

**THE CONCEIVABILITY ARGUMENT AND
TWO CONCEPTIONS OF THE PHYSICAL***

Daniel Stoljar
The Australian National University
The University of Colorado at Boulder

I

The conceivability argument (CA) against physicalism¹ starts from the premises that:

- (1) It is conceivable that I have a zombie-twin, i.e., that there is someone who is physically identical to me and yet who lacks phenomenal consciousness; and
- (2) If it is conceivable that I have a zombie-twin, then it is possible that I have a zombie-twin.

These premises entail that physicalism is false, for physicalism is the claim—or can be assumed for our purposes to be the claim²—that:

- (3) For any subjects S and S^* and worlds W and W^* , if S in W and S^* in W^* are physically identical, then they are psychologically identical.

If (3) is true, it is not possible that I have a zombie twin: anybody physically identical to me is psychologically identical, and must therefore also have phenomenal consciousness. If (1) and (2) are true, however, it *is* possible that I have a zombie twin. Hence, if physicalism is true, either (1) or (2)—or both—is false.

One response for the physicalist is to reject (1). The notion of conceivability that is at issue in (CA) is most usefully thought of as a kind of rational intuition or intellectual presentation of a possibility: a clear and distinct idea that something could be physically identical to me and yet psychologically different.³ But the problem is that such ideas are very hard to pin down. It's not at issue that we sometimes *have* these intuitions—that we have them is an assumption of many discussions in philosophy, and will not be questioned here. It is

rather than introspective judgments concerning the character and content of these intuitions are notoriously subject to influence. As Descartes says in the *Meditations*, in order to be clear that we are having the intuitions we say we are, we need to withdraw to a quiet place and meditate. Physicalists who respond to (CA) by rejecting its first premise are in effect saying that we haven't meditated enough.

A different response for the physicalist is to reject (2). In general it is plausible to suppose that intuition or conceivability should be thought of as a guide to possibility—just as perception is a guide to actuality—in the following sense: if I conceive that p is possible (or have the rational intuition that p is possible) then I am *prima facie* justified in believing that p is possible. But the trouble is that conceivability is not an infallible guide. From the mere fact that you conceive that p is possible it does not *follow* that p is possible.⁴ In effect, then, what (2) says is that in the zombie case there are no defeating conditions for the inference. However, it is open to physicalists to reply that there *is* a defeating condition. They might say that inferences from conceivability to possibility are always subject to defeat, and thus never legitimate—this would be to endorse a form of skepticism about knowledge of modality. More plausibly, they might say that such inferences are not legitimate in certain important cases, some of which are the cases relevant to (CA).

In this paper I defend a third response. I will argue that (CA) involves a fallacy of equivocation: there is a sense of 'physical' according to which the premises of (CA) are true, and there is a sense of 'physical' according to which physicalism is true; the trouble for (CA) is that these are distinct senses. In my view, this sort of response captures what is attractive in the previous two responses but evades what is controversial. What it captures and what it evades will become clearer as we proceed.

II

I will begin by making a two preliminary remarks about the strategy I want to consider. First, our strategy raises a set of considerations which are distinct from, but might be confused with, those raised by what is often called Hempel's Dilemma.⁵ Hempel's Dilemma is that if the notion of the physical is defined via reference to contemporary physics, then physicalism is false; but if it is defined via reference to future or ideal physics then physicalism is trivial. The conclusion of Hempel's Dilemma is supposed to be either that there is no notion of the physical, or that there is no clear notion, or that there is no notion that can play the role that philosophers want to the notion of the physical to play.

Now, the strategy I want to consider here presupposes that this conclusion is mistaken. I will assume not only that there is *a* clear notion of the physical but that there are *two* clear notions (or that there are two spellings out of a single clear notion). What then has gone wrong in Hempel's Dilemma? This question demands a paper in its own right, but to put it rather briefly, Hempel's

Dilemma forgets that we all have and understand an ordinary or commonsensical concept of the physical. What Hempel's Dilemma shows, if it shows anything, is that a particular proposal about how to analyse this concept—viz., that it should be analysed via reference to scientific physics at a particular stage of its development—is mistaken.⁶ But from this one cannot conclude that one has no clear concept at all. After all we all have and understand plenty of concepts that we don't know how to analyse (Lewis 1970). Indeed, in this paper, I will not attempt to analyse the notion of the physical at all. Rather, I will simply present two ways in which that concept might be understood. Whether that means that we in fact have two rather different concepts, or that we have two spellings out of the same concept, is a question I will set aside.

The second preliminary is this. One interesting aspect of the strategy I want to consider is that, if successful, it allows us to answer (CA) in a way that avoids taking up controversial issues in the epistemology of modality. Many philosophers are attracted by the idea that the right answer to (CA) appeals to necessary a posteriori truths such as those discussed by Kripke (1980). Such truths, it is widely held, make available the possibility that the second premise of the conceivability argument—(2)—is false. Now, the mere existence of such truths is not controversial or at least should not be. Hence there should be nothing controversial if a response to (CA) appeals to necessary a posteriori truths—indeed, as I will note briefly later on, our strategy is at one point naturally paired with such an appeal (cf. fn. 13).

Nevertheless, there *is* a large controversy over the correct *theory* of the necessary a posteriori. According to one side of this controversy, every necessary a posteriori truth is a logical or a priori derivation from a contingent a posteriori truth and a necessary a priori truth; according to the other side, at least some necessary a posteriori truths are not derived in this way. Many contemporary philosophers who appeal to necessary a posteriori truths—*a posteriori physicalists*, as I will call them—are committed not only to the idea that necessary a posteriori truths can be marshalled in a response to (CA), they are committed to the stronger claim that such truths must be underived. But this *is* to take up a controversial position; indeed, the question of whether there are underived a posteriori necessary truths is one of the most controversial questions in the epistemology of modality. So an advantage of our approach to (CA) is that it allows us to sidestep this issue.⁷

III

Having made these preliminary points, we can now turn to the distinction between the two conceptions of what it is to be a physical property. According to the first—which I will call *the theory-based conception*—a physical property is a property which *either* is the sort of property that physical theory tells us about *or* else is a property which logically (or a priori) supervenes on the sort of property that physical theory tells us about. According to this conception, for example, if physical theory tells us about the property of having mass,

then having mass is a physical property. Similarly, if physical theory tells us about the property of being a rock—or, what is perhaps more likely, if the property of being a rock supervenes on properties which physical theory tells us about—then it too is a physical property. Let us say that any property which is physical by the lights of the theory-based conception is a *t-physical property*.⁸

According to the second conception of the physical—which I will call the *object-based conception*—a physical property is a property which *either* is the sort of property required by a complete account of the intrinsic nature of paradigmatic physical objects and their constituents *or* else is a property which logically (or a priori) supervenes on the sort of property required by a complete account of the intrinsic nature of paradigmatic physical objects and their constituents.⁹ According to this conception, for example, if rocks, trees, planets and so on are paradigmatic physical objects, then the property of being a rock, tree or planet is a physical property. Similarly, if having mass is a property of paradigmatic physical objects—or if the property of having mass is required in a complete account of the intrinsic nature of such objects—having mass is a physical property. Let us say that any property which is physical by the lights of the object-based conception is an *o-physical property*.

Since the concept of the physical is one of the most fundamental in our conceptual scheme, many philosophical issues arise in connection with it. But for our purposes, the important thing is that the two conceptions are not co-extensive: some o-physical properties are not t-physical. The following passage from Frank Jackson provides a convenient way to develop the point.¹⁰

When physicists tell us about the properties they take to be fundamental, they tell us about what these properties *do*. This is no accident. We know about what things are like essentially through the way they impinge on us and on our measuring instruments. It does not follow from this that the fundamental properties of current physics, or of ‘completed’ physics, are causal cum relational ones. It may be that our terms for the fundamental properties pick out the properties they do via the causal relations the properties enter into, but that at least some of the properties so picked out are intrinsic. They have, as we might put it, relational names but intrinsic essences. However, it does suggest the possibility that i) there are two quite different intrinsic properties, *P* and *P**, which are exactly alike in the causal relations they enter into, ii) sometimes one is possessed and sometimes the other, and iii) we mistakenly think that there is just one property because the difference does not make a difference (as the point is put in information theory). An obvious extension of this possibility leads to the uncomfortable idea that we may know next to nothing about the intrinsic nature of our world. We know only its causal cum relational nature ... I think we should acknowledge as a possible, interesting position one we might call Kantian physicalism. It holds that a large part (possibly all) of the intrinsic nature of our world is irretrievably beyond our reach, but that all the nature we know about supervenes on the (mostly or entirely) causal cum relational nature that the physical sciences tell us about. (1998, p. 26-7)

This description of Kantian physicalism allows us to distinguish two classes of properties. The first class comprises the causal cum relational properties

that physical sciences tell us about—to put things in our terms, these are the t-physical properties. The second class comprises the properties that duplicate physical objects will share—to put things in our terms, these are the o-physical properties. As Jackson in effect points out, it is easy to imagine two objects *O* and *O** which are identical in respect of their causal cum relational properties but differ in their intrinsic properties. To put things in our terms, it is easy to imagine two objects *O* and *O** which are alike in respect of their t-physical properties, but differ in respect of (some of) their o-physical properties. But this is just to say that there are some o-physical properties which are not t-physical, and that the two conceptions of what it is to be a physical property are not co-extensive.

In suggesting that there are two conceptions of the physical, I do not mean to suggest that there are no controversial assumptions at play in the argument for their distinction suggested in this passage from Jackson. On the contrary, there are at least three such assumptions. First, there is the issue of why physical theory should tell us only about causal cum relational properties. Second, there is the issue of why, in addition to the causal cum relational properties, there are the intrinsic properties which underlie them. Third, there is the analysis of the notion of ‘tells us about’—clearly, for a theory to tell us about a property cannot simply mean for the expressions of the theory to refer to that property, since (as Jackson points out) it is clear that expressions of physical theory could refer to the intrinsic properties if there are any.

However, while these assumptions are certainly controversial, I will not attempt to defend them here. I have three reasons for this. First, so far as I am aware, proponents of (CA) accept the assumptions, so there can in the context be nothing objectionable about the procedure of adopting them so as to draw the distinction at issue.¹¹ Second, the assumptions have nothing at all to do with philosophy of mind, so one can appeal to them without begging any of the questions that arise in the context of (CA). Finally, and most importantly, the metaphysical picture described by Jackson, and our development of that picture, is in the end only the scaffolding of our answer to the conceivability argument. As I will argue in §VI, we can do without the scaffolding and nevertheless defend the view at issue.

IV

In addition to the two conceptions of the physical, there is one further piece of machinery that is required before we can turn our full attention to the conceivability argument. This is a distinction, due in its modern form to James Van Cleve (1983), between weak and strong conceivability. As I noted earlier, the notion of conceivability that is in play in (CA) is the notion of a clear and distinct idea or presentation of a possibility. Now, this is a very hard notion to clarify. The Cartesian phrase ‘clear and distinct idea’ is obviously only a gesture at a fuller treatment. I will not attempt any fuller treatment here beyond endorsing Van Cleve’s remark that “there is such a thing as just ‘seeing’—by a

kind of intellectual vision—that a proposition is true. Each of us can ‘see’ in this way that $2 + 3 = 5$, that nothing is both round and square, and much else besides” (1983, p. 36).

However, it *is* possible to distinguish a stronger and a weaker claim of what conceivability might consist in, as follows:

- (4) *S* strongly conceives that *p* just in case *S* has a clear and distinct idea that *p* is possible.
- (5) *S* weakly conceives that *p* just in case it is not the case that *S* has a clear and distinct idea that *p* is impossible.¹²

One might bring out the distinction here by pursuing the analogy between conceiving and perceptual experience. There is a difference between being visually under the impression that there is a lamp on the table, on the one hand, and *not* being visually under the impression that there is *not* a lamp on the table on the other. Likewise, there is a difference between conceiving that something is possible, on the one hand, and *not* conceiving that it is *not* possible on the other. The first of these is strong, the second weak, conceivability.

Weak and strong conceivability differ from each other in a number of ways, but only two of these are important for us. First, the slogan ‘conceivability is a guide to possibility’ is only plausible when the notion of conceivability at issue is strong conceivability. If it visually appears to me that there is a lamp on the table, I seem thereby to be *prima facie* justified in believing that there is; hence one might reasonably say being visually appeared to is a good guide to actuality. But suppose it *is not* the case that it visually appears to me that there is *no* lamp on the table. We would not normally say that that in such cases I am thereby *prima facie* justified in believing that there *is* a lamp on the table. Perhaps if episodes of this kind were repeated often enough I would have some sort of justification. But this is not what we mean when we say that perceptual experience is a guide to actuality. Likewise, while it is true that if I conceive that *p* is possible I am thereby *prima facie* justified in believing that *p* is possible, it is not true that if I do *not* conceive that *p* is *not* possible I am thereby *prima facie* justified in believing that *p* is possible. And this means that it is safe to assume that (2) of (CA) is only plausible if ‘conceivability’ means ‘strong conceivability’.

Second, strong conceivability is an epistemically more demanding notion than weak conceivability. One can strongly conceive a proposition—that is, have a clear and distinct idea that the proposition is possible—only if one has a sufficiently rich understanding of the proposition in question. But one can weakly conceive a proposition even in the absence of such understanding. Thus, for example, one cannot strongly conceive that there might be creatures with eyes that are not creatures with ears (a Van Cleve example) unless one has a sufficiently rich understanding of what it is to be a creature with eyes and a creature with ears. To take a different example, suppose the ontological argument is sound, and the honest and sincere atheist says that he can conceive of a God-less universe (cf. Yablo 2000). A conceivability report of this kind cannot be

taken as a report of strong conceivability. For, if the ontological argument is sound, Godlessness is self-contradictory—you cannot conceive that it is possible if you understand it.

It is sometimes said that the distinction between strong and weak conceivability is, while real enough, one which is likely to be of considerably less interest once we factor in a different distinction, namely, that between ideal and *prima facie* conceivability (Chalmers 2000a). The idea is that if we consider conceivability in ideal conditions, then negative conceivability *can* be taken to be a guide to possibility, or, anyway, can be taken to be as good a guide as strong conceivability. I want to set this consideration aside, however. For one thing, the idealizations involved in this suggestion are extremely tricky. For another, and more importantly, our argument will go through even if we stick to positive conceivability. Again, I will return to this issue in §VI.

V

So far, I have introduced two conceptions of the physical, and a distinction between strong and weak conceivability. The point to be developed now is that these distinctions provide an answer to the conceivability argument.

The first step is to define, on the basis of our two conceptions of what a physical property is, two conceptions of what it is for a subjects S and S^* to be physically identical. This of course is a simple matter. In general, S is physically identical to S^* iff S and S^* share every physical property. Hence S is t-physically identical to S^* iff S and S^* share every t-physical property; *mutatis mutandis* for o-physical identity. Given our assumption that the class of o-physical properties is distinct from the class of t-physical properties, it follows that the notion of t-physical identity is distinct from o-physical identity.

The next step is to distinguish two versions of both the premises of (CA) and the statement of physicalism with which it operates. The two versions of the premises are:

- (1-t) It is conceivable that I have a t-zombie-twin, i.e. that there is someone who is t-physically identical to me and yet who lacks phenomenal consciousness;
- (1-o) It is conceivable that I have an o-zombie-twin, i.e. that there is someone who is o-physically identical to me and yet who lacks phenomenal consciousness.
- (2-t) If it is conceivable that I have a t-zombie-twin, then it is possible that I have a t-zombie-twin.
- (2-o) If it is conceivable that I have an o-zombie-twin, then it is possible that I have an o-zombie-twin.

And the two versions of physicalism are:

- (3-t) For any subjects S and S^* and worlds W and W^* , if S in W and S^* in W^* are t-physically identical, then they are psychologically identical.

- (3-o) For any subjects S and S^* and worlds W and W^* , if S in W and S^* in W^* are o-physically identical, then they are psychologically identical.

We might say that (3-t) articulates a version of physicalism called t-physicalism, while (3-o) articulates a version of physicalism called o-physicalism.

The problem for the proponent of (CA) can now be stated with some precision. There are two versions of the conceivability argument. The first version takes you from (1-t) and (2-t) to the falsity of (3-t)—let us call this (CA-t). The second version takes you from (1-o) and (2-o) to the falsity of (3-o)—let us call this (CA-o). But neither version is a threat to physicalism.

The reason (CA-t) is not a threat to physicalism is that, even if it is sound, it only establishes that (3-t)—t-physicalism—is false. However, since there are some o-physical properties which are not t-physical, (3-t) can be false while (3-o) is true. It follows that the soundness of (CA-t) does not contradict physicalism, for that argument does not rule out the possibility of o-physicalism. The most (CA-t) shows is that, if one wants to be a physicalist one should endorse o-physicalism rather than t-physicalism, (3-o) and not (3-t). It does not show that one should give up physicalism altogether.¹³

The reason (CA-o) is no threat to physicalism a little more complicated, and depends on the distinction between weak and strong conceivability. Given this distinction we can further differentiate between two readings of the initial premise of (CA-o). These are:

- (1-o-w) It is weakly-conceivable that I have an o-zombie-twin, i.e. that there is someone who is o-physically identical to me and yet who lacks phenomenal consciousness.
- (1-o-s) It is strongly-conceivable that I have an o-zombie-twin, i.e. that there is someone who is o-physically identical to me and yet who lacks phenomenal consciousness.

As we pointed out in §IV, if (2)—the premise of the original conceivability argument taking us from conceivability to possibility—is to have any force, it had better be interpreted in terms of strong, rather than weak, conceivability. What this means is that it is (1-o-s) rather than (1-o-w) which is the relevant premise of (CA-o). In other words, the success or failure of the argument turns on whether it is strongly conceivable that I have an o-zombie twin.

However, this claim can be reasonably denied by the physicalist. The reason, as we also noted in §IV, is that strong conceivability is an epistemically more demanding notion than weak conceivability. One can only strongly conceive of a proposition if one has a sufficiently good understanding of the proposition. On the other hand, on the Kantian picture, I know next to nothing—as Jackson puts it—about the intrinsic physical nature of our world. But if I know next to nothing about the physical nature of my world, then I know next to nothing of *my* physical nature and thus have no genuine understanding of what this nature consists in. However, if I have no understanding of my own o-physical

nature, I cannot truthfully be said to strongly conceive the possibility of something which is o-physically identical to me. But then it cannot truthfully be said that I can strongly conceive of an o-zombie, because an o-zombie is something that is o-physically identical to me *and* who lacks phenomenal consciousness. However—and this is the last step—if I fail to strongly conceive an o-zombie, there is no reason why (1-o-s) cannot be denied by the physicalists. But if (1-o-s) is denied, (CA-o) is no threat to physicalism.

Earlier we saw that (CA-t) is no threat to physicalism because even if it is successful it only refutes one version of physicalism. What we have just seen is that (CA-o) is likewise no threat to physicalism. These two points together constitute our response to the conceivability argument.

Before moving on, I want to pause to introduce an analogy¹⁴ that I hope will make clear the basic shape of the view that I am pursuing. Imagine a mosaic *M* constituted only by two sorts of tiles, triangles and pieces of pie. *M* may of course have many different shapes in it, so long as those shapes are constructed by transparent principles from the basic ones: diamonds, circles, half-moons, and so on. Now imagine that our only epistemological access to *M* is via two shape-detecting systems, a triangle-detecting system and a circle-detecting system. In other words, given our epistemological access to *M*, it would be natural to think that it was constituted *only* by triangles and circles—of course this is a mistake, but it would be a natural one in the situation I am imagining.

Now, in that situation one might well develop a conceivability argument concerning *M* to the effect that circles are something over and above the rest of the mosaic. This argument would proceed from two premises:

- (1-m) It is conceivable that there is a zombie-mosaic-twin, that is, a mosaic exactly like *M* in non-circle respects but which lacks circles.
- (2-m) If it is conceivable that there is a zombie-mosaic-twin, it is possible that there is zombie-mosaic-twin.

On the basis of these two premises one might argue that, since it is possible that there is a zombie-mosaic-twin, circles must be something quite over and above the non-circular aspects of *M*. When God created *M*, he first had to build it out of non-circular respects, and then he had to add circles—circles are (as Descartes might have said) a special creation.

It is obvious that this argument is a bad one. It is palpably *not* possible that there be a zombie-mosaic-twin in the sense we have defined. A mosaic exactly like *M* in non-circular respects would be exactly like it both in terms of triangles *and* pieces of pie—but such a mosaic would *also* be exactly like the actual mosaic in terms of circles. In other words, circles are *not* a special creation of God, and we stand in need of a plausible diagnosis of what has gone wrong in this argument. I think it is plausible to say that in the first premise we are at most *weakly* conceiving of a zombie-mosaic. Since we clearly do not understand central aspects of *M*, we cannot be said to strongly conceive of a zombie mosaic, to have a clear and distinct idea of zombie mosaic. So what is wrong

with the argument is that it is invalid: (1-m) is only true if ‘conceivable’ means weakly conceivable, but (2-m) is only true if ‘conceivable’ means strongly conceivable.

In effect, my proposal is that the same thing goes for (CA-o). In both (CA-o) and the mosaic argument, the first premise is only plausible if conceivability means ‘weak conceivability’: (CA-o) is only plausible if (1-o) is interpreted as (1-o-w). Interpreted in that way, however, no troubling modal conclusion will follow.

V

Our conclusion to this point is that, if the two conceptions of the physical can be kept apart, the conceivability argument fails to refute physicalism. The most that is refuted is t-physicalism. But o-physicalism emerges unscathed. In this section of the paper I want to consider a number of objections to our proposal. In the course of responding to these objections, we will be able to expand on and develop what has gone before.¹⁵

Objection #1: Are o-Zombies only Weakly Conceivable?

The first objection takes up our suggestion that o-zombies are only weakly conceivable. Is this really true? Perhaps one *can* strongly conceive of o-zombies? After all, many people report that they can conceive of the possibility of zombies. Why am I so sure that these reports of conceivability should not be interpreted as reports of strong conceivability of o-zombies?

As I have said, I think it is doubtful that one can strongly conceive of the possibility of an o-zombie. It needs to be appreciated that conceivability reports such as (1-o) are empirical reports about what people can and cannot conceive. But if we interpret those claims as involving strong conceivability, they will have the consequence that people have a rich understanding of their own physical nature. On the other hand, if the Kantian picture or something like it is true, this consequence is false: people who report themselves as conceiving zombies may know next to nothing about it.

However, even if one insists that one *can* strongly conceive of o-zombies, and hence that (1-o-s) is true, our argument will still go through. The central consideration so far has been that ignorance of various aspects of the physical world makes it implausible to suppose that we can strongly conceive that o-zombies are possible. But a similar consideration will serve to show, that even if we can strongly conceive of the possibility of o-zombies, the inference from this to their genuine possibility will fail.

In order to illustrate this, it is useful to have before us a model of how ignorance can defeat the inference from (strong) conceivability to possibility. The best account that I know of is that developed in a number of papers by Stephen Yablo. Yablo writes:

If X finds it conceivable that E, then she is *prima facie* justified in believing that E is possible. That justification is defeated if someone can provide her with reason to suspect the existence of a D such that (i) D is true, (ii) if D is true, then E is impossible, and (iii) that X finds E conceivable is explained by her failing to realize (i) and/or (ii). (2000, p. 121)

On this model of modal error, it is clear that ignorance (“failing to realize”) can sometimes defeat the inference from conceivability to possibility. To see this, look again at the example of the mosaic. Suppose you insist that, in the case I imagined, you *can* strongly conceive a zombie-twin-mosaic, that is, that you can strongly conceive of a mosaic identical to *M* in non-circular respects but which lacks circles. In that case you are *prima facie* justified in believing that the zombie-twin-mosaic is possible. But now suppose you come to suspect the existence of the pieces of pie. Then it seems true both that what you conceived to be possible is in fact not possible, and also that you only conceived it to be possible because you did not know about the pieces of pie. More generally, if you insist that you can strongly conceive of a zombie-twin-mosaic, the pieces of pie will defeat the inference from this to the genuine possibility of such a mosaic. In sum, then, we seem to be faced with two possible suggestions about what has gone wrong in the mosaic argument. On the one hand we can say what we said before, that we can only weakly conceive of a zombie-twin-mosaic, and that in consequence the argument is invalid. On the other hand, we can say that we *do* strongly conceive of zombie-twin-mosaic but that this does not support the inference to the possibility of a zombie-twin-mosaic.

Once again, my suggestion is that parallel remarks apply in the case of the (CA-o). On the one hand, one might say what we said before: that the argument is invalid because the first premise can only be read as involving weak conceivability. On the other hand, one might concede that the first premise involves strong conceivability, but insist that, in that case, the second premise is false. The second premise is false because the inference from conceivability to possibility is subject to defeat. The defeaters are the o-physical properties which are distinct from t-physical. Suppose *D* is the statement which specifies the o-physical properties which are distinct from my t-physical properties. Then *D* will be true of my (putative) zombie-twin as much as me, and will also (if physicalism is true) entail that my zombie-twin does not lack consciousness. From this point of view, our suggestion is that one only finds o-zombies conceivable because we don’t know the nature of the relevant o-physical properties. But if that is so, the second premise of (CA-o) is false, and again the physicalist has an answer to the conceivability argument.

Objection #2: What role does the distinction between the two conceptions of the physical play in this account?

The second objection asks what role the distinction between the two conceptions plays in the account. We have just seen that the idea that we are igno-

rant of the physical world plays an important role in our response to (CA). But one might point out that we are certainly ignorant of many t-physical properties as well. For example, why couldn't the defeater suggested by Yablo's account of modal error simply be a t-physical truth, and not an o-physical truth? Or, again, why couldn't one simply argue against the strong conceivability of t-zombies?

One thing to say about this objection is that it seems to miss an important feature of the dialectical situation. Suppose somebody says that they clearly and distinctly conceive of the Eiffel Tower painted yellow (a Yablo example). It does not follow from this that they have to imagine every *detail* of the Eiffel Tower.¹⁶ It is not to the point to ask whether they imagined every rivet yellow, for example. Nevertheless, it seems clear that they must conceive of the Eiffel Tower as being yellow to *some* degree of detail, otherwise their conceivability claim seems at best forced. Now, how much detail is required? To answer to this question fully would doubtless be extremely difficult, but in general, it seems reasonable to say that one needs to imagine situations in sufficient detail so that we are reasonably confident that further detail will not defeat the initial conceivability claim or the inference from conceivability to possibility. Thus, in the case of the Eiffel Tower, we can imagine it painted yellow not because we can imagine it in every detail, but because we can imagine it in enough detail that we are confident that further detail does not defeat the conceivability claim.

Now, in the case of the conceivability argument the situation is the same. If somebody says that they clearly and distinctly conceive of having a zombie-twin you cannot defeat that by pointing out that they haven't conceived of it down to the last detail. For it doesn't matter that they haven't conceived of it in every detail. All that matters is that they have conceived of it in sufficient detail that they can be confident that further detail will not defeat the claim.

However, once it is clear that we do not need to conceive of every detail of conceivable situations, the role of the two conceptions of the physical in our account becomes clear. If you say that you can conceive of the possibility of a t-zombie, it is certainly true that you do not need to conceive of this possibility in every detail. Nevertheless, in the absence of further information, it is plausible to say that further detail will not defeat the conceivability claim, for further detail will simply be roughly similar to the same sort that we have now. On the other hand, if you say you can conceive of the possibility of an o-zombie, it is open to you to say that further detail will or might defeat the conceivability claim, for in that case the further detail is detail about which we have no knowledge. If it is really true that the sort of knowledge we derive from physical science does not tell us anything about the intrinsic properties of physical objects, then further knowledge of that sort won't help us. In short, the distinction between the o-physical and t-physical properties of physical objects is crucial in our account because while we are certainly ignorant of both, our ignorance of o-physical properties is of a different order, and it is this different sort of ignorance that is in play in our response to (CA).

Of course, I certainly do not want to rule out the possibility that our ignorance of the t-physical properties might be as extreme as, or of the same order

as, our ignorance of the o-physical properties. If that were so, then I certainly think that we could operate with only t-physical properties. But the important point is that, in that case, the essential structure of the account remains the same. We have two conceptions of physical properties, one which introduces a class of properties that are fairly familiar and one which introduces a class of properties some of which we do not understand. And recognizing this fact affords a response to the conceivability argument.

Objection #3: Rejecting the Kantian Picture

The final objection I want to consider focuses more directly on the plausibility of the Kantian picture that we took over from the passage from Jackson. So far we have been presuming that the picture is right, and pursuing its consequences. However, there is a certain frame of mind from which the picture seems preposterous. Why should we believe that there are these strange properties underlying ordinary physical properties? And, even if we do believe this why should we believe that these properties have anything to do with consciousness? (Notice that Jackson makes no such assumption.) Finally, even if there are such properties, and even if they have something to do with consciousness, the proposal seems to have a devastating epistemological consequence: since we are never going to know the nature of these intrinsic properties, we will never be able to give a scientific account of consciousness. In short, the account seems to be simply a metaphysical fantasy, and, worse, a fantasy with disastrous results.

One way to respond to this objection is to confront it directly—defend the metaphysics of the Kantian picture, and downplay its epistemological import. Another response is to bite the bullet—the world is a strange place, one might say, you’ve just got to learn to live with it. I have some sympathy with these responses, but I think that there is a better response available. The better response is to show that the Kantian picture is in the end only the scaffolding for our account. The picture makes very vivid a certain account of what physical objects and our access to them are like, and this account is very useful in the expression and development of our view. But the structure and plausibility of view survives, even if that account is dispensed with.

The best way to illustrate that our account does not require the Kantian picture is to consider another famous conceivability argument: Descartes’ argument about language in section 5 of the *Discourse*. There we find Descartes’ arguing that a certain version of physicalism is false given certain obvious facts about what we would nowadays call linguistic competence—the ability to effortlessly and appropriately speak a language in the way that typical humans do. He writes:

...we can certainly conceive of a machine so constructed that it utters words, and even utters words which correspond to our bodily actions causing a change in its organs (e.g. if you touch it in one spot, it asks what you want of it, if you touch it in another it cries out that you are hurting it, and so on). But it is not conceivable

that such a machine should produce different arrangements of words so as to give an appropriately meaningful answer to whatever is said in its presence, as the dullest of men can do. (AT VI 56: CSM I I40, 1988, p. 44).

There is no doubt a number of different themes lying below the surface of this passage. But in general it seems reasonable to interpret Descartes as mounting a conceivability argument here that is very similar in structure to the one we have been considering from the outset. One may present this argument as proceeding from two premises:

- (6) It is conceivable that I have a machine-twin, that is, something like me in all physical respects but which lacks linguistic competence.
- (7) If it is conceivable that I have a machine-twin, then it is possible that I have a machine-twin.

On the basis of this argument Descartes concludes that physicalism is false. If a machine in the sense defined is possible, competence cannot in the end be a physical matter. However, since we *are* linguistically competent, physicalism is false.

Now, how should one react to Descartes' argument? In my view, the best response is to say that Descartes was right *given his understanding of physicalism*. According to John Cottingham (1992), Descartes' argument in this passage "starts from the observation that a machine, or a *bête machine*, is essentially a stimulus and response device" (p.247) and combines this with the further observation that "the human language-user has the capacity of respond appropriately to an indefinite range of situations, and this capacity seems *toto caelo* different from anything that could be generated by a "look-up tree...correlating inputs and outputs" (p. 248). It seems fair to summarize this by saying that Descartes' argument against physicalism assumes a certain conception of physical identity, namely that, at least as far as language goes, physical identity just is *behavioral identity*, i.e. identity in terms of actual and potential stimulus and response. On this interpretation, in other words, the first premise of Descartes argument appears to be more properly rendered as:

- (6*) It is conceivable that I have a machine-twin, that is, something like me in all (actual and potential) behavioral respects but which lacks linguistic competence.

However, this premise *is* plausible. For the fact is it *does* seem plausible to suppose that someone identical to me in terms of actual and potential stimulus and response might lack linguistic competence. Moreover, *if* one were assuming, as Descartes apparently is, that behavioral identity just is physical identity, then reasoning from this premise does plausibly support Descartes' overall conclusion, namely, that physicalism is false.

Why is the first premise of the Cartesian argument so plausible once it is rendered as (6*)? The answer is that, on this interpretation, the premise looks strikingly similarly to the first premise of arguments which we *know* (or at least

justifiedly believe) to be plausible, namely the refutations of behaviorism suggested by Ned Block (1981, 1998) and others. Block invites us to imagine a machine so designed that it mimics the linguistic dispositions of his Aunt Bubbles, but which operates only on the basis of a giant look-up tree. If the mimicry is clever and complete enough, the machine will be identical to Aunt Bubbles in terms of actual and potential behavior—in short, it will be Aunt Bubbles' machine-twin. But, as Block points out, if the causal transactions which mediate this behavior are of the wrong sort, the Aunt Bubbles' machine-twin will have no linguistic competence at all—"the intelligence of a toaster", as Block famously put it (1981, p. 21). Block convincingly argues on this basis that behaviorism, even in its most sophisticated form, is false. And once again the structure of Block's argument is exactly the same as the conceivability argument we have been considering from the outset: If I can conceive of myself (or Aunt Bubbles) having a machine-twin, then I am *prima facie* justified in believing that this is possible. But if it possible that I have a machine-twin then behaviorism is false. For if behaviorism in the form Block means to discuss it is true, then, necessarily, anything behaviorally identical to me would have to be psychologically identical.

Of course, the difference between Block and Descartes is that Block is aiming only to refute behaviorism, while Descartes is aiming to refute physicalism. Nevertheless, there remains an important sense in which their arguments are the same, and hence that if Block's argument is convincing (and I think we should agree that it is) there is an important sense in which Descartes' argument is convincing also. Given Descartes' assumption that physical identity just is behavioral identity (an assumption obviously denied by Block) his argument has whatever plausibility Block's has.

But of course, even if Descartes' argument in the *Discourse* is successful given his understanding of physicalism, it does not follow that the argument is successful *tout court*. Looking back on Descartes' argument we can see that he was operating with what is from our point of view a greatly impoverished conception of the physical world. From our point of view, what Descartes is missing, among other things, is the idea that internal information processing can assume highly complex computational forms. With that notion in place the whole idea that I can conceive of something physically identical to me but which lacks linguistic competence seems quite unmotivated. Something physically identical to me *would* be identical to me in terms of internal information processing. And, in turn, something identical to me in terms of information processing *would* be identical to me in terms of linguistic competence. In other words, with the monumental advantage of hind-sight, one might suggest that the reason that Descartes' argument seemed plausible to him is only that he lacked certain concepts of the physical world, namely the idea of internal information processing. Moreover, the reason that his argument seems to us fine as a refutation of behaviorism which anticipates Block, but hopeless as a refutation of physicalism is precisely that the concept of internal information processing seems to be a physical notion broadly construed, even if it is not a behaviorist one.

Now, however, we are finally in a position to say why our account does not require the Kantian picture. We need to sharply distinguish two aspects of the Kantian picture that we have so far been running together. On the one hand, there is the claim that, if that picture is true, we are ignorant of certain aspects of the physical world. On the other hand, there is the claim that this ignorance can be recruited to answer the conceivability argument. But for our purposes it is really only this *second* claim that is important. The Kantian picture presents to us a certain vivid model of what our ignorance might be like, but it is really the ignorance, and not the model, that matters. To put it slightly differently, there are a number of different ways of making it plausible that we are ignorant of certain features of the physical world, and that this ignorance can induce conceivability arguments. One way is to appeal to the Kantian picture. But a different way is to appeal to historical cases such as that of Descartes' discussion of linguistic competence. In that case our answer to the conceivability argument is simply to suggest that our position with respect to consciousness is rather like Descartes's position with respect to linguistic competence. Just as he was ignorant with respect to features of the physical world, and just as this ignorance allowed him to formulate a conceivability argument concerning language, so too we are ignorant with respect to features of the physical world, and this allows us to formulate conceivability arguments concerning phenomenal consciousness. But in neither case is the denial of physicalism plausible.

If our proposal can be effectively de-coupled from the Kantian picture, then the final objection to our position can be answered. It remains true on our account that we can distinguish two conceptions of the physical, one which introduces a class of properties of which are ignorant, and another which does not. But it is mistaken to confuse this claim about ignorance with the full-blown Kantian picture. In particular, the claim that we are ignorant with respect to the physical world does not entail that any view about the metaphysics of relational or intrinsic properties, nor that our ignorance concerns or is limited to intrinsic properties, nor does it entail that our ignorance is chronic, in the way that the Kantian picture suggests. Indeed, for all our proposal requires, our ignorance might simply be a function of the present state of science. But of course, it is a mistake to infer from this ignorance that physicalism is false.¹⁷

VI

My argument has been that, if we distinguish between two conceptions of the physical, the physicalist has nothing to fear from the conceivability argument. Operating with the theory-conception of the physical, one might concede that the argument is successful, but on that interpretation, it does not refute physicalism in its most general form. Operating on the object-conception, on the other hand, the argument does not go through. For on the object-conception it is open to the physicalist to say that we are ignorant of various features of the physical world, and this ignorance will block the conceivability argument. But of course, even if what I have suggested is right, this only

shows that our account is one way to answer the conceivability argument. It does not show that it is the best or the only way. In this final section, therefore, I will make some very brief remarks how our proposal fares when compared to the alternatives.

At the heart of the proposal is a certain hypothesis which one might call *the missing concept hypothesis*. This hypothesis says that we are missing certain concepts of the physical world, and that this fact explains our present predicament with respect to consciousness. It is hard of course to defend this hypothesis directly. In the nature of the case, one cannot say what the concepts are, or even what they are like. One must appeal to made-up analogies (such as the mosaic example) or analogies from history of philosophy and science (such as the Kantian picture and the Cartesian argument about linguistic competence). Moreover, the hypothesis is a theoretical one and may of course be wrong. Nevertheless, if it explains the data in way which is less costly than alternatives then it should be preferred. So the crucial question is whether the missing concept hypothesis is more attractive on balance than the alternatives.

Contemporary philosophy seems to present us with three main alternatives to the missing concept hypothesis.¹⁸ One option is eliminativism, according to which there are no experiences as we usually think of them, but there might of course be replacement concepts. Another option is primitivism or dualism, which says that experiences are a primitive or non-supervenient feature of the world. And a third option is a posteriori physicalism—the position I considered briefly in §II. According to this view, statements connecting phenomenal consciousness and the physical are necessary and a posteriori but irreducibly so, in the sense that involves a controversial issue in the epistemology of modality. Demonstrating the superiority of our account over these accounts would require an extensive comparative analysis. Obviously no such analysis can be undertaken here. What I will do, however, is suggest one consideration which favors our alternative over the others.

The consideration looks again the Cartesian argument about language that we considered in the previous section. It is important to see that analogues to all four positions (our own plus the three alternatives) are available as responses to Descartes' argument. A primitivist response would say that linguistic competence is an irreducible fact. (Descartes himself opted for something like this with his postulation of a rational soul.) An eliminativist response says that there is no linguistic competence as we usually conceive of it, and articulates a substitute notion. The a posteriori physicalist would say that the truths about linguistic competence are necessarily connected to truths that involve actual and potential stimulus and responses. Finally, the missing concept view says that we are missing certain concepts of the physical world, and that these concepts will defeat the argument at issue.

Now, in the case of the Cartesian argument, it is obvious in hindsight that the missing concept view is the right view, though of course it is not the view taken by Descartes. Both primitivism and eliminativism are positions which to contemporary philosophers seem obvious overreactions in the case of linguis-

tic competence. And it is not usually thought that an appeal the necessary a posteriori is going to help in this case. Notice, for example, that in discussions of Ned Block's argument against behaviorism, it was never assumed to be open to the behaviorist to appeal to the necessary a posteriori, and say that linguistic competence is only a posteriori related to behavioral dispositions. That would be regarded as an unmotivated departure from usual practice. Since we do not pursue such options when other conceivability arguments are at stake it is not clear why we should pursue such options when considering the argument with which we began.

On other hand, if the missing concept hypothesis is obviously the right answer in the case of the Cartesian argument about language, we are confronted with a powerful reason for supposing that it is likewise the right answer in the case of the conceivability argument that concerns phenomenal consciousness. In fact, the analogies here are rather uncanny. Of course we have a natural tendency to think that, at least in principle we know everything, and that that it is unlikely that new concepts of the physical world will be developed. But I assume that that tendency was strong in Descartes also.

Notes

*A version of this paper was given to the Metaphysics and Consciousness in Philosophy of Mind Conference, Chateau Hotel, Sydney. Thanks to all who took part on that occasion. For comments and discussion about previous drafts, I am also indebted to Andrew Botterell, David Chalmers, Frank Jackson, and Philip Pettit.

1. For the most prominent recent defence of the conceivability argument, see Chalmers 1996. By phenomenal consciousness, I mean the property experiences have when there is something is like to have those experiences. For more on this terminology and for a good discussion of the relation between consciousness in this sense and other senses, see Block 1995. The first premise of the conceivability argument has variations that might be more plausible to some philosophers—for example, one might claim only to conceive someone physically identical to me who is phenomenally *different* (e.g because of inverted or alien experiences). But I will concentrate on the simplest version of the argument here.
2. Strictly speaking physicalism is not quite (3). First, physicalism is usually thought of as a contingent thesis whereas (3) is necessary. Second, physicalism is better stated as a global supervenience thesis; but (3) is not a global supervenience thesis. I will ignore these differences here. For some further discussion, see Jackson 1998 and the references therein.
3. For a very good discussion of the different senses of conceivability, see Van Cleve 1983, Yablo 1993, 2000, and Chalmers 1999.
4. But see Chalmers 2000a for arguments that, at least in some cases, conceivability *does* entail possibility. Here I will stick with the more standard assumption that allows modal error.
5. See Hempel 1970. For a more recent discussion see Crane and Mellor 1990, and Melnyk 1997.
6. The idea that the notion of the physical should be defined via reference to physical theory at a particular stage of its development is mistaken for another reason. Whether

a physical theory is true or not is a function of the contingent facts; but whether a property is physical or not is a not function of the contingent facts. For example, consider medieval impetus physics. Medieval impetus physics is false (though of course it might not have been) and thus it is irrational to suppose it true. Nevertheless, the property of having impetus—the central property that objects have according to impetus physics—is a physical property, and a counterfactual world completely described by impetus physics would be a world in which physicalism is true. But it is hard to see how any of this could be right if physicalism were defined by reference to the physics that we have now or by the physics that happens to be true in our world.

7. For more on this controversy, see Chalmers 1996, Jackson 1998, Stoljar 2000, and in particular the exchange on Chalmers's book in *Philosophy and Phenomenological Research*, Vol. LIX, No. 2, June 1999.
8. The theory-based conception bears some relation to the notion of physical₁ discussed in Feigl 1967; more explicit defense is found in Smart 1978, Lewis 1994, Braddon-Mitchell and Jackson 1996, and Chalmers 1996. For further discussion and defense of this notion of the physical, see Stoljar 2001.
9. The best example of a philosopher who operates with the object-conception of the physical is Feigl 1967; it is also a position that one encounters regularly in discussion. For further discussion and defense of the notion, see Stoljar 2001.
10. For a similar recent development of the themes articulated in this passage from Jackson, see Langton 1998.
11. See, e.g., Chalmers 1996, pp. 153–4, and n.29 on p.375.
12. For further discussion, see Bealer 1999, Chalmers 2000a, Van Cleve 1983, and Yablo 1993.
13. It is important to note that I have not here assumed that (CA-t) is sound, only that, even if it is, it would not defeat physicalism in all its forms. In fact, it is quite possible for us to deny that (CA-t) is sound by taking a leaf out of the book of the a posteriori physicalist—the position I considered briefly in §II. If we suppose that the connection between the t-physical truths and the phenomenal truths is a necessary and a posteriori one, then (CA-t) will not go through. However, this response to the argument does not mean that we need to take up any controversial position in the theory of necessary truths—for the o-physical truths which are not also t-physical will supply the truths from which the a posteriori necessary truths are logically derived.
14. This analogy is also discussed in Stoljar 2001.
15. In Stoljar 2001 I develop a similar strategy as a response to the knowledge argument. In that paper I also consider three further objections to the strategy under discussion. I am not supposing that the objections discussed either here or in that paper amount to all of the objections one might bring against the strategy I am pursuing.
16. This point is well made in Chalmers 2000a.
17. It is interesting to note that this way of stating our proposal makes it vivid that the concepts of the physical world that we are missing need not be concepts of the structure of matter itself. One problem with the Kantian picture is that it encourages the idea that our ignorance is limited to the very small, that is, to matters that pertain the structure of matter itself. To the extent that one might be sceptical that ideas drawn from this level of reality are applicable to consciousness, our answer to the conceivability argument will seem misguided. However, as the example of Des-

cartes makes clear, there is no reason to assume that our ignorance is limited in this way. Notions such as computation and information processing, for example, are notions which apply to the physical world, and are physical notions broadly construed, but they are not about the ultimate structure of matter.

18. One might think that there are many more options than this. However, as far as I can see, most of the apparently different options resolve themselves into one or other of the options discussed in the text. To take one example, a common idea in contemporary philosophy of mind is representationalism, the doctrine that the phenomenal truths supervene on intentional truths. When this position deals explicitly with arguments such as the conceivability argument it usually becomes a version either of eliminativism (as in Harman 1990) or a posteriori physicalism (as in Tye 1999).

References

- Bealer, G. 1998. 'A Theory of Concepts and Concept Possession'. In Villanueva, E. (ed.) *Concepts: Philosophical Issues*, Vol. 9, Ridgeview, pp. 261–301.
- Block, N. 1981. 'Psychologism and Behaviorism', *Philosophical Review*, XC, 1, 1981, pp.5–43.
- Block, N. 1995. 'On A Confusion about a Function of Consciousness', *Behavioural and Brain Sciences*. Reprinted in Block et alia (eds) 1997. *The Nature of Consciousness*. MIT Press. References to the reprinted version.
- Block, N. 1998. 'The Mind as the Software of the Brain'. In Osherson, D et al. (eds.) *Invitation to Cognitive Science Vol 3*. (2nd Ed.), Cambridge, MA: MIT Press.
- Block, N. and Stalnaker, R. 1999. 'Conceptual Analysis, Dualism, and the Explanatory Gap', *Philosophical Review*, Forthcoming.
- Braddon-Mitchell, D and Jackson, F. 1996. *Philosophy of Mind and Cognition*. Blackwell.
- Chalmers, D. 1996 *The Conscious Mind*, Oxford.
- Chalmers, D. 1999. 'Materialism and the Metaphysics of Modality', *Philosophy and Phenomenological Research*, Vol. LIX, No. 2, June 1999 Web version at: <http://www.u.arizona.edu/~chalmers/papers/modality.html>
- Chalmers, D. 2000. 'The Content and Epistemology of Phenomenal Belief'. Forthcoming Web version at: <http://www.u.arizona.edu/~chalmers/papers/belief.html>
- Chalmers, D. 2000a. 'Does Conceivability Entail Possibility'. Forthcoming. Web version at: <http://www.u.arizona.edu/~chalmers/papers/conceivability.html>
- Cottingham, J. 1992. 'Cartesian Dualism: Theology, Metaphysics and Science'. In Cottingham, J. (ed.) 1992 *The Cambridge Companion to Descartes*, Cambridge: Cambridge University Press, pp. 236–257.
- Crane, T. and Mellor, D.H. 1990. 'There is no Question of Physicalism', *Mind* 99:185.
- Descartes, R. 1988. *Selected Philosophical Writings* (ed. J. Cottingham) Cambridge: Cambridge University Press.
- Feigl, H. 1967. 'The "Mental" and the "Physical"' (Minneapolis: University of Minnesota Press. Original Publication: 1958)
- Harman, G. 1990. 'The Intrinsic Quality of Experience', in Tomberlin, J. (ed.) 1990 *Philosophical Perspectives*, Vol. 4, Ridgeview, pp.xx-51. Reprinted in Block et alia, (eds.) 1997 *The Nature of Consciousness: Philosophical Debates*, MIT Press.
- Hempel, C (1970) 'Reduction: Ontological and Linguistic Facets'. In Morgenbesser et al. (eds) *Essays in Honor of Ernest Nagel*. New York: St Martin's Press.
- Hill, Christopher S. and McLaughlin, B. 1999. 'There are Fewer Things in Reality than are Dreamt of in Chalmers' Philosophy', *Philosophy and Phenomenological Research*, forthcoming.
- Jackson, F. 1998. *From Metaphysics to Ethics: A Defense of Conceptual Analysis*, Oxford.
- Kripke, S. 1980. *Naming and Necessity*, Harvard.
- Langdon, R. 1998. *Kantian Humility* Oxford: Clarendon.
- Lewis, D. 1970 'How to Define Theoretical Terms' *Journal of Philosophy* 67:427–46.

- Lewis, D. 1994 'Reduction of Mind'. In Guttenplan, S 1994 (ed) *A Companion to the Philosophy of Mind* Oxford: Blackwell.
- Loar, B. 1997. 'Phenomenal States'. In Block N. et. alia (eds.) 1997. *The Nature of Consciousness*, MIT Press, pp. 597–616.
- Melnick A "How To Keep The 'Physical' in Physicalism", *Journal of Philosophy*, XCIV, (1997), 622–637.
- Searle, J. 1992. *The Rediscovery of the Mind*, Cambridge, MA: MIT Press.
- Smart J.J.C. 'The Content of Physicalism', *Philosophical Quarterly* 28 (1978) pp. 239–41.
- Strawson, G. 1994. *Mental Reality*. Cambridge, MA: MIT Press.
- Stoljar, D 2000. 'Physicalism and the Necessary A Posteriori', *Journal of Philosophy*, XCVII, 1, (January 2000), pp. 33–54.
- Stoljar, D 2001. 'Two Conceptions of the Physical', *Philosophy and Phenomenological Research*, forthcoming.
- Tye, M. 1999. 'Phenomenal Consciousness: The Explanatory Gap as a Cognitive Illusion', *Mind*, Vol.108, 432, October 1999.
- Van Cleve, J. 1983. 'Conceivability and the Cartesian Argument for Dualism', *Pacific Philosophical Quarterly* 64 (1983) 35–45.
- Yablo, S. 1993. 'Is Conceivability a Guide to Possibility?', *Philosophy and Phenomenological Research* 53 (1993), pp. 1–42
- Yablo, S. 2000 'Textbook Kripkeanism & the Open Texture of Concepts'. *Pacific Philosophical Quarterly* 81 (2000), pp. 98–122