

## **Chalmers and the Self-Knowledge Problem**

Robert Bass  
Department of Philosophy  
Coastal Carolina University  
Conway, South Carolina 29526  
rbass@coastal.edu

In *The Conscious Mind: In Search of a Fundamental Theory*, David Chalmers poses an interesting and powerful challenge to materialism or physicalism. Further, he goes a long way towards providing a proof by example that the rejection of materialism need not commit one to scientifically suspicious “ghost in the machine” doctrines, but can be wedded to a generally naturalistic perspective. As an (as yet) unpersuaded physicalist and functionalist, his case against physicalism seems an appropriate target for criticism. However, it would be beyond the scope of the present paper – or anything of similar length – to attempt a full-scale reply. A full-scale reply could be executed in brief compass only if Chalmers were guilty of (and rested his case entirely upon) some relatively simple mistake. But he does not. If he is mistaken, the mistakes are more subtle and difficult to bring to light. Instead, I shall outline the essentials of his case against materialism, attempt to point to *one* problem that appears to infect his views and argue that, if his position is to be acceptable, he needs to deal with it more adequately than he has so far.

### *Some Terminological Distinctions*

In order to present Chalmers’ central argument against materialism, some terminology needs to be introduced in order to make it clear what he is (and isn’t) arguing. First, he distinguishes between what he calls the *psychological* and the *phenomenal* concepts of mind. The psychological concept of mind is “as the causal or explanatory basis for behavior. A state is mental in this sense if it plays the right sort of causal role in the production of behavior, or at least plays an appropriate role in the explanation of behavior.” (Chalmers, 11) By contrast, the phenomenal concept is of mental states as consciously experienced states; to be in a phenomenal mental state is, borrowing from Nagel, for there to be “something it is like” to be an organism in that state. “On the phenomenal concept, mind is characterized by the way it *feels*; on the

psychological concept, mind is characterized by what it *does*.” (Chalmers, 11)

Mental states in both the phenomenal and in the psychological senses, of course, often appear in tandem and, for at least some cases, may be perfectly correlated. It may be, for example, that all cases of intense *felt* or phenomenal pain are associated with states that, *ceteris paribus*, control aversive behavior. Psychological mental states, Chalmers is prepared to concede, may be entirely explained in terms of the properties, entities and relations countenanced by (completed) physics or physical sciences, where physics is to be taken as “the fundamental science developed to explain observations of the external world.” (Chalmers, 119) It is a further question, however, whether the same physical sciences can explain phenomenal experience.

An additional necessary distinction is that between *logical* and *natural* supervenience. We can say first that supervenience is the relation that holds if and only if properties of one type – “B-properties,” generally considered to be of a higher level, are related to properties of another type – “A-properties,” generally considered to be of a lower level, in such a way that “no two possible situations are [qualitatively] identical with respect to their A-properties while differing in their B-properties.” (Chalmers, 33) Filling in some details, B-properties would *logically supervene* on A-properties if a full specification of the A-properties uniquely determined the B-properties as well in all logically possible worlds. No two logically possible worlds would be identical in A-properties without also being identical in B-properties.<sup>1</sup> As an example, Chalmers suggests that there are no two situations which are identical with respect to their physical properties but which differ biologically – as to whether the same entities are or are not alive. Being alive logically supervenes upon certain physical properties. There is another sense in which there might be no two possible worlds or situations “identical with respect to their A-properties while differing in their B-properties.” It might be that there are no two situations

identical with respect to A-properties but differing with respect to B-properties in worlds in which the same natural laws hold. In such a case, we can say that the B-properties *naturally supervene* upon the A-properties. To illustrate, Chalmers points out that though a stable cubic mile of uranium-235 seems logically possible, it is not naturally possible in any world in which the relevant laws of nature are the same as ours because instability naturally, though not logically, supervenes upon being a cubic mile of uranium-235 in all possible worlds that are identical to ours with respect to laws of nature. (Chalmers, 37)

When B-properties logically supervene upon A-properties, then there is a relation of *entailment* between a full specification of the A-properties and a corresponding full specification of the B-properties. It will not be logically possible that the A-properties are identical in two cases while the B-properties are not. When B-properties naturally supervene upon A-properties, there is no entailment from a full specification of the A-properties to a full specification of the B-properties, unless laws of nature are added to the specification of the A-properties. Or, to put it the other way around, the A-properties *plus* the laws of nature entail the B-properties.

The relation between these two sorts of supervenience is one of inclusion: All cases of logical supervenience are also cases of natural supervenience, but not all cases of natural supervenience are cases of logical supervenience.

### *Chalmers' Central Argument*

With the help of these distinctions, we can state Chalmers' central argument. Suppose that we take the relevant A-properties to be all and only those properties, entities and relations countenanced by a completed physics and to include any natural laws necessary to such a completed physics. Since physics is, for his purposes, defined as "the fundamental science

developed to explain observations of the external world,” there will be no reference to conscious or phenomenal states in this completed physics. Second, let us take the relevant B-properties to be the facts of phenomenal or conscious experience that occur in our world.

Chalmers’ view is that if physicalism is true – that is, if it is true that consciousness has no features that are not identical to or accountable for in terms of physical properties, entities, facts and relations – then all facts about phenomenal states would have to logically supervene upon a complete microphysical description of the universe. The argument for this can be put simply: If phenomenal states do not logically supervene upon the physical (including physical laws), then the physical world could be the same while the phenomenal states could be different or absent. Therefore, settling the physical state of the world is not sufficient to settle what phenomenal states it contains, if any. Something else – something *ex hypothesi* non-physical – would have to be added.<sup>2</sup>

As noted earlier, Chalmers has agreed that all psychological properties, those causally relevant to the control and modulation of behavior, can be accounted for in terms of a completed physics and thus are logically supervenient upon the physical. Now, we are in a position to consider whether the same holds for phenomenal properties. Chalmers argues that it does not. It appears that we can conceive of a world physically identical to ours in which the phenomenal properties are different or absent.

To make this more concrete, he offers the example of a world (or situation) in which there are zombies. Zombies are physically identical with beings (such as ourselves) who possess phenomenal consciousness. However, zombies do not themselves have any phenomenal consciousness. Since psychology (in Chalmers’ sense) logically supervenes upon the physical, they will behave exactly as we do, make the same verbal reports, modulate their behavior in the

same ways that we do, and so on – but there will, in all this, be no phenomenal consciousness on their part. “There is nothing it is like to be a zombie.” (Chalmers, 95)

Are such cases conceivable? Chalmers claims that they are and that such conceivability is a good guide to what is logically possible. If he is correct in both these claims – and I shall not directly challenge either – then it appears that he has provided a powerful argument against physicalism. If zombies are even logically possible, then fixing all physical facts, properties and relations is not enough to fix phenomenal properties. And if it is not, then the phenomenal properties cannot be accounted for by even a completed physics. Something else needs to be said to account for them, so physicalism is mistaken.<sup>3</sup>

Even if Chalmers is mistaken about either claim – that zombies are conceivable or that their conceivability shows that they are logically possible – his careful development of the arguments should at least shift the burden of proof to those of a materialist persuasion: It seems to be up to materialists to show either that the alleged conceivability is somehow an illusion or that the conceivability of zombies does not show that they are logically possible.

### *The Problem of Self-Knowledge*

I am not at all sure that the materialist (at least, this materialist) can discharge that burden of proof. Instead, I will attempt something more oblique. In outline, I shall maintain that although Chalmers has a valid argument and one that may be sound, that does not yet entitle us to believe its conclusion for we may not be in a position to know that its premises are true. One of the crucial presuppositions of the argument is that there are phenomenal properties that stand in need of explanation. However, it seems that there is a difficulty, if Chalmers’ account is correct, in our knowing that there *are* any phenomenal properties.<sup>4</sup>

Consider a possible world that is physically identical to ours down to the smallest microphysical detail but which is inhabited (as counterparts of ourselves) solely by zombies. Since our psychological features are, as Chalmers grants, logically supervenient upon physical facts, the psychological features of the zombies, all those relevant to the control and modulation of their behavior, will be identical to ours. On every occasion in our world in which any mental event, state or process makes a difference to the physical course of events, exactly the same thing, making precisely the same difference, will occur in the zombies' world. There is even, in their world, a zombie edition of David Chalmers, disconcerting zombie materialists with powerful arguments that the reality of phenomenal experience and the fact that it does not logically supervene upon the physical shows that materialism is false.

Note now that everything in the zombies' world (we may assume) *does* logically supervene upon the physical. Thus, materialism is true of them; there are, in their world, no facts, entities or relations that are not identical to or that cannot be completely accounted for in terms of the facts, entities, relations and laws of a completed physics.<sup>5</sup> The zombie materialists are right, even if they cannot find a convincing answer to the zombie David Chalmers.

At this point, I think the problem should be apparent: How does David Chalmers (*our* David Chalmers) know that ours is not the zombie world? If it were, it appears that neither he nor anyone else would be able to tell. Insofar as belief, knowledge, consideration of evidence and argument can be understood to be properties or states relevant to the modulation of behavior,<sup>6</sup> they will occur in exactly the same way in the zombie world as in ours. Zombies will report their beliefs, experiences and inner states just as we do. They will argue that some mental events shape their behavior without being conscious and argue about the extent to which their conscious lives are the product of such unconscious influences. They will distinguish the

conscious from the unconscious and the phenomenal from the psychological. Beings in the zombie world will hold exactly the same (“merely behavioral”) beliefs about phenomenal experience that we do. But they will be mistaken, for there are no phenomenal experiences in their world.

To repeat the question, how does (our) David Chalmers (or anyone else in our world) know that we are not zombies – however unbelievable we may find that thought? (Remember, zombies in a physically identical world *also* find it unbelievable.) It is a question Chalmers has considered. In fact, he provides a forceful statement of two different forms of it:

[The argument] might proceed directly from the possibility of my zombie twin. My zombie twin makes the same phenomenal judgments that I do. Where I judge that I am conscious, he judges that he is conscious. Further, his judgments are produced by the same *mechanisms* as my judgments. If *justification* accrues to judgments solely in virtue of the mechanisms by which they are formed, as is often supposed, then the zombie’s judgments will be as justified as mine. But surely his judgments are not justified at all. After all, they are utterly and systematically false. It seems to follow that *my* judgments cannot be justified either. They are produced by the same mechanisms that are responsible for deluded judgments in a zombie, and so they surely cannot qualify as knowledge. (Chalmers, 192)

and

It is often held that the crucial factor in justifying a belief is an appropriate causal connection between the belief and the entity it is about. My beliefs about the table I am looking at, for example, are justified at least in part by the fact that

the table is causally responsible for the beliefs ....

But experience is causally irrelevant, or so I am conceding for now. A conscious experience plays no causal role in the formation of a judgment about that experience. If a causal theory of knowledge is correct, it follows that we cannot know anything about our experiences. (Chalmers, 193)

His response is what it appears that it would *have* to be: “A property dualist [such as himself] should argue that a causal theory of knowledge is not appropriate for our knowledge of consciousness, and that the justification of our judgments about consciousness does not lie with the mechanisms by which those judgments are formed.” (Chalmers, 193) That much seems clear. The interesting question at this point is what he proposes to put in place of such accounts and how satisfactory that alternative is.

His proposal is that “it is *having* the experiences that justifies the belief .... My experiences are *part* of my epistemic situation, and simply having them gives me evidence for some of my beliefs.” (Chalmers, 196) If this is correct, he can say that he is in a different epistemic situation than his zombie twin, for his zombie twin does *not* have the experiences. Thus, he is justified in believing in his own phenomenal experiences while the zombie twin is not.

Will this do as an account of Chalmers’ (or anyone’s) knowledge of his own experiences? I cannot see that it will. More precisely, it seems to me that the notion of justification is slipping through our fingers. It has been known, at least since the time of Plato, that knowledge cannot be identified with true belief. Some condition about justification in having a true belief is also necessary.<sup>7</sup> Chalmers says that having the experiences is what justifies the belief, but, in this

case, the belief in question is just that the experiences are being had. That seems to amount to nothing more than saying that the belief is true, so I fail to see what Chalmers has that his zombie twin lacks, besides a true belief.

Let me try to make this more pointed. Giving an adequate account of what it is for belief to be justified is no easy task. (If it were, epistemologists would be out of business!) The justification condition seems to require at least this much, however – there must be some connection between the belief and the truth which is its content. It must not simply be accidental that the belief is true.<sup>8</sup> I will not, for present purposes, insist that the connection must be either causal or via a reliable belief-forming mechanism or mechanisms.<sup>9</sup>

Now, if zombies are logically possible, it follows that the belief that one has conscious experience (or the presence of that belief) is not identical with having conscious experience nor does it entail it. The belief (or its presence) cannot be identical with having conscious experience or entail the having of conscious experience because, if either of those were the case, zombies, who have the belief, would also have the experience. Nor can the belief (or its presence) render it more probable than not that one has the experience. Since, for every logically possible (physical) world in which there are beings with conscious experience, there is, if Chalmers is correct, a physically identical possible world none of the denizens of which have conscious experience, there can be no better than a fifty percent chance that a belief that one has conscious experience is true.<sup>10</sup> This seems tantamount to saying that the belief that one has conscious experience *is* only accidentally related to the fact and, therefore, that the justification condition is not met.

Actually, however, matters are still worse.<sup>11</sup> To see how, let us take a closer look at the belief in one's own consciousness. As I have described it so far, it is rather colorless. Since

holders of the belief do not (typically) think themselves to be continuously conscious throughout their lives, their belief has to be that they are conscious at some time or times. We can call this the General Claim. Our actual beliefs about our own conscious experience (mine, anyhow) hardly ever take this form. Instead, they are much richer and more detailed, having to do with particular episodes of conscious experience undergone on particular occasions. I am now, for example, conscious of looking at a computer screen, feeling pressure from the keys on my fingertips as I type, hearing an air conditioner in the background – or so I believe. These are Particular Claims – each of which has the content that one has or undergoes some specific conscious experience.<sup>12</sup> Further, I take it that I am not at all unusual in this respect.

What relations are there, for an individual, between the General Claim and Particular Claims? First, and obviously, if any Particular Claim is true, then so is the General Claim. Second, and equally obviously, if the General Claim is false, then all Particular Claims are false as well. Not quite so obvious, but to be elaborated below, are the third and fourth points, that from the falsity of any one or any set of Particular Claims, nothing follows about the truth or falsity of the General Claim and that from the truth of the General Claim, nothing follows about the truth or falsity of any Particular Claims.

Now, in order to be consistent with Chalmers' premises, we have to assume that the probability that we can justifiably assign to the truth of any Particular Claim or to the truth of all the members of any proper subset of any set of Particular Claims is less than unity, for if we were justified in assigning a probability of unity (to any one or any set), then we would also be justified, since the truth of any Particular Claim entails the truth of the General Claim, in denying that any being physically identical to a holder of that belief or set of beliefs could be a zombie. In fact, for any one or any set of Particular Claims, for reasons indicated above, we could not

justifiably assign a probability greater than fifty percent.

There is a further question to be pressed: Can we justifiably assign a probability to any Particular Claim or set of Particular Claims *as high as* fifty percent or, equivalently, a probability of unity conditional upon the truth of the General Claim? We can simplify this to the question whether (again, conditional upon the truth of the General Claim) any single Particular Claim can justifiably be assigned a probability of unity. For it is impossible, for a finite set,<sup>13</sup> for any subset to have a probability of unity unless at least one of its members does. The probability that all are false may be vanishingly small, but so long as no single member has a probability as great as unity, there is a non-zero probability that *all* are false.

So, given the truth of the General Claim with respect to some person, could there be any Particular Claim she accepts that can properly be assigned a probability of one hundred percent? It is hard to see how any could, unless Chalmers' more general views are mistaken. Any Particular Claim to the effect that she is conscious at some time, in some particular way, since it is part of her psychology (in Chalmers' sense), will supervene logically upon her physical constitution at the time she holds the belief. That she is conscious at the relevant time,<sup>14</sup> however, will supervene *naturally* (if at all) upon her physical constitution during the (supposedly) conscious episode. In other words, its supervenience will depend upon some contingent law of nature that relates her physical state to the occurrence of the conscious episode. Hence, the law of nature governing the supervenience of conscious states upon physical states could be different – could connect conscious states to different physical parameters – in such a way that, though she is indeed conscious at some time (so the General Claim is true of her), she is not conscious on that occasion.

So, given the truth of the General Claim with respect to some person, no Particular Claim

of hers can justifiably be assigned a probability of one hundred percent. And, since that high a probability cannot be assigned to any single Particular Claim, there is also no set of Particular Claims such that there is a probability of one hundred percent that it contains at least one true member. Since there is no such single claim or set of Particular Claims, then, for the general case in which the stipulation that the General Claim is true is withdrawn, for any Particular Claim, it is more likely to be false than true, and also for any set of Particular Claims, it is more likely that they are all false than that any of them is true.

Moreover, for anyone who issues a large number of Particular Claims, the probability that any one of them is true must be, not just less than fifty percent, but very small. Suppose, for instance, that I make two such claims each of which has a probability of thirty percent. That's equivalent to saying that each of them has a seventy percent chance of being false. Given that, however, the probability that *both* are false is only forty-nine percent – which is equivalent to a fifty-one percent chance that at least one of them is true.<sup>15</sup> But that violates the condition that no set of Particular Claims can have a probability of containing a true member greater than fifty percent. This is a serious limitation because the condition is easily violated even when the probabilities attaching to each claim and the number of such claims is relatively small. If, say, each Particular Claim is assigned a probability of one percent, then there can be no more than sixty-eight such claims, because, if each of sixty-nine claims has a ninety-nine percent chance of being false, the probability that they are all false drops below fifty percent (since  $.99^{69} < .5$ ). For people who make hundreds or thousands of Particular Claims in the course of their lives, it must be virtually certain that each is false.

This is troubling in the following way. To this point, I have spoken, rather coyly perhaps, of Particular Claims in a way that abstracts from two of their features. The first has to do with

the way in which Particular Claims enter the discussion. They are not simply proposals offered for our assent or dissent that we can hold at arms' length until we decide what to make of them. They enter the discussion *as beliefs* and, of course, as beliefs, are taken to be true. It may be disturbing to realize that, on Chalmers' account, these claims, which do not come up for assessment at all unless they are believed to be true, are probably all false.

But there is a further source of disquiet having to do with a second feature of Particular Claims. It is not only that they are believed to be true when they probably are not, but that they play a crucial role in what we take to be our justification for accepting the General Claim. The General Claim, as it applies to a given person, holds that that person is conscious at some time or times. Now, I think it likely that some people do not accept the General Claim, not because they believe it to be false but because they have never considered the matter in that form and thus have never formed any judgment as to whether or not it is true. However, for those of us who do accept the General Claim with respect to ourselves, we infer it *from* one or more of the Particular Claims that we accept.<sup>16</sup> Of course, this would be eminently reasonable if either any Particular Claim was certain to be true or if Particular Claims typically had some non-trivial probability of being true.

Since neither of these last two conditions is met by Chalmers' account, we are presented with a theory which, on one hand, assigns a probability of fifty percent to a certain claim, but on the other, implies that none of the evidence from which we in fact infer the truth of that claim, can justify a probability even as high as fifty percent. In addition, in the present case, the claim to which the fifty percent probability is assigned is one that all of us – all who have considered the question, anyhow<sup>17</sup> – accept. It is not, as one might expect would be warranted by that kind of evidential situation, that we are undecided about the matter nor (another possible response)

that some of us think it true, while others think it false. Now, if there is any way of showing that this combination of claims is inconsistent, I do not know what it is, but surely, even if they are consistent, the combination is sufficiently implausible that we should rethink the reasoning that leads to it.

### *Conclusion*

What, then, is the justifying relation between the belief that one is conscious and the fact that one is conscious? Whatever Chalmers thinks it is, it appears that the state of believing that one is conscious would not be different if one were not conscious nor does being in that state (or in any state of belief from which one might infer that one is conscious) stand in any positive probabilifying relation to being conscious.

Why then does this relation deserve to be called “justification”? If Chalmers is right about the logical possibility of zombies, then any person in any possible world physically identical to ours who believes that he is conscious is no more likely to be correct than mistaken. It appears to me that, despite a few suggestive but undeveloped remarks about the “intrinsically epistemic” (Chalmers, 196) status of experience, he has really added nothing to the claim that, in his case, the belief is true. But since true belief is not equivalent to knowledge, he has not given any grounds for supposing that he knows what he claims (nor has he suggested any way in which we *can* know or be justified in the belief that we are conscious).

A possible reply that Chalmers could make – in fact, I think that in various ways, he *does* make it or rely upon it – is that it is just incredible to claim that we do not know that we have conscious experience. I agree that it is incredible, but I think this counts against him rather than in his favor. To the extent that we are certain that we *really* have conscious experience, we

should be suspicious of any account, such as Chalmers', according to which, in the first place, we are no more likely to be correct than incorrect in believing that we have conscious experience and, in the second place, the evidence upon which we actually rely to support belief in our own consciousness – namely, the Particular Claims we accept – is radically defective, since, (a) taken individually, each one is almost certainly false, and (b) taken jointly, they are more likely to be false than true.

If what I have argued is correct, it appears, not necessarily that Chalmers is mistaken, but that if he is correct, neither he nor we are in any position to know or reasonably conclude that he is. A stripped-down version of his argument can be put this way:

1. There are phenomenal experiences.
  2. Phenomenal experiences are not logically supervenient upon the physical.
  3. If anything actual is not logically supervenient on the physical, then physicalism is false.
- 

Therefore, physicalism is false.

That is certainly valid. It may also be sound. I have said nothing to directly argue that any of the premises are false. However, if it is sound, I have argued, we do not know and, to all appearances, *cannot* know that it is. Therefore, unless something (consistent with Chalmers' other central claims) can be done to shore up the claim that we can know or be justified in believing that we have phenomenal experience, it is not rationally incumbent upon us to reject physicalism even if Chalmers' argument is sound.<sup>18</sup>

## References

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

---

<sup>1</sup> Actually, this needs some further qualification (which Chalmers provides, 38-41) to deal with cases in which there are *additional* B-facts or properties in some possible world to those that appear in some other world that is identical with respect to A-facts or properties. For present purposes, I think we can ignore these qualifications, though that detracts somewhat from precision. I shall not, in any case, base any criticisms upon claims about such possible worlds.

<sup>2</sup> “If B-properties supervene logically on A-properties, then once God ... creates a world with certain A-facts, the B-facts come along for free as an automatic consequence. If B-properties merely supervene naturally on A-properties, then after making sure of the A-facts, God has to do more work in order to make sure of the B-facts: he has to make sure there is a law relating the A-facts and the B-facts.” (Chalmers, 38)

<sup>3</sup> Chalmers also considers the position, suggested by some, that there is some other relevant sense of possibility than logical or natural possibility. This – ‘metaphysical possibility,’ perhaps – might provide the materialist with a way to say that zombies are logically possible but not metaphysically possible and thus that there are no metaphysically possible worlds in which phenomenal experience does not logically supervene upon the physical facts. Chalmers argues convincingly against such notions of metaphysical possibility, 136-138.

<sup>4</sup> Of course, I do not mean to deny that there are any phenomenal properties or even that we know that there are any. My position is rather that if Chalmers’ position is correct, then we would not know that there are any phenomenal properties and therefore, would not be justified in accepting his argument.

<sup>5</sup> Consider the objection, brought to my attention by Sara Worley, that there is something true of the zombies’ world that does not logically supervene upon its physics, namely, that it contains no conscious experiences. That, of course, has to be true if zombies are logically possible. However, it can hardly be Chalmers’ point to argue, as this might suggest, that if physicalism is not logically necessary, then it is logically impossible. The summary statement that there is no conscious experience in the zombies’ world is not entailed by any of the facts of zombie physics, but still, every true statement that asserts that something has some mental property is entailed (trivially, for the case of phenomenal properties, since they do not figure in any true statements affirming their presence in the zombies’ world) by some fact or facts of zombie physics.

<sup>6</sup> It might be objected that zombies do not have real beliefs, real knowledge or make real judgments (etc.). Some phenomenal experience is essential for the subject’s states to qualify as believing, knowing or judging. I think that’s not true in general. First, we all have beliefs that are never consciously formulated, and it is difficult to see how phenomenal experience can be essential to *their* identity. Second, even for the clearest case, beliefs expressed as occurrent thoughts, it doesn’t seem that any *particular* phenomenal experience is essential to them. The same belief may have differing phenomenology on different occasions on which it is expressed as an occurrent thought, so it can’t be that its phenomenal concomitants on either occasion are essential to it. But even if it is essential to our being genuine believers that there be some phenomenal concomitants (or others) of some of our beliefs, the issue can be by-passed. It can be granted that zombies do not have real beliefs, only pseudo-beliefs that are identical with respect to all physical and behavior-modulation features to real beliefs, and the arguments I shall present can be reformulated in those terms. (“Zombies pseudo-believe that they are conscious; conscious beings really believe in their own consciousness. What warrants the conclusion that we have real beliefs rather than pseudo-beliefs?”)

<sup>7</sup> Since Gettier, it has been known that “justified true belief” would not do the job either. Nonetheless, I think (and shall assume) that some kind of justification is necessary if not sufficient for knowledge.

<sup>8</sup> This much is implied just by the claim that true belief is not sufficient for knowledge.

<sup>9</sup> What other kinds of connections there might be, I don't know. Perhaps God knows things without standing in causal relations to them or employing any belief-forming mechanisms.

<sup>10</sup> Here and below, I speak of the probability assignable to certain claims about consciousness, where the relevant population consists of possible worlds that are physically identical (including identity with respect to physical laws) but which differ as to the physical features upon which consciousness does or does not supervene. Such claims about probabilities may be thought to be troublesome if, plausibly, there are infinitely many distinct possible supervenience laws connecting physical parameters to conscious experience. I do not think this objection works in the end, but let us suppose for the moment that it does. Then, for any infinite set,  $S$ , which can be exhaustively partitioned into two proper infinite subsets,  $A$  and  $B$  (whether discrete or continuous will not matter, so long as both  $A$  and  $B$  are either discrete or continuous), we may not be able to say, of a randomly selected member of  $S$ , that it is more probable that it is a member of  $A$  or that it is more probable that it is a member of  $B$  or that it is equally likely to be a member of  $A$  as to be a member of  $B$ . No determinate (even ordinal) probability judgment will be true. But then, the claim that one is conscious must either be reasonable in some other way (than by having some assignable probability) or not. If it is not reasonable (or more reasonable) in some other way, then the justification condition is not met. If it is, then some account is needed of what this other way of being reasonable is.

<sup>11</sup> In an earlier version of the present paper, I essayed a mistaken argument to the effect that, given Chalmers' assumptions, the belief that one was conscious was *probably* not true. This could not work because, so long as it is being assumed that, for the class of persons (defined psychologically, in Chalmers' sense), consciousness supervenes naturally but not logically upon some feature or features of the physical constitution of persons, then, if in some possible world,  $\alpha$ , a given person lacks the feature or features upon which consciousness supervenes, there is a physically identical possible world,  $\beta$ , in which the relevant natural law so differs that exactly the persons who lacked consciousness in  $\alpha$  are conscious in  $\beta$ , and vice versa. In other words, in  $\alpha$ , the natural law governing the supervenience of consciousness upon physical constitutions will specify some (possibly complex) feature of persons,  $F$ , as necessary and sufficient for them to be conscious, while, in  $\beta$ , the corresponding law will specify that it is the absence of the feature  $F$  that is necessary and sufficient for consciousness. So, the probability must remain at fifty percent.

Further reflection convinced me of my mistake but also led to the recognition that there was a different way – the argument to be developed in the text – in which it was problematic to think that, on Chalmers' assumptions, we could be justified in believing in our consciousness.

<sup>12</sup> I assume, here and elsewhere, that belief in the General Claim does not count as a Particular Claim. Also, I am concerned in this section only with *one* way in which Particular Claims may be false, namely, whether there is a conscious episode (of some type) when there is believed to be. A Particular Claim might, of course, fail to be true in other ways, e.g., by misidentifying the character of a conscious experience.

<sup>13</sup> I take it to be obvious that the set of anyone's Particular Claims, however large, is still finite.

<sup>14</sup> The timing of the belief need not, of course, be the same as the timing of the conscious state. I can believe now, for example, that I was conscious of a brief rainstorm last night.

<sup>15</sup> This assumes that the probabilities that two or more Particular Claims are true are independent. I'm not sure what to make of the suggestion that the probabilities may not be independent. Normally, absence of probabilistic independence is a sign of some kind of direct or indirect causal relation, but conscious events, on Chalmers' account, do not enter into the right kinds of causal relations.

<sup>16</sup> I confess that I have little argument for this except that it seems entirely obvious to me. I think that if I did not first believe that I was conscious on some particular occasions, it would never occur to me to suspect, much less to believe, that I was nonetheless conscious at some time or other. Perhaps it *would* occur to me (and even be believed by me) if I could treat consciousness as having some kind of causal or explanatory relevance to my behavior. (Something like that is, after all, why I believe in the occurrence of unconscious mental events.) But that kind of causal or explanatory role for consciousness is just what is excluded on Chalmers' account.

<sup>17</sup> As mentioned above, it may be that there are people who neither believe nor disbelieve that they are conscious by virtue of not having formulated such a proposition.

<sup>18</sup> Of course, that is not equivalent to saying that there is any rational requirement that we accept physicalism.