

Chapter 9

Evidence, explanation, and experience:

On the harder problem of consciousness

(Originally published in *The Journal of Philosophy* CI (5): 242-254, 2004)

Preview

Creatures that have different physical realizations than human beings may or may not be conscious. Ned Block's 'harder problem of consciousness' is that naturalistic phenomenal realists have no conception of a rational ground for belief that they have or have not discovered consciousness in such a creature. Drawing on the notion of inference to the best explanation, it appears the arguments to these conclusions beg the question and ignore that explanation may be a guide to discovery. Thus, best explanation can both validate an interpretation of the evidence and lead to the discovery of consciousness.

Chapter 9

Evidence, explanation, and experience:

On the harder problem of consciousness*

I. Introduction

A new type of problem is taking hold in the philosophical debate about consciousness, one that adds to the so-called ‘hard problem of consciousness’.¹ It concerns our epistemic relation to the possible phenomenal states of creatures that are different from us in their physical realization. Ned Block has recently given a thorough account of this ‘harder problem of consciousness’,² and David Papineau, David Chalmers, Joseph Levine, and Thomas Nagel have discussed similar but less substantial versions of this type of problem.³ It arises for naturalistic phenomenal realists, that is, those who have good but defeasible reasons for believing that consciousness has a scientific nature, and who on the basis of first-person grasp of consciousness believe that phenomenal experiences are instantiations of real properties, not something that can be conceptually reduced away in the manner of analytical functionalism.

Block's harder problem of consciousness (HPC, hereafter) is that naturalistic phenomenal realists face an epistemic tension: they have no conception of a rational ground for believing that other creatures, who do not relevantly share our physical nature, are conscious or not. Even if they assume that the other creature is conscious, they have no conception of a rational ground for believing that consciousness in this creature has a physical basis, even though physicalism is the default view. Furthermore, it is the subjective default that other creatures are not conscious, but the phenomenal realist must leave it an open question whether they are.

Block discusses HPC in terms of a conditional whose antecedent is the conjunction of naturalism, phenomenal realism and anti-scepticism about other minds (the latter is necessary to distinguish HPC from the other minds problem). The consequent is the epistemic tension. I discuss two aspects of Block's arguments that this conditional is true.

First, I assess the argument to the conclusion that naturalistic phenomenal realists have no conception of a rational ground for belief in disjunctive physicalism. I argue that it is possible to resist the conclusion because, given a plausible notion of inference to the best explanation, the argument begs the question (Section 2).

Second, I assess the argument that naturalistic phenomenal realism produces an epistemic obstacle to discovering phenomenality in other creatures. This argument ignores the

priority of explanation over discovery in inferences to the best explanation (section 3). I further explore the relationship between evidence, explanation and conscious experience: some of the hardness of HPC may stem from doubts about the occurrence of the evidence, or from doubts about the interpretation of the evidence. But the first of these doubts is misplaced, and the second may be overcome, again with help from the notion of inference to the best explanation (section 4).

My criticism is not directed at the grounding intuition that there are epistemic obstacles in the study of consciousness; it is pretty clear that we have this intuition. It is directed at the arguments that Block mounts, on the basis of this intuition, to describe the nature and consequences of our current epistemological predicament.

II. The argument for meta-inaccessibility of disjunctive physicalism

Throughout, Block makes use of the science fiction case of Commander Data, who is a merely superficial functional equivalent to us, equivalent to us in respect of folk-psychology and what folk-psychology entails, but not in other psychological and neuropsychological functions. Importantly however, Data is unlike us in the physical nature and organisation of the control mechanisms of the folk-psychological functions (pp. 401–402).

We have good but defeasible grounds for believing that Data is conscious because he acts like us, and we act the way we do in part because we are conscious. The grounds are defeasible because we might find that Data's physical constitution shares none of our neural correlates of consciousness (pp. 402–403).

One of Block's central premises is that we have no conception of a rational ground for the belief that Data is or is not conscious. He is stipulated to be physically different from us, so any science of consciousness that is based on us will be powerless to determine whether he is conscious or not. And Block argues further that we cannot base a science of consciousness on him unless we already know he is conscious (I return to this argument in section 3). So we have no idea what would be evidence for the belief that Data is conscious or not; Data's consciousness is meta-inaccessible (p. 405).

If the phenomenal realist is a physicalist who thinks that consciousness has a deep, non-disjunctive physical basis, then Data is not conscious because he does not have this basis. But this seems unduly chauvinist. A weaker disjunctivist physicalism allows that Data can be conscious. Disjunctivism says that consciousness is the heterogeneous physical property of having either our or Data's realization of our shared functional state.

Disjunctivism is attractive to the naturalistic phenomenal realist because it allows that we and Data can have overlapping phenomenality with a common, albeit disjunctive,

scientific basis. The problem is that, even though it is a default that consciousness has a scientific nature, we have no conception of a rational ground for inferring that disjunctivism is true. It is very odd that someone who believes that consciousness is real and has a scientific nature must also acknowledge that there is no conception of what the evidence for the naturalist position would be (pp. 407–8, 413–414).⁴

Block's argument to this conclusion begins by noting that one needs to arrive at disjunctivist physicalism via inference to the best explanation. The identification of phenomenality with the one disjunctive property of our or Data's realizations must contribute to the best explanation of facts about phenomenality, such as why there is phenomenal overlap between us and Data (pp. 410–412). This is similar to the way in which the identification of water with H₂O contributes to the best explanation of facts about water, such as that it expands when freezing (pp. 409–410). Of course, this only makes sense if it is assumed that Data is conscious: there is no best explanation of the phenomenal overlap, if he is not conscious (p. 411).

There are now two arguments against arguing via inference to the best explanation to the conclusion that consciousness is a heterogeneous physical disjunction.

The argument about fundamental difference. The first argument in effect begins by stipulating that our and Data's physical realizations are fundamentally different in a sense

that make us balk at a disjunctive explanation of facts about phenomenality. When it comes to two fundamentally different properties, distinct explanations are better than one disjunctive explanation. Block illustrates the point by saying that we could, in theory, explain solid-like formation in terms of crystallization or continuous hardening into an amorphous substance (which would make glass a solid). But we reject the disjunctive explanation and prefer two distinct explanations because the properties of crystallization and continuous hardening into an amorphous substance are fundamentally different. Block doesn't specify the basis for judgements of fundamental difference, he just notes that we have a vague notion of it and then proceeds to the above stipulation (p. 411).

Block's argument leads up to the conclusion that disjunctivism is meta-inaccessible: that we have no conception of a ground for rational belief in disjunctivism. If our and Data's realizations are fundamentally different, then there can be no argument via inference to the best explanation to a physical disjunction (though disjunctivism may of course still be true).

Nevertheless, this use of the notion of fundamental property difference is questionable. To begin with, it seems doubtful whether a merely intuitive notion of fundamental difference should guide our grounds for rational belief. Intuitively (i.e., in the absence of knowledge of the scientific truth), diamonds and graphite are fundamentally different, but

they are surely fundamentally similar, since science tells us they are different crystallizations of carbon.

Then there is the case of solid-like formation. It is my understanding that water can freeze into a glassy, amorphous solid when it is cooled rapidly,⁵ just as in the case of window glass, which is sand that is melted and rapidly cooled. So here we have solid-like formation—freezing into glassy ice—that is explained in terms of continuous hardening, not crystallization. This opens the possibility of a disjunctive explanation, at least by the indicators employed by Block: glassy ice is *not* a super-cooled liquid (glassy solids are more viscous, though still metastable, counterparts of super-cooled liquids⁶), and its formation *is* a kind of freezing. But then a reasonable case can be made that we have an explanation of solid-like formation in terms of a disjunction of intuitively fundamentally different properties, viz. continuous hardening-or-crystallization. So intuitive fundamental difference is not in itself sufficient to close off the feasibility of disjunctive explanation.⁷

The point is not that the discovery of glassy ice straight off should make us accept the disjunctive explanation; it is that the discovery makes it reasonable to suggest that the explanation is disjunctive, in spite of the intuitive fundamental difference. An epistemically defensible notion of fundamental property difference may therefore depend on empirical work and on explanation, and is thus not so easily up for uncontroversial

stipulation. In particular, it seems plausible that judgements of fundamental property difference and similarity are, at least in part, based on how well assumptions of fundamental similarity or difference contribute to overall explanatory integration.

The underlying thought is that we wouldn't rationally care much about fundamental similarity and difference unless it had an impact on the overall explanatory picture. If we achieve the best explanatory integration of a wide range of phenomena by one grouping of properties rather than another, then we go with the first. Allowing the disjunctive explanation of solid-like formation may be messier, and less conducive to overall integration than going with the two distinct explanations (it may be less messy in the overall scientific picture to count glass as a liquid). Conversely, once we find out about the glassy solid state of water we might well decide that it is the disjunctive picture that is less messy. It is hard to think of an alternative basis for judgements of fundamental property difference, short of some kind of aesthetic preference that we can hardly allow to rationally inform these epistemic issues.

On this view, the stipulation that our and Data's realizations are fundamentally different isn't independent of issues concerning best explanation. Rational judgements of fundamental property difference come *after* we evaluate which explanations are best; we cannot avail ourselves of them first and then use them to tell us which explanations are best. But then the best way to make sense of the stipulation presupposes that there is no

argument from inference to best explanation to a physical disjunction. This begs the question for that is the conclusion Block is aiming at.

Objection: it is possible to obtain scientifically informed views on fundamental property difference on the basis of all *other* explanations than explanations of consciousness, and then to pass judgement on the fundamental difference of our and Data's realizations.

Reply: this is unduly conservative: even if they were judged different on the basis of other explanations, the best-making qualities of explanations of consciousness may override this judgement (just as the best-making qualities of the explanation of solid-like formation may override previous judgements, once we learn about glassy solids).

Objection: Block's argument explicitly concerns our epistemic predicament in a pre-explanatory situation (p. 406, 424); so it begs the question against Block to invoke what we may say about property difference once explanations become available. *Reply:* the argument is not that HPC will disappear once explanations of consciousness become available, it is rather that Block faces a dilemma: be irrationally guided by pre-explanatory judgements about fundamental property difference or beg the question by presupposing that there is no best disjunctive explanation.

The argument about bad and good questions. Block's second argument against arguing via inference to the best explanation to the conclusion that consciousness is a

heterogeneous physical disjunction concerns the neatness of best explanations.

Considerations of neatness say that best explanations should solve more explanatory puzzles, and not give rise to bad explanatory puzzles. The theoretical identification of water and H₂O, and heat in gases and mean molecular kinetic energy, prevent bad questions about why it is that water is *correlated* with H₂O and why it is that heat is *correlated* with mean molecular energy from arising; there is no explanation of why things correlate, if they are identical. According to Block, this is part of the reason why we accept these identities.

Now consider the question of why we overlap with Data in this or that phenomenal respect. This question does not arise if we accept disjunctivism. If there is phenomenal overlap, then this just *is* our having the disjunctive property of our or Data's realization—and identities do not have explanations. Block now claims that the question does not seem bad, and therefore should not be ruled out, so, accordingly, we should not accept the theoretical identification. (This is not evidence that Data is not conscious since, for the phenomenal realist, it is still an open question whether he in fact is). If inference to the best explanation cannot do the job, then there seems no source of a conception of a rational ground for belief in disjunctivism.

However, we have to ask what makes us deem that certain why-questions are good and others bad. In an argument about the epistemic viability of physicalism, we cannot let us

be rationally guided merely by how questions *seem*; that is too arbitrary. Moreover, it seems plausible that such judgements can be wrong: what we once deemed a good question may later transpire as a bad question, and vice versa. My view is that it is explanation that can make us change our minds about the quality of why-questions. We say that the question about why it is that water is correlated with H₂O is bad because the best explanation of facts about water involves acceptance of the theoretical identification of water and H₂O. If we had no such explanation it is very plausible that it would be deemed a good question, at least by some (such as staunch phlogiston-believers who are just coming to terms with Lavoisier's notion of *oxygen*).

Some explanations rule out some questions as bad and other explanations other questions. But it is only after we have the overall explanatory picture that we can make an informed judgement about which questions were good and which bad. What this shows is that the ruling out of bad questions, in Block's sense, is not an independent best-maker for explanations, and thus should not in itself be used to plead acceptance or rejection of theoretical identities. But then, under the assumption that there is as yet no explanation, arguments based on these notions are not available, unless the question is begged in favour of there not being a best disjunctive explanation.

Objection: What you argue is that as long as there is no explanation of consciousness, there will be no resolution of these problems, but isn't that precisely the situation Block

describes, viz. that in our current epistemic situation the questions about physicalism about consciousness are meta-inaccessible? *Reply*: Block is arguing *for* the claim that physicalism is meta-inaccessible; it is not a premise in his argument. My argument is that this argument doesn't go through without begging the question. The arguments in this section do not concern the *premise* that we have no conception of a rational ground for belief that Data is conscious or not.

III. Explanation and discovery

Block's premise about the meta-inaccessibility of Data's consciousness plays an important role in the overall argument. The argument is that any science of consciousness that can generalize to other creatures must be based in part on them in the first place, but that this cannot be done unless we already have discovered whether they are conscious or not (pp. 406–407). The premise about Data prevents this discovery. So, as things stand, there will be no generalizable science of consciousness.

Behind this argument is the following epistemic principle about the priority of discovery over explanation:

[DE] There can be no explanation of a phenomenon unless it is already discovered that it occurs.

I think DE gets some of its intuitive plausibility from the following metaphysical principle:

[OE] There can be no explanation of a phenomenon that does not occur.

For example, there can be no explanation of the occurrence of global peace if it doesn't occur (though of course there can be explanations of the non-occurrence of a phenomenon, like the non-occurrence of global peace). OE is weaker and less controversial than DE because it doesn't require any discovery of the occurrence of the phenomenon; it is consistent with a range of epistemic states vis a vis the phenomenon, and with discovering its occurrence after having explained it.

In his influential discussion of inference to the best explanation, Peter Lipton notes that the point of explaining is not just to explain *why* something is the case; in the shape of inference to the best explanation it is also our primary tool for discovering *what* is the case.⁸ One of Lipton's examples illustrates this. Consider the red shift observed in the light spectrum of some stars. This was taken to be evidence of some unknown, undiscovered phenomenon. Various hypotheses can be proposed to explain the evidence, and the one that best explains the evidence is the one we come to believe, in this case that

the star recedes with a certain velocity. So we used best explanation to discover something about the world. The principle is:

[ED] Sometimes best explanation can lead to discovery.

This principle does not contradict DE. DE says that the explanandum phenomena must already be discovered. In the recession example, the explanandum phenomenon is the red shift, which *is* discovered. The discovery that is relevant to ED concerns the explanans, that is, the recession. The question is, however, whether ED will help circumvent the role that DE plays in Block's argument. Can explanation of various kinds of phenomena lead to discovery of consciousness in other creatures?

What is the best explanation of the evidence of folk psychological functionality in Data? We would normally reason from like effects to like causes and say the explanation is the occurrence of phenomenal states. But this inference is not sound here because we know that Data's realization is different from ours, and this defeats the argument by analogy.

Consider instead the possibility of inferences to hypotheses that explain the organising principles of the physical properties of Data's silicon brain and how they give rise to the folk-psychological functionality. One of these hypotheses may explain some of this evidence partly in terms that clearly identify the phenomenal property of pain with a

particular pattern of silicon brain activity. Another may explain it partly in terms of the identity of a non-conscious, merely cognitive property and that particular pattern of silicon brain activity. We should accept the hypothesis that best explains the evidence. If this hypothesis is the one that identifies phenomenal properties with silicon brain properties, then that is the one we should believe. That is, best explanation can lead to the discovery of phenomenality.

Objection. This discussion of HPC presupposes that we have a conception of explanations of how a physical property can be identical to a phenomenal property but HPC is explicitly set in a pre-explanatory context (pp. 406, 424). *Reply:* This is irrelevant to the issue at hand, which concerns an argument about the relative priority of discovery and explanation. It begs the question in favour of the thesis that discovery must precede explanation if we cannot assume for the sake of argument that explanations are available.

Objection: HPC is not a principled problem, it doesn't attempt to say that we *could* not have a conception of a rational ground for belief in Data's being conscious or not. Rather, it is contingent on the fact that we have not yet managed to explain how a physical property can be identical to a phenomenal property, and thus it merely argues that our current epistemic predicament about consciousness is of a certain nature and has its origin in naturalistic phenomenal realism. It is no criticism of this view to note that HPC is resolved once we arrive at such explanations. *Reply:* It is one thing to note our strong

intuition that, if confronted with Data, we would have no conception of a rational ground for belief that he is conscious or not. I do not deny that we have this intuition. It is another thing to *argue* that we need to discover phenomenality before we can have a generalizable science of consciousness, or indeed that disjunctivism is not epistemically viable. Perhaps our intuitions support these conclusions (though I am not sure about that), but it doesn't follow that the arguments leading up to the conclusions are convincing. My criticism so far is that these arguments are not convincing. So we are left with the intuitions.

IV. Doubts about the occurrence and interpretation of the evidence

The argument in the preceding section establishes that explanation can lead to the discovery of consciousness in other creatures, and thus we can circumvent Block's claim that there cannot be a science of consciousness for creatures like Data. It is however possible to resist this conclusion by changing the context of doubt so that it no longer concerns what phenomena we have and have not discovered, but instead concerns our justification for believing in the occurrence of the evidence. This kind of doubt may be the source of the hardness of the harder problem, and it does indeed lead to a very hard problem, but the underlying doubt is misplaced. Another kind of doubt, about the interpretation of the evidence, is not misplaced but can be dealt with by once again employing properties of inference to the best explanation.

Doubts about the occurrence of the evidence. Here is how the doubt arises. We can grant that best explanation can lead to discovery, but only if there are no doubts about the occurrence of the evidence for the explanation. For example, explanation of a red shift can lead to discovery of recession, but only if there are no doubts about the occurrence of the red shift (was it a genuine red shift, or an artefact of our measuring devices?). Similarly, the doubt goes, ‘silicon brain science’ explanations of the evidence concerning the organising principles of the physical properties of Data’s silicon brain and how they give rise to the folk-psychological functionality can lead to discovery of consciousness, but only if there are no doubts about the occurrence of the evidence. But there are such doubts, for perhaps the organising principles are designed to mimic genuine folk-psychological functionality and are therefore no more evidence of phenomenality than a tape-recording of introspective reports.

We cannot respond by claiming that if there *is* a best explanation in terms of phenomenality then there are no grounds for doubting the evidence. This kind of argument leads to vicious circularity in cases where the phenomena that are being explained are themselves essential parts of the evidence for the explanation. If someone doubts the occurrence of the red shift in the light spectrum of a star, then we cannot justify belief in its occurrence with reference to the best explanation (i.e. in terms of recession), since the red shift is evidence for the recession. Similarly, if someone doubts

the occurrence of the evidence concerning the organising principles of Data's brain and the functionality, then we cannot justify belief in its occurrence with reference to a best explanation (i.e. in terms of phenomenality), since it is an essential part of the evidence for that explanation.

This is a general feature of scientific explanation, and in particular of inference to the best explanation. Some explanations are self-evidencing, that is, the phenomena they explain are themselves part of the evidence for the explanation. There is circularity here, but it is benign. It turns vicious, however, when the aim is to use best explanation to establish the occurrence of the evidence.⁹

Hence, in the context of doubts about the occurrence of the evidence, it seems there is no way of forming a conception of a ground for rational—that is, non-circular—belief in Data's being conscious or not, even if there are available explanations that identify phenomenal properties with some physical properties of his brain! In this version, the problem posed by Block's premise about the meta-inaccessibility of Data's phenomenality is therefore very hard. In contrast to the hard problem, it will not immediately go away once we have a conception of what explanations of phenomenality will be like.

Doubts about the interpretation of the evidence. Luckily, this is not a good way to characterise the doubt that one may have about the evidence. For in a very real sense there is no doubt about the occurrence of the evidence. Data's folk psychological behavior, and the silicon circuitry in his brain, is in principle there for all to observe. We have no reason to introduce some fault in our methods for acquiring this kind of evidence.

I think there is a better way to capture the relevant doubt about the evidence, but that the problem posed by this doubt is solvable such that we can have a plausible conception of a rational ground for belief, also post explanation. It is possible to view the problem as concerning the interpretation of the evidence rather than the occurrence of the evidence. Sometimes people disagree about the acceptability of various hypotheses because they disagree about the standards of interpretation of the evidence, not because they disagree about whether the evidence confirms the hypothesis. Such disputes seem intractable because any further evidence that gets drawn in to settle the matter is itself subject to the different standards of explanation (which is why this kind of dispute also plays a role in Kuhn's notion of incommensurability).

By once again employing the resources of inference to the best explanation it is nevertheless possible to decide such disputes.¹⁰ The standards of interpretation are partly driven by theories, and if the explanation of the evidence is much better on one

interpretation of it than on another, then this vindicates the theory behind the better explanation. For example, we can imagine two different interpretations of the red shift evidence: one in terms of a largely modern astronomical theory, and one in terms of a terracentric, epicycle theory. On the first, it is evidence for and is explained by recession, on the second it is evidence for and is explained by, let's say, some new kind of epicycle. No new astronomical evidence can decide between the hypotheses, but we can nevertheless decide between them because the modern theory allows a much better explanation than the terracentric theory. So the dispute is not intractable, if we have at least one good enough explanation, and if there are independent standards of best-ness of explanation (in Kuhnian cases of incommensurability we may not have shared best-ness standards either).

I think a good part of the epistemic worries about naturalistic phenomenal realism can plausibly be sourced to this kind of doubt about the evidence. Anything a proponent of phenomenality in Data interprets as evidence of phenomenality (or of a particular phenomenal state), the doubter can subject to interpretation under the theory that there is no phenomenality, only functional equivalence (or that it is a different phenomenal state), and vice versa. This explains why we have no conception of what kind of evidence could decide the matter.

Assume now that two different silicon science hypotheses about Data's brain are offered to explain what goes on when his silicon brain produces his apparent introspective reports of, for example, "being in pain". The first hypothesis involves good, well-integrated explanations that show how it can be true that the phenomenal property of pain is identical to particular physical properties of silicon brain activity. The second involves bad, poorly integrated explanations that show how some other phenomenal state such as the experience of jealousy, or some non-conscious merely cognitive state, is identical to particular physical properties of silicon brain activity.

If we should choose the best explanation, then we are given reason to believe that Data's reports of "being in pain" are reports of that phenomenal state and, further, that there is phenomenal overlap with us when we are in pain. If the second hypothesis had been the better one, then that would be reason to believe that the reports are not of the phenomenal state of pain, and that there is no such overlap.

Hence, if we construe the problem as one concerning the interpretation of the evidence, then inference to the best explanation can again come to the rescue.¹¹

Objection: this gets us nowhere because a best explanation must among other things explain the phenomena, and we have no conception of a rational ground for believing that one explanation of phenomenality is better than another (cf. p. 426). *Reply:* this

misconstrues the debate. What is true is that we have no conception of what a scientific explanation of phenomenality may be like: of how physical properties could be identical to phenomenal properties, or how subjective concepts and objective concepts could pick out the same property. This is what remains of the hard problem of consciousness, according to the conceptual dualism espoused by Block (pp. 395–398). If we did have such a conception, then we could employ our general notion of inference to the best explanation to reach a decision on Data's being conscious or not; naturalistic phenomenal realism does not throw up a further, separate issue about conceptions of rational grounds for belief.

V. Concluding remarks

The question is whether naturalistic phenomenal realism produces some epistemic problems for the science of consciousness, on top of the problem that it is fiendishly difficult to explain consciousness, and that we currently have no real idea how to go about it. The answer is that consciousness is a singularity because it is so evidentially insulated and that this makes things difficult, but that it doesn't put the science of consciousness in a league of its own. The epistemic problems associated with naturalistic phenomenal realism are tractable, using the normal resources of scientific practice, in particular as provided by the notion of inference to the best explanation. The trouble with seeing this is that debate of the resources of inference to the best explanation works best

on the presupposition that there *are* explanations to infer to, and we resist making that presupposition in the case of consciousness. But then it is no wonder epistemic worries crop up, because the relevant epistemic questions are best assessed in terms of properties of inference to the best explanation.

* Thanks to the audience at the Annual Meeting of the Danish Philosophical Association, February 2003, especially to Ned Block; and to the audience at a seminar at the Philosophy Program, RSSH, ANU. The research for this paper was supported by a grant from the Danish Research Council for the Humanities to the research project *Naturalised Mind – Cognisant Nature* (namicon.au.dk).

¹ See David Chalmers, *The Conscious Mind* (Oxford: Oxford University Press, 1996).

² Ned Block, “The harder problem of consciousness,” *The Journal of Philosophy* XCIX (2002): 1–35. Page references are to this article unless otherwise noted.

³ David Papineau, *Thinking about Consciousness* (Oxford: Clarendon Press, 2002), Chapter 7; David Chalmers, “Availability: The cognitive basis of experience?” in N. Block, O. Flanagan, and G. Guzeldere, (eds.), *The Nature of Consciousness*. (Cambridge: Mass: MIT Press, 1995); Joseph Levine, *Purple Haze: The Puzzle of Consciousness* (Oxford: Oxford University Press, 2001), Chapter 3; and Thomas Nagel, “The psychophysical nexus,” in P. Boghossian and C. Peacocke (eds.). *New Essays on the A Priori* (Oxford: Clarendon Press, 2000), fn21.

⁴ Block also considers another physicalist doctrine, superficialism (p. 412), which I will ignore here.

⁵ Philip Ball, *H₂O – A Biography of Water*. (London: Weidenfeld & Nicholson, 1999), p. 187. Block notes this phenomenon as well, and says we do not regard it as freezing (p. 411)

⁶ Osamu Mishima and H. Eugene Stanley, “The relationship between liquid, supercooled and glassy water,” *Nature* 396 (1998): 329–335.

⁷ We could take this to show that *solid-like formation* is a merely nominal category, like *jade*. But this would beg the question since it in effect says that only nominal categories can have disjunctive explanations.

⁸ Peter Lipton, *Inference to the Best Explanation* (London: Routledge, 1991), pp. 57–58, 68–69.

⁹ Carl G. Hempel, *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science* (New York: Free Press, 1965), pp. 370–372; Lipton, *op. cit.*, pp. 26–27.

¹⁰ Cf. Lipton, *op. cit.*, pp. 71–72.

¹¹ This type of strategy is further explored, in the context of a concrete neuroscientific example, in Jakob Hohwy and Chris Frith “Can neuroscience explain consciousness?” forthcoming in *Journal of Consciousness Studies*, June-July (2004).