concept strategist cannot accept an analysis of phenomenal concepts which shows that, for every difference known that way, there must be an observable difference in behaviour at the public mental level[24]. If there were such an analysis, a difference in physics sufficient to explain these publicly observable differences would be sufficient to explain the difference in qualia[25] (on the model of the explanation of the properties of water). The connection between the physics and the phenomenal level would not be *a posteriori*, after all.

It might be thought that the phenomenal concept strategist could still claim that, whilst there can be no conceptually necessary difference in behaviour corresponding to a difference in qualia, there still might be a conceptually necessary difference in behaviour corresponding to an agent *knowing* one thing as opposed to another about their own qualia. But actually, they cannot accept this either. Even if qualia are 'covert' when not known about, and only become 'overt' when known about, the normal model of explanation can get a grip. Any physical description which shows why there are these behaviourally observable differences (in the cases where the differences are overt) and why there are no behavioural differences (in the cases where the differences are not overt) will once again explain the physical nature of qualia (on the model of the explanation of water). Once again, the connection between the physics and the phenomenal level would not be *a posteriori*, after all.

I don't think any of this pushes the phenomenal concept strategists to a position which they would be unwilling to accept. It seems very close to (and perhaps actually) explicit in the approach that certain phenomenal differences (and, equally, certain differences in phenomenal knowledge) will not result in any behaviourally detectable difference.

---

paper) can, I think, feel very sympathetic to much of what Chalmers (2003) says about the nature of acquaintance; that is, can feel that very much of it ought to be naturalisable (for a little more on this, see section 6).

[23] The quick argument given below is very closely related to the central argument towards the same conclusion presented in Chalmers (2006). The main difference is that I proceed directly in terms of explicability, rather than via conceivability.

[24] They could perhaps accept the bizarre position that whilst there is no reason (which we could ever understand) for there to be such a difference in every case, it nevertheless turns out that there is such a difference in every case.

[25] It is important to the argument that I specified that for *every* difference known, there is a difference in behaviour – this is what the phenomenal concept strategist cannot accept.

The trouble with all this is that it makes quite clear that the phenomenal concept strategy is entirely incompatible with an explanation of the status of qualia along the lines outlined in section 2. Not only are qualia themselves not naturalisable along these lines, but the special type of knowledge which is supposed to save physicalism is (and must remain) inexplicable for the very same reasons. We seem to be back to square one[26], with no third-person reason to believe that knowledge of this type exists. Even if we do have a first-person reason to believe this (and the final sections of this paper argue against that claim), we are left with an unsatisfying, purely 'ontological' physicalism in which we can have no explanation of why certain things are part of the physical world, merely an acceptance that they are.

In fact, I wonder whether things are not worse than this, for the phenomenal concept strategists. Their claim is that the existence of this type of phenomenal knowledge is itself not entailed by anything which physics can teach us (however well we understand the physics, and the concept of phenomenal knowledge). If this is correct, then surely Chalmers (Chalmers, 1996) has been right all along? Surely all the physical facts might have been exactly the same, and the phenomenal facts might have been different, or absent altogether? At least, if this is not so, physics can't explain why it is not. As such, it looks to me as if Chalmers has been the most honest here, all along. *If* you start from the assumption that there is a pure (i.e. in both directions) *a posteriori* relation between the phenomenal and the physical, or *if* you start from the assumption that behaviourally undetectable inverted spectra are possible, *then* you should end up where Chalmers ends up: you *should* accept that phenomenal properties, and any principles bridging them to normal physical properties, are fundamental facts about our universe.

In the final part of the paper I want to ask two questions. First, what justifications are there for taking the problematic strong phenomenal realist starting point? Second, if the relevant justifications are found wanting, what could we use as a replacement starting point, if we still want to naturalise qualia?

5. The Properties of Sensory Experience

Qualia are properties of sensory experience broadly construed to include states such as seeing, hallucination, sensory memory, sensory imagination, and so on. Furthermore, as we have seen above, if there is any reason to believe that qualia are problematic in the way in which the strong phenomenal realist claims they are, this reason must be introspective.

But there is very little agreement about what sensory experience consists in, and even less agreement as to what the introspectible properties of sensory experience are (c.f. Crane, 2008; Gertler, 2008). I know that I am seeing a scarf on the desk in front of me (it is cold round here, right now!); but can I know that I am seeing this in virtue of some more direct kind of acquaintance with sense data? Sense data theorists certainly thought so, but this view is now widely agreed to be false. Can I know that I am seeing the scarf in virtue of, or at least accompanied by, qualia which can vary free of the physical facts? Chalmers and many others have thought so; but many

[26] Actually, as Chalmers notes (2005, Section 4) the phenomenal concept strategy has at least made the genuine contribution of clarifying that strong phenomenal realism requires an account of this type of knowledge. I would argue (and again, I think most phenomenal concept strategists would be quite happy to agree with me) that the main aim of such accounts must therefore be to convince us that we are wrong to want an explanation of the type I have described, in the case of qualia or of phenomenal knowledge: that physicalism does not require this.

others again *don't* share this certainty. On a related note, the reductionist approach taken by the Churchlands entails that what we know in introspection (of pain states, of colour experience, and so on) includes opaque knowledge of the *physical* nature of certain sub-personal states which underpin these sensory experiences; this, too, is impossible according to many other theories of introspection[27].

Note that all of the above mentioned claims about *perception* (that it involves sense-data; that it entails the possibility of introspective knowledge of the physical states underlying it; that it is accompanied by behaviourally undetectable qualia) constrain our eventual theory of *introspection*, which has to be such as to allow for introspective knowledge of the problematic states in question. Moreover – arguably in all cases, and certainly in the case of the view which is being critiqued here (strong phenomenal realism) – whatever plausibility these starting points have itself derives from introspection.

Sense-data theorists certainly did take themselves to have introspective knowledge of sense-data. It strikes me as highly plausible that this assumption was an input to the sense-data theory, not an output from it; that the theory made explicit what already seemed introspectively obvious. But, it is widely agreed, the theory was false – we have no such knowledge for there are no sense-data.

Equally, as we have seen, at least some physicalist advocates of the inverted spectrum have taken themselves to have opaque introspective knowledge of the *physical* nature of certain of their internal states. Again, is this input or output? With certain implicit, but theoretical, assumptions about introspection under one's belt, it can seem more or less obvious that we do have introspective knowledge of the physical states which constitute us. But actually, the claim that introspection is like this is a major theoretical assumption. It cannot be justified as a starting point, unless we *already* (i.e. entirely pre-theoretically) have introspectively based knowledge, which entails that it is true. Do we have such knowledge? It seems to me very hard to see how we can decide the case either way, simply by introspecting 'harder' or 'more carefully', and very easy to become misled by one's theoretical commitments.

The same points certainly apply to strong phenomenal realism. As we have seen, the *starting* point of the view is this: there is something which we know by introspection, which is a valid basis for the claim that phenomenal facts cannot be deduced from publicly observable facts[28]. Viewed with some perhaps healthy scepticism, this looks very like an implicit, *not* necessarily justified, *theoretical* claim about introspection, which has managed to work itself into the framework of all strong phenomenal realist theories.

With such a wide range of intuitions about introspection, and with an apparent tendency to interpret what we find, when we look inwards, in the light of our (perhaps implicit) theoretical assumptions, it is far from clear whether we are on safe ground, if we make *any* proclamations about what it is that we know when we introspect the features of our sensory states, including qualia.

On the other hand, if we make no proclamations here at all, then we have no way of specifying our target of explanation as we try to understand qualia. Is there a middle

[27] Indeed, this is impossible on any theory in which the facts introspected are all at a conceptually independent mental level, e.g. Sellars (1956), Shoemaker (1996). This follows as long as the conceptual independence of the mental level from the physical is at least as strong as (but it need be no stronger than) the conceptual independence of the water-level from the H_2O level.

[28] And, as we have seen, the view also builds in the claim (which again must be introspectively based, if true) that this non-deducibility is so in a significantly stronger sense than the agreed, but far less surprising, sense in which it is so for indexical facts (note 17).

ground? Is there a way to say *anything*, whilst remaining neutral as between competing theories of introspection? In the final section, I will argue that there is.

6. Some Moderate Subjective Properties

For my part, I am much more certain that there is *something* subjective about my mental life, and that I know this 'something' by introspection, than I am that what I know in this way transcends all physical and functional truths. Therefore, I am proposing that we allow ourselves to be guided, in our quest for qualia, by looking for an independently plausible account of introspection; specifically, we should look for qualia amongst the properties which are introspectible on such an independently plausible account[29].

I have just said that qualia are 'subjective' properties, but of course anything introspectible is subjective in a certain sense, for introspection consists in the ability of a subject to come to know properties of itself in a fundamentally first-person way.

However, I am prepared to concede that some 'subjective' properties, in this sense, are the wrong type of thing to be qualia. Imagine, for instance, a subject seeing a red ball as a red ball (where red, in this case, should be thought of as a public, if gerrymandered[30], property). Essentially any account of introspection must allow that the right kind of subject can introspectively know *that* she is seeing a red ball when she is. This is a specific example of a general type of introspection, whereby a subject becomes aware that they have some 'propositional attitude'-type relationship (believing *x*, desiring *x*, seeing *x*, remembering *x*, imagining *x*, etc.) to some (perhaps only counterfactually existent) *public* object(s) or state of affairs *x*. I will be at least this much of a phenomenal realist: if independently plausible theories of introspection *only* allow that we have introspective knowledge of this type, then such theories do not have the materials to naturalise qualia. If things were to turn out thus, I should (and I think would!) accept that there *are* no qualia, and that I am as much in need of Dennettian therapy (Dennett, 1988) as are all those who maintain that qualia have non-naturalisable properties in the ways discussed in the earlier parts of this paper.

But there seems a very natural next step to take, which is to wonder whether there might not be introspectible properties which are subjective in a slightly stronger sense: to wit, introspectible properties which cannot be fully specified, simply by specifying any number of the non-controversially introspectible properties just mentioned.

So now, imagine two subjects each seeing a red ball *as* a red ball. Imagine, also, that both have agreed on a common language for referring to public properties (red, ball, etc.) and to the 'propositional attitude' type states (including seeing *x*, etc.). Evidently things could be thus, even whilst there are facts about each subject's relation to the world which differ on a perfectly naturalistic account; for example, affective or motivational facts, and facts about the learnt associations between properties (e.g. red reminds one agent of blood and pain, and the other of celebration and good fortune). Now, these facts are subjective in yet a third sense: they are partly

[29] This does not amount to the requirement that qualia should always be introspectible. Whether or not non-introspectible qualia exist will hinge on the details of our theory of introspection, and on the details of any plausible candidate-properties for qualia within such a theory. For instance, on Shoemaker's account of introspection, mental states whose nature is to be introspectible can nevertheless exist in creatures which lack the resources to introspect them (Shoemaker, 1988, Section 3).

[30] This is Dennett's usage, it means that the outlines of what is and isn't red may depend on the constitution and interests of creatures like us, rather than on anything more fundamental about the world.

constitutive of the subject's relationship to the world. But what is not yet clear (at least, until we have an independently motivated account of introspection) is whether any such further facts can be known (perhaps, opaquely) in introspection. If they can be, then they are subjective facts in all three senses: subjective *qua* partially constitutive of the subject; subjective *qua* introspectible; and subjective in the sense just defined, of going beyond the most non-controversially introspectible facts.

In stating that the above is possible, I have not endorsed behaviourally undetectable inverted spectra: for the differences I have mentioned would all be behaviourally detectable. Even so, the situation described is not entirely unlike the standard inverted spectrum starting point. There could indeed be two subjects who see a red ball as a red ball (who even agree, in a shared language, that it is a red ball, and that each is seeing it) whilst there are bona fide introspectible facts about their experience which differ. As such, this seems to me a moderate approach with the potential to explain, rather than completely explain away, the widely held belief that qualia are invertible.

The suggestion that we concentrate on 'motivational, associative and affective' facts is just one proposal, intended to be compatible with the idea of being guided by an independently plausible theory of introspection. But there is a general problem with any proposal of this type, directly related to the two ways of understanding *a priori* analysis noted earlier (section 3.1). It could be taken to imply that the properties in question have been thoroughly "operationalized": that is, expressed in *fully* non-mental terms (setting aside the issue of whether or not this is truly possible). I have already suggested that that approach to *a priori* analysis leads to an overly strong reductionism which should be resisted. Indeed, if qualia are truly mental-level facts, then there is no reason to expect that anything which we know introspectively about them need entail any fully non-mental facts, *even if qualia can be explained on the normal scientific model* (remember that the water facts do not entail the H_2O facts). So the "operationalized" proposal is not the kind of proposal I am making.

Instead, the associative, motivational and affective facts (or whichever facts turn out to best fill the required role) should be read as properties at the independent mental level of description. The question at issue, when the proposal is read this way, is whether there is a conceptual independence between one type of mental description (a thinking, introspecting agent in a certain motivational state, say) and another (an agent having introspectible qualia, say). My suggestion is that we may well be able to find a two-way conceptual interrelation between qualia and the right set of not-so-obviously-qualitative *mental* concepts. If there is, then we would have a coherent account of the entire mental level, including introspectible qualia; and this whole account might yet map onto *some*[31] appropriate description of the physical in the normal way.

Of course, a standard response here is to claim that it is quite conceivable that our qualia are independent of any such (motivational, associative, affective, etc.) facts. Perhaps so, but I am not sure how (or indeed whether) I know that. I have suggested that the prior 'knowledge' of this 'fact', which many presume themselves to have, may be grounded in (implicit) endorsement of perhaps mistaken theories of introspection.

The strategy proposed here may also offer the possibility of explaining, rather than explaining away, other intuitions about the nature of qualia. I am thinking here,

[31] Lest I be misunderstood, I explicitly want to leave open the possibility that the currently popular information processing and representational descriptions may not be best suited for the low-level role in such an explanation.

particularly, of Shoemaker's defence of a "moderate Cartesianism" (Shoemaker, 1988), which looks to be an entirely naturalisable account of a rather direct type of acquaintance we should *expect* to have with *any* introspectible property, on at least one independently plausible, apparently naturalisable, account of introspection[32].

Of course I need, for consistency's sake, to allow that my own presuppositions can be overruled. For any given prior intuition about the nature of qualia, if there are no facts which explain why this intuition was broadly (or even roughly) correct, then qualia do not have the intuited property. And, as I have already conceded, if none of our intuitions about qualia could be naturalised (not even the intuition that there *are* introspectible subjective properties, in the above sense), then there would be no qualia. But there does not yet seem to be any good reason to rule out the suggestion that we may find such properties, within some independently plausible account of the mental level in general, and of introspection in particular.

7. Conclusion

I have argued that what Chalmers calls *phenomenal realism* (Chalmers, 2003) (and what I have called *strong phenomenal realism*) automatically rules out a certain standard form of scientific explanation. I have agreed with Chalmers that the modern phenomenal concept strategy cannot prevent this conclusion. Therefore, if Chalmers is right that the only way to "take consciousness seriously" (Chalmers, 1996) is to be a strong phenomenal realist, then a physicalist account of consciousness cannot succeed. This is certainly the case if physicalism is conceived of as a quest for this type of explanation of the nature of qualia, as I think it should be. But I have also briefly given reason to agree with Chalmers that physicalism cannot succeed on *any* reasonable interpretation, given these starting points.

I have then tried to throw doubt on the strong phenomenal realist starting point which leads to these objectionable conclusions. I have argued that whatever we know about the problematic aspects of qualia, which is supposed to lead us to strong phenomenal realism, must be known through introspection. I have noted that there is much evidence that we are entirely unclear about what we can introspect. I have also suggested that, historically, many theories of *perception* have built into themselves unjustified theoretical commitments as to the nature of *introspection*. I have argued that strong phenomenal realism (an account of the nature of conscious perception) may well be guilty of this same sin.

I have therefore proposed that we take a different approach, and have suggested that, as theorists, we should look for qualia amongst the properties introspectible on some independently plausible theory of introspection. I have noted that on essentially any theory of introspection, we can introspect certain 'propositional attitude'-style states, including "seeing *x*" and "experiencing *x*", where *x* is some (at least counterfactually) public state of affairs. I have therefore defined 'subjective' properties, as those introspectible properties (if any) which can still vary (within or between agents), however many of the basic, uncontroversially introspectible propositional attitude style properties have been fixed. It follows directly from this

[32] Nothing in the present paper establishes that there *are* any properties which are introspectible on Shoemaker's model of introspection, and which are subjective in the sense outlined above. Equally, there are certainly those who argue that Shoemaker's account of introspection is not plausible (Kind, 2003). Finally, I should emphasize that the approach to naturalising qualia which I am proposing remains fundamentally different from Shoemaker's present approach (1994a, 1994b), even though I think his model of introspection is well-suited to be integrated with the approach I offer.

definition that if there *are* such properties, they are *ipso facto* the right kind of thing to explain, rather than explain away, the inverted spectrum intuition. Not, that is, to explain the classic inverted spectrum, which remains incompatible with physicalism, but to explain how something which sounds very much like it is physically quite possible. I have also noted that such properties may be able to explain, rather than explain away, other apparently problematic intuitions about our epistemic relationship to qualia.

If we can find introspectible properties which are subjective in the above, moderate, sense, then we would have achieved some kind of phenomenal realism: there would be introspectible facts which at least come free of the standard propositional attitude facts about an agent. For the reasons given, it strikes me that such properties, if they exist, *are* plausible and adequate naturalizers of qualia. This is clearly not phenomenal realism as Chalmers defines it, but it does seem reasonable to call the present approach moderate phenomenal realism.

In sum, my proposal is that it is plausible and workable to *define* qualia as subjective, introspectible properties in the above moderate sense. Adopting this proposal allows us to be guided, in our attempt to understand qualia, by whatever independently plausible accounts of introspection we have. There currently seems no reason to rule out the suggestion that we may find such properties within an independently plausible, naturalisable, account of introspection.

Acknowledgments

I would like to thank Ron Chrisley, Steve Torrance, Chris Nunn, Tom Beament, Marieke Rohde and Tom Froese for invaluable feedback.

References

Beaton, M. (2005). What RoboDennett Still Doesn't Know. *Journal of Consciousness Studies* 12(12): 3-25.

Carruthers, P. and B. Veillet (2007). The Phenomenal Concept Strategy. *Journal of Consciousness Studies* 14(9–10): 212–36.

Chalmers, D. (1995). Facing Up to the Problem of Consciousness. *Journal of Consciousness Studies* 2(3): 200-19.

Chalmers, D. (1996), *The Conscious Mind: In Search of a Fundamental Theory* (New York: Oxford University Press).

Chalmers, D. (2003), 'The Content and Epistemology of Phenomenal Belief'. *Consciousness: New Philosophical Perspectives*. Q. Smith and A. Jokic (Oxford: OUP).

Chalmers, D. (2006), 'Phenomenal Concepts and the Explanatory Gap'. *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*. T. Alter and S. Walter (Oxford: OUP).

Chalmers, D. and F. Jackson (2001). Conceptual Analysis and Reductive Explanation. *Philosophical Review* 110: 315-361.

Churchland, P. M. and P. S. Churchland (1982), 'Functionalism, Qualia and Intentionality'. *Mind, Brain and Function*. J. I. Biro and R. W. Shahan (Norman, Oklahoma: University of Oklahoma Press)**:** 121-145.

Churchland, P. M. and P. S. Churchland (1998), *On the Contrary: Critical Essays, 1987-1997*: MIT Press).

Crane, T. (2008), 'The Problem of Perception'. *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*. E. N. Zalta.

Dennett, D. C. (1988), 'Quining Qualia'. *Consciousness in Modern Science*. A. Marcel and E. Bisiach: Oxford University Press).

Gertler, B. (2008), 'Self-Knowledge'. *The Stanford Encyclopedia of Philosophy (Winter 2008 Edition)*. E. N. Zalta.

Kind, A. (2003). Shoemaker, Self-Blindness and Moore's Paradox. *The Philosophical Quarterly* 53(210): 39-48.

Levine, J. (1983). Materialism and Qualia: The Explanatory Gap. *Pacific Philosophical Quarterly* 64: 354-61.

Lewis, D. (1980), 'Mad Pain and Martian Pain'. *Readings in the Philosophy of Psychology: Volume I*. N. Block (Cambridge, MA: Harvard University Press).

Loar, B. (1997), 'Phenomenal States (Revised Version)'. *The Nature of Consciousness*. N. Block, O. Flanagan and G. Güzeldere (Cambridge: MIT Press)**:** 597–616.

Papineau, D. (2002), *Thinking About Consciousness* (New York: Oxford University Press).

Perry, J. (2001), *Knowledge, Possibility, and Consciousness* (Cambridge, MA: MIT Press).

Sellars, W. (1956), 'Empiricism and the Philosophy of Mind'. *Minnesota Studies in the Philosophy of Science, Volume I: The Foundations of Science and the Concepts of Psychology and Psychoanalysis*. H. Feigl and M. Scriven (Minneapolis, MN: University of Minnesota Press)**:** 253-329.

Shoemaker, S. (1975). Functionalism and Qualia. *Philosophical Studies* 27: 291-315.

Shoemaker, S. (1988). On Knowing One's Own Mind. *Philosophical Perspectives, 2, Epistemology*: 183-209 (reprinted in Shoemaker, 1996, pp.25-49).

Shoemaker, S. (1990). First Person Access. *Philosophical Perspectives, 4, Action Theory and Philosophy of Mind*: 187-214 (reprinted in Shoemaker, 1996, pp.50-73).

Shoemaker, S. (1994a). Phenomenal Character. *Noûs* 28(1): 21-38.

Shoemaker, S. (1994b). Self-Knowledge and 'Inner-Sense'. *Philosophy and Phenomenological Research* LIV(249-314 (reprinted in Shoemaker, 1996, pp.201-268)).

Shoemaker, S. (1996), *The First-Person Perspective and Other Essays* (Cambridge: Cambridge University Press).