

Exclusion Again¹

Karen Bennett
Princeton & ANU

1. Introduction: Two Uses of The Exclusion Argument

I think that there is an awful lot wrong with the exclusion problem. So, it seems, does just about everybody else. But of course everyone disagrees about exactly *what* is wrong with it, and I think there is more to be said about that. So I propose to say a few more words about why the exclusion problem is not really a problem after all—at least, not for the nonreductive physicalist. The genuine *dualist* is still in trouble. Indeed, one of my main points will be that the nonreductive physicalist is in a rather different position vis à vis the exclusion problem than the dualist is. Properly understanding nonreductive physicalism—and clearly recognizing that it is, after all, a form of *physicalism*—goes a long way toward solving the exclusion problem.

Let us begin by reminding ourselves how the exclusion problem is supposed to go. The basic idea is that if everything that happens can be accounted for in purely physical terms, then the mental seems to get shunted to one side and left with nothing to do. More precisely, the problem is that the following five claims are incompatible with each other:

Distinctness: Mental properties (and perhaps events) are distinct from physical properties (or events).²

Completeness: Everything that happens has a sufficient physical cause.³

Efficacy: Mental events sometimes cause physical ones, and sometimes do so in virtue of their mental properties.

¹ I initially drafted a version of this paper for a conference on mental causation put on by the NAMICONA group at Århus University. Many thanks to all the conference participants for very helpful discussion, particularly Helen Beebe, Frank Jackson, and Barry Loewer. Thanks also to the Corridor reading group, constituted on that occasion by Cheryl Chen, Adam Elga, Liz Harman, Sean Kelly, Jim Pryor, and Ted Sider. I presented a more recent version at the AAP meetings and at a workshop on mental causation at Macquarie University. Thanks especially to David Chalmers, Kit Fine, Peter Menzies, Adam Pautz, Denis Robinson, and Daniel Stoljar.

² Different versions of the exclusion problem arise depending upon whether it is type identity, token identity, or both that is denied. I shall be as neutral as possible on this question.

³ Two quick points. First, not much of substance is affected by weakening completeness to the claim that everything that happens has its probability fixed by entirely physical antecedents. Second, it is worth noting that completeness does not say that everything that happens has *only* physical causes. That claim is much stronger, and is for obvious reasons not a good way to start out the exclusion argument (Kim flirts with using it in 2003, 162-164, but rightly decides not to).

Nonoverdetermination: The effects of mental causes are not systematically overdetermined; they are not on a par with the deaths of firing squad victims.⁴
Exclusion: no effect has more than one sufficient cause unless it is overdetermined.

This way of presenting the problem is neutral about which claim should be rejected. It is not always presented like this—as a package of inconsistent claims, some one of which has to give, and any one of which will, in principle, do. It is sometimes instead presented as an argument against Distinctness. Framed as an *argument*, then, it is supposed to show that the mental is not distinct from the physical after all.

But that conclusion is ambiguous, and so is the role of the exclusion argument in the philosophy of mind literature. The argument gets used for two rather different purposes. Sometimes it is used as an argument for physicalism, as against property or substance dualism (see, e.g., Papineau 1995, 2001).⁵ And sometimes it is used as an argument for *reductive* physicalism, as against *nonreductive* physicalism (see, e.g., Kim 1989, 1993, 1998). That is, sometimes it is used to defend physicalism, and sometimes to defend a *version* of physicalism. These two uses blur into each other, and it is not always obvious which a given author is doing. But they need to be carefully distinguished. As a proud card-carrying physicalist, I heartily approve of the former role. But as a proud card-carrying *nonreductive* physicalist, I think the latter needs to be rejected.

We nonreductive physicalists should try a lot harder than we have to claim that the exclusion problem is not a problem for us, *while nonetheless maintaining that it is a real problem for the nonphysicalist*. That is, we should do our best to deny that we are in the same boat as emergentists vis à vis the exclusion problem and the commitment to so-called ‘downward

⁴ It is purely a terminological matter whether this is formulated as stating that the effects of mental causes are not overdetermined at all, or as stating that they are not overdetermined in some particularly *bad* way. The important point is the insistence on the disanalogy.

⁵ Sometimes people instead appeal to Ockham’s razor to get from the completeness of physics to physicalism. But this is really just a version of the exclusion problem. After all, the idea would be to say that the fact that the physical

causation' (*pace* Crane 2001, Kim e.g. 1989, 1993). We should do our best to deny that the exclusion argument is a good argument for reduction, while nonetheless insisting that it *is* a good argument for the claim that the mental is nothing over and above the physical. After all, actual arguments for physicalism are rather hard to come by, and we should not throw the baby out with the bathwater. Defending nonreductive physicalism should not require defending full-blown dualism, too.

It is therefore clear what the nonreductive physicalist's goal should be—find a solution to the exclusion problem that is only available to a physicalist. This immediately entails that she should not 'solve' the exclusion problem by endorsing epiphenomenalism or claiming that systematic overdetermination of the everyday firing squad sort is perfectly fine. Neither move is particularly palatable anyway, but what matters at the moment is that both of them are *equally available to the dualist*. There is nothing physicalist in them; people like Descartes and Chalmers could endorse them too. Neither denying Efficacy nor denying Nonoverdetermination can achieve the nonreductive physicalist's goal.

That leaves Distinctness, Completeness, and Exclusion. This is more than enough to do the job. I shall argue that certain ways of putting pressure on those three claims opens up two promising routes for the nonreductive physicalist—two responses to the exclusion problem that are not available to the emergentist or substance dualist. One of these responses turns on denying Exclusion; it is a strategy I call 'compatibilism', and have defended in more detail elsewhere (2003). The other response is new, and locates the mistake at the intersection of Distinctness and Completeness rather than with Exclusion. On this line, the exclusion problem simply does not arise for physicalists—at least not in the same form as for dualists. At the end of

is doing all the causal work renders the mental unnecessary and *extra*. But the further claims of Nonoverdetermination, Efficacy, and Exclusion are needed to render that argument formally valid.

the day, this move needs to be supplemented with compatibilism to defuse the version of the problem that remains. Nonetheless, it is important, because it forces us to get straight about what, if anything, the problem really is. Doing so requires gaining a proper understanding of nonreductive physicalism—an understanding that occasionally seems to have been lost.

The rest of the paper divides roughly into two halves. In the first half—sections 2, 3, and 4—I will try to clarify nonreductive physicalism, and sketch how properly understanding the view affects the way the exclusion problem looks. It is not obvious that the combination of Distinctness and Completeness leads to any trouble at all for a nonreductive physicalist. In the second half—sections 5, 6, and 7—I will turn my attention to compatibilism, the claim that it is Exclusion that is false. This view has been motivated in two different ways. However, I will argue that only one of them is successful, and that it is only open to a physicalist. In doing so, I will argue that, *pace* occasional claims to the contrary, the force of the exclusion problem does not rest on any particular account of causation.

The upshot will be that there are at least two types of solution that are only available to a physicalist. Both are promising and, I think, much more plausible than the moves a dualist can make. I shall not really argue for that latter claim here. But my overall claim is nonetheless that we nonreductive physicalists can indeed meet our goal. We can claim, *contra* Kim, that the exclusion argument provides no reason at all to return to reduction, while leaving plenty of room to claim that it very much does provide overwhelming reason to be a physicalist. The alleged problem looks rather different for a physicalist than for a dualist.

2. Nonreductive Physicalism Redux

The first item of business, then, is to get a firm grip on what nonreductive physicalism does and does not say. Let us begin with the ‘physicalism’ part. It is notoriously hard to define it adequately, but I can at least offer up the same slogans as everyone else. Physicalists not only endorse the completeness of physics, but also think that all the facts are physical facts—that there is nothing ‘over and above’ the physical. Physicalists believe that everything globally supervenes⁶ on the physical as a matter of *metaphysical* necessity (see Lewis 1983, Chalmers 1996, and Jackson 1998; see Hawthorne 2002 for interesting challenges to their definitions). That is, physicalists deny that it is merely *nomologically* impossible for there to be a world physically just like this one but mentally different. There are no special psychophysical laws that *link* or *tether* the mental to the physical, and that can be broken. The connections between the physical and the mental are a lot more analogous to the connections between *being cerulean* and *being blue*, or between *being blue* and *being blue or green* than they are to the connections between fire and smoke, or between *being water at 100° C at sea level* and *being boiling*.⁷ Physicalists think that mental events and properties are not truly distinct existences that can be snipped away from their physical bases; the connecting laws simply are not breakable. There is no room for any wedge. That is why the metaphysical necessity of the supervenience claim is of crucial importance to their view.⁸

⁶ I invoke global supervenience because it is both standard and convenient. As I have argued elsewhere (2004), however, any claim made with global supervenience can also be formulated in terms of strong supervenience.

⁷ I have only chosen these examples because they are particularly clear. Nothing turns on the fact that they are cases of *a priori* entailment; I am officially neutral on whether or not the physical’s guarantee of the mental is knowable *a priori*.

⁸ It has been suggested that supervenience with metaphysical necessity, though required, is not enough for physicalism. However, I am not convinced that there is any real reason to think that supervenience with metaphysical necessity is compatible with dualism. For example, Andrew Melnyk argues that if everything supervenes with metaphysical necessity on the physical, there must be some further explanation of why that relation of covariance holds. Melnyk appeals to the relation of ‘realization’ here (2003). This may well be right. However, it does not quite show that metaphysically necessary supervenience is not sufficient for physicalism—after all, metaphysically necessary supervenience might guarantee realization, which in turn is sufficient for physicalism. It at best shows that there are more informative *characterizations* of physicalism to be had.

Jessica Wilson (forthcoming) argues that necessitarianism about the laws of nature (Shoemaker 1980, Swoyer 1982) collapses the distinction between nomological and metaphysical necessity, and that this means that supervenience with metaphysical necessity is not sufficient for physicalism. But not only is the Shoemakerian

Compare, for example, the natural reaction of the person who believes in ordinary composite objects like tables when faced with the nihilist claim that there are no tables, only simples arranged tablewise (c.f. Unger 1979, van Inwagen 1990, Merricks 2001). The natural reaction is to say that there is no room for any wedge. Once you buy the simples arranged tablewise, you *ipso facto* buy the table. This position is similar to that of the physicalist, who thinks that once you buy the existence and distribution of instantiated physical properties, you *ipso facto* buy the mental ones as well. (Of course, I have not argued that the natural reaction is the *correct* reaction, but then I have not argued that physicalism is true either. I am at this stage merely analogizing the positions.)

Next, *nonreductive* physicalism. We nonreductive physicalists accept everything in the above paragraph, but reject certain sorts of identity claims. However, it is important to be clear about just *what* sorts. Nonreductive physicalism is often characterized as the conjunction of physicalism with the rejection of type identity claims, but that is not quite accurate. For one thing—which isn't important for present purposes—there are plenty of people who consider themselves physicalists but reject *token* identity claims as well.⁹ For another thing—which *is* important for present purposes—nonreductive physicalists simply do not reject all type identities. We only reject *some*.

premise controversial, it is also not clear why it would entail that supervenience with metaphysical necessity is not sufficient for physicalism. It at best supports the claim that the modal force of the postulated supervenience relation cannot be the central difference between dualism and physicalism. Arguably, that is, the proper conclusion to draw is that property dualism with nomological supervenience (à la Chalmers 1996) is not a coherent view.

Adam Pautz has tried to convince me that there is room for a form of dualism according to which the mental is both genuinely distinct from and metaphysically necessitated by the physical. I must confess that I simply do not understand such a view. I cannot understand what makes it properly *dualist*, because I cannot understand how genuinely distinct mental properties (let alone full blown *souls*!) could be *metaphysically necessitated* by the physical. Like Hume, I cannot understand how there can be metaphysically necessary connections between fully distinct existences. Clearly, the emphasis is on 'fully' and 'genuinely' distinct; part of the point here is that dualism requires more than an insistence on numerical distinctness (c.f. Stoljar forthcoming).

In what follows, I shall continue to take the metaphysically necessary supervenience of the mental on the physical as sufficient for physicalism.

⁹ They do so because they think mental and physical events differ modally, like statues and lumps of clay allegedly do. Many such people nonetheless consider themselves physicalists. See, e.g., Burge 1979, Boyd 1980, Pereboom and Kornblith 1991, Yablo 1992.

Nonreductive physicalists do not think—or, at any rate, *should* not think—that mental events and properties really are not identical to any physical ones. All we think is that they are not identical to any *standard* physical ones. We have no reason to deny that they are identical to physical events and properties reachable by extension or analogy with standard ones. Let me try to put this marginally more carefully, by loosely distinguishing between a narrow and a broad sense of ‘physical’.

Say that a property, object, or event-type counts as physical in the narrow sense just in case it is invoked in the laws and generalizations of a clearly physical science like physics, chemistry, or neuroscience. (Or maybe even just physics, though for present purposes I am content to count that as physical in a super-narrow sense.) Narrowly physical events and objects tend to be very small—hence the popularity of the term ‘microphysical’—but I leave this officially open. Now, there are of course lots and lots of well-known difficulties coming up with any adequate definition of ‘physical’ (e.g. Hempel 1980, Crane and Mellor 1990, etc.), and I don’t pretend to have a new or interesting response to them. Just pick your favorite answer, and plug it in here. I am merely gesturing at an answer of that sort—that’s the idea of ‘physical’ in the narrow sense.

In contrast, an event, object, or property counts as physical in the broad sense just in case it is constructible out of narrowly physical ones in a clearly articulatable way. So, for example, broadly physical objects are constructed from narrowly physical ones by mereological operations and other forms of composition.¹⁰ Broadly physical properties are those constructed from narrowly physical ones by means of property-forming operations like disjunction, conjunction, and quantification (though presumably not negation!). Broadly physical events are those

¹⁰ The “and other forms of composition” is crucial; I am not saying that ordinary middle-sized objects are mereological fusions of narrowly physical objects.

constructed from narrowly physical ones by means of various forms of spatio-temporal, mereological, and modal gerrymandering. And so on. This list is merely supposed to give the general idea, and presumably needs to be expanded and tweaked in various ways.

But let me nonetheless immediately deflect two potential misunderstandings about the operations I have mentioned. First, it is no part of my view that ordinary middle sized objects and events are mereological fusions of narrowly physical ones. Composition as defined by classical mereology is *a* form of composition, but it may not be the only form; that is why I said “mereological operations and other forms of composition.” Second, notice that I have not said that the set of broadly physical properties, events, or objects is the result of *closing* the set of narrowly physical ones under the property-, object-, and event-forming operations I have listed. My claim is weaker than that. My claim is perfectly compatible with the view that only *some* disjunctions of narrowly physical properties count as broadly physical properties, and with the view that only *some* mereological sums of narrowly physical objects count as broadly physical objects, and so forth. My claim is just that *if* a particular disjunction of narrowly physical properties (or mereological sum of narrowly physical objects) counts as a property (object) *at all*, it counts as a broadly physical property (object). There are two separate questions in the vicinity—one about which properties and events *exist*, and another about which of the existing properties and events count as physical. I am only addressing the latter question at the moment; my answer is neutral on the former question.

The property *being an electron or a blender* is a broadly physical property, if it exists at all. Indeed, so is *being a blender*. So are functional properties of the form *having some narrowly physical property or other that plays the pain role*. A leaf is a broadly physical object, and a particular leaf’s falling is a broadly physical event. Whatever exactly the compositional

principles needed to get from the behavior of individual molecules to the movements of a *leaf*, those principles generate a broadly physical event from a lot of narrowly physical ones. Broadly physical events also include things like the mereological sum of a particular leaf's falling and a particular bird chirping, as well an event otherwise just like that perfectly ordinary leaf falling except for the fact that it essentially takes place at the precise moment it actually does.

Here is the crucial point. The nonreductive physicalist is only committed to denying that mental events and properties are physical in the *narrow* sense. She denies that mental properties like *being in pain* are identical to any neurophysiological property, or any other property that gets picked out by a narrowly physical predicate. *Pain* is not identical to *C-fibers firing*. But she needn't, and shouldn't, deny that pain is a physical property in the broad sense. Mental events and properties do not count as physical in exactly the same sense, and for exactly the same reason, that properties like *being either an electron or a blender*, and events like a modally fragile leaf-falling do not count as physical. They just aren't *narrowly* physical. If I may quote from a somewhat unlikely ally: "physicalism need not be, and should not be, identified with microphysicalism" (Kim 1998, 117).

Slogans like "mental properties aren't identical to physical ones" might make for nice bumper stickers, but they make for a misleading characterization of nonreductive physicalism. Nonreductive physicalists need to be careful not to get carried away by their own rhetoric. Mental properties *are* physical properties—just not the ones invoked in the laws of physics, chemistry, or neuroscience. That is all the standard arguments against reduction purport to show. Nonreductive physicalism, whether about events, or properties, or both, is primarily a claim about classification and individuation. The properties and events picked out by mental terms *cross-classify* those picked out by narrowly physical terms. That is the only lesson of multiple

realization, and of all the modal arguments for the distinctness of mental and physical event tokens. When nonreductive physicalists deny that the mental is physical, they are saying something much, *much* weaker than dualists are.

Perhaps that is obvious. But perhaps it will help to lay out what this way of thinking about nonreductive physicalism says about what the interesting questions about reduction are, and about which battles are worth fighting. It is an interesting question, one that Kim has pressed (1992, 1998), whether mental properties in fact survive multiple realization. That is, it is an interesting question whether there in fact are any overarching mental properties, like *being a belief*, or *being a pain*, or whether there really only are more narrowly circumscribed, species- or structure-specific properties like *pain-in-humans*, *pain-in-dogs*, etc. And, if we decide that the overarching, second order properties do exist, then there is an interesting semantic question—one that Lewis has pressed (see xx, and especially 2000b)—whether our mental kind terms pick out the overarching second-order properties or their first-order realizers. And it is obviously an interesting question whether physicalism is true at all. But *assuming* physicalism, and *assuming* willingness to countenance non-structure-specific mental properties like *being in pain*, it is just *not* an interesting further question whether those properties count as physical in any sense. All I'm saying here is that it would not be an interesting victory for the reductionist if pain turns out to be a disjunction of narrowly physical properties, or a second-order property defined in terms of narrowly physical ones. That is not the right battleground for the fight about reduction.

3. Equivocation?

This is an important point in its own right. It is a mistake to assume, as an earlier self of mine did, that nonreductive physicalism would be falsified if we could identify mental events

and properties with *any* physical ones, no matter how disjunctive, extrinsic, and otherwise complex (2003, 486-487).¹¹ And it is a mistake to slide, as Thomas Crisp and Ted Warfield do, from the claim that “the nonreductive physicalist is committed to the idea... that mental properties are not identical with physical properties” to the claim that “the nonreductive physicalist denies the reductivist claim that mental properties are ‘nothing over and above’ physical properties” (2001, 305). And it is a mistake to say, as Kim recently has, that the nonreductive physicalist is committed to denying the claim that “mind-body supervenience is logically or metaphysically necessary, since such a view is essentially a reductionist view” (2003, 162).

But the point is even more important in the context of the exclusion argument, because there is a case to be made for the claim that the argument *equivocates*. It is not at all obvious that the sense of ‘physical’ invoked in Completeness is the same as the sense of ‘physical’ invoked in Distinctness.¹² As I’ve just been emphasizing, we nonreductive physicalists think that mental properties and/or events are distinct from physical ones only in the *narrow* and *microphysical* sense. But there is a case to be made that the correct version of Completeness only says that everything that happens has a sufficient physical cause in the *broad* sense. If that’s right, the nonreductive physicalist can perfectly well claim that some events have causes that are physical in the weaker, broader sense, but not in the stronger, narrower sense. That is, there would be room to claim that some events only have mental, or otherwise higher-level physical causes. Completeness would not guarantee that the effects of mental causes always have another physical cause lurking nearby.

¹¹ I.e., it is a mistake to worry, as I once did, about whether my version of compatibilism—to be discussed later—is a poison pill that undercuts nonreductive physicalism.

¹² Homework assignment: compare and contrast Sturgeon 1998, who claims that the argument equivocates on the sense of ‘physical’ used in Completeness and the sense of ‘physical’ used in Efficacy. He claims that the sense in

Now, I'm not sure that it *is* right. I'm not by any means sure what set of properties Completeness is true of. (I rather doubt the issue can be settled from my armchair.) However, I'm quite sure that we don't need to include special ectoplasmic properties, and I'm equally sure that we must include more than the intrinsic properties of quarks. We must at least include relations among them; we must at least supplement the set with configurational properties. And once we count configurational properties and the like as physical, the category expands dramatically, and it is no longer clear that a nonreductivist, as opposed to a genuine emergentist, would deny that mental properties are physical in the relevant sense—the sense invoked in Completeness.

So I am not going to claim that I have a knockdown argument for the equivocation. But I certainly do not think it is obvious that Completeness and Distinctness rely on the same notion of the physical. Thus I hereby put it forth as a challenge to those who want to push the exclusion argument against nonreductive physicalists—find a notion of the physical such that

- a) the physical is causally complete, and
- b) typical *nonreductive physicalists* will deny that the mental is physical.

The proponent of the exclusion argument needs to tell us what he means by 'physical', and convince us that there is no mismatch. As matters stand, it is not at all obvious that the standard version of the exclusion problem really gets off the ground against a nonreductive physicalist. To assume without argument that it works exactly the same as it does against a dualist requires mistakenly assuming that the dualist and the nonreductive physicalist mean the same thing when they deny that mental properties are physical properties.

I have already quoted Kim as a somewhat unlikely ally. It is worth mentioning in passing that the point I am making here turns out to have a lot in common with a view that Kim

which mental causes can have physical effects is different from the sense in which every physical thing that happens

expresses in the last few pages of *Mind in a Physical World*, though he does not characterize the point as an equivocation.¹³ Kim broadens the notion of the physical to include various macrophysical and functional properties, partly to accommodate intuition, and partly to ensure that Completeness comes out true (1998, 112-115). He consequently winds up admitting that the only mental properties that are threatened by the exclusion argument are those that “resist functionalization” (116), and thus that the only people who are in trouble with the exclusion argument are those who think that qualia are nonfunctionalizable. Now, Kim puts more emphasis on functionalizability than I do, but the basic picture is the same. Only mental properties that are not in *any* sense physical—not functionalizable, not broadly physical, not metaphysically necessitated by particular narrowly physical properties, etc.—are threatened by the exclusion argument.

Now, it may look rather odd to push the exclusion problem against the nonreductive physicalist as hard as Kim has over the past 15 years, and then to claim in the end that it is really only a problem for those who “embrace a serious form of dualism” (119). Then again, Kim clearly does not use ‘nonreductive physicalism’ in the way I have been arguing we should use it. I take it that I am, in his sense, a reductivist. Indeed, he seems to blur nonreductive physicalism into dualism. Consequently, it is not really all that surprising that he both claims that the exclusion problem is a problem for nonreductive physicalists, and claims that it is only a problem for dualists.

4. But the Problem Remains

has a sufficient physical cause.

¹³ Thanks to Carsten Hansen for pointing this out to me.

But even if all that is right, it would be a mistake to blithely assume that the nonreductive physicalist is home free. Suppose the exclusion argument does indeed equivocate, and thus that the mismatch between Distinctness and Completeness means that the standard version of the exclusion problem does not quite stick. Nonetheless, alternate versions will likely arise. After all, the very most that we can get from the above discussion is that *Completeness* does not guarantee that each effect of a mental cause always has another, somehow more decidedly physical cause as well. But it is still open that some *other* principle guarantees it. In particular, mental events may compete with more or less coincident physical events with different essences, and mental properties may compete with their more narrowly physical realizers. To borrow an example of Stephen Yablo's (1992), consider a pigeon who is trained to peck at scarlet buttons. But of course every particular button that is scarlet is also *red* (not to mention colored and spatially extended). So at least on the face of it, there might be a question about whether it is the scarletness or the redness that causes the pigeon to peck in any particular case. Yet it isn't *Completeness* that guarantees that the instantiation of scarletness is always accompanied by another potential cause; it's just the instantiation of scarletness itself.

The analogy is not perfect, in perhaps telling ways, but the point should be clear enough. The nonreductive physicalist does accept that the mental is physical, but she will usually *deny* that the mental event or property mentioned in some particular example is identical to the physical one mentioned *in that example*. And as long as both are apparently causes, and as long as she denies the identity between them, worries about causal competition are off and running. So even though there is definitely something funny about invoking *Completeness* here, given that the *mental is itself physical*, some version of the standard problem surely remains. However, if anything in the ballpark of my worries about *Completeness* is right, it is not quite the *same*

version. And whatever happens, all apparent competition has been reduced to competition among *physical* events and properties.

Thus far, then, we have seen that once we are careful about what nonreductive physicalism amounts to, the exclusion problem starts looking rather different for a physicalist and a nonphysicalist. But we are not yet done; the problem has not totally evaporated. So let us move on to the second half of the paper. I shall now turn away from focusing on Distinctness and Completeness, and take a look at the strategy of denying Exclusion. As I shall argue, this is a particularly physicalist maneuver too.

5. Denying the Exclusion Principle: Compatibilism

Exclusion, recall, says that no effect has more than one sufficient cause unless it is overdetermined. The idea behind the name is just that the existence of one sufficient cause somehow blocks or excludes all others, except in the presumably rare case of overdetermination. To deny this claim is to insist that there is more to overdetermination than that—to claim that in certain circumstances effects can have several sufficient causes and nonetheless not count as overdetermined in the paradigmatic firing squad sort of way.¹⁴ The resulting view is both popular and plausible. It is the view that I have called Compatibilism, and defended in more detail elsewhere (2003).

As I argued there, however, we cannot simply *proclaim* Exclusion false; we have to at least make a modicum of an effort to say *why* it is false. There are two strategies for doing so, based on two (compatible) diagnoses of what overdetermination requires. To see this, think

¹⁴ Again, it does not matter whether I say that sometimes effects can have more than one sufficient cause without being overdetermined, or whether I say that sometimes effects can have more than one sufficient cause without being overdetermined in some particularly bad way (see note 4). It is a purely terminological matter whether I say that not all double causing is overdetermination, or whether I say that all double causing *is* overdetermination, but not all overdetermination is on a par.

about some paradigm cases of overdetermination: firing squads, two children (presumably named Billy and Suzy) simultaneously throwing rocks through a window, your alarm clock sounding at just the same moment as a jackhammer starts up outside your window. Clear cases like these seem to involve a) two completely distinct b) transfers of energy. That is, there two separable components in there—that there are *two* causes, and that the causes are *oomphy* in some sense that I will try to say a bit more about later. A compatibilist can focus on the latter, and reject the oomphy notion of causation in favor of something along the lines of a pure counterfactual dependence notion. Or he can focus on the former, and claim that in the relevant cases the requirement of twoness is not quite met—i.e., he can claim that in the relevant cases the causes are tightly related in some way that defuses the threat of overdetermination.

So one strategy for denying the exclusion principle focuses on the notion of causation in play, and the other focuses on the relation between the causes. Only one of them works, and only one of them is available to a dualist. Unfortunately, it's not the same one. The one that works is the strategy that claims that certain kinds of tightly related causes can both be causally sufficient for the same effect without overdetermining it. In section 6, I will briefly sketch my own preferred version of this strategy, and argue that no dualist can avail himself of it. In section 7, I will argue that although the dualist certainly *can* move to a pure dependence notion of causation, doing so does not defeat the Exclusion principle. Doing so consequently cannot itself ground compatibilism.

6. Compatibilism I: Tight Relations Between the Causes

The thought that the mental and physical are so tightly related that they do not overdetermine their effects is a natural one, and I am far from the first person to have it (see in particular Yablo 1992 and Shoemaker 2001). But unlike some people who push this line, I think

that we both must and can say more about just *why* certain kinds of tight relation moot the threat of overdetermination. I think it is unsatisfactory to say, while emphasizing one's nonreductivism, that mental event or property *m* is not identical to narrowly physical event or property *p*, and then to say in practically the same breath that of course *m* and *p* do not causally compete in any way. If they are distinct, the threat of competition must be *argued against*. We cannot just *assert* that we can have it both ways. We nonreductive physicalists must properly shoulder the burden of proof and say *why* these intimately-related-but-*distinct* causes do not overdetermine their effects. This is what I tried to do in my earlier paper (2003).¹⁵

What I claimed is that overdetermination requires the nonvacuous truth of certain counterfactuals. In order for two causes, *m* and *p*, to overdetermine some effect *e*, it must be nonvacuously true that

- (O1) if *m* had happened without *p*, *e* would still have happened: $(m \ \& \ \sim p) \ \square \square \ e$, and
(O2) if *p* had happened without *m*, *e* would still have happened: $(p \ \& \ \sim m) \ \square \square \ e$.¹⁶

This is not meant to involve commitment to a counterfactual *analysis* of causation; it is simply supposed to reflect our everyday reasoning about causation and overdetermination. Take any case you like. If only one of the putative causes really was a cause, only one of the counterfactuals will be true. If they were joint causes, both of the counterfactuals will be false. If *m* and *p* are in fact the very same event, both counterfactuals will be vacuous. And if *e* really is overdetermined by *m* and *p*—think firing squads, Billy and Suzy throwing rocks at a window, etc—both counterfactuals will be nonvacuously true.

¹⁵ If I may take a moment to situate myself with respect to some other compatibilisms on the market... I take it that in trying to say more about *why* certain pairs of causes do not overdetermine their effects, I am addressing a question that Stephen Yablo (1992) does not try to address. With some minor tweaking, he can more or less take my view on board if he likes. In contrast, Sydney Shoemaker *is* addressing the same question as me (2001). However, his recent twist on Yablo's approach requires commitment to a variety of substantive metaphysical views. My approach is much more neutral on questions about the nature of properties, causal powers, and the like. Nonetheless, all three of us are compatibilists; our views are in the same rough vicinity.

If I am right that the nonvacuous truth of these counterfactuals is necessary for overdetermination, the next step is clear.¹⁶ The question is whether the dualist, the physicalist, or both, can claim that one of the counterfactuals is not nonvacuously true—is either vacuous or false. Messing with (O1) is tricky, and I shall not get into it here, except to say that it definitely does not look vacuous, and it is hard to see how it could be false without undermining the assumed causal efficacy of the mental. So the real question is about the status of (O2). I've argued elsewhere (2003) that the *physicalist* gets to say that (O2) will come out either false or vacuous in all cases of mental causation, depending on what sort of physical events or properties she takes to be causally sufficient for the effect in question. What I would like to argue now is that the same is not true of the dualist.

Let's start with the easy bit. It is clear that only the physicalist can say that (O2) *ever* comes out vacuous. The dualist can't, because he does not think that there are any physical events or properties that metaphysically necessitate mental ones. He precisely thinks that are—at best!—contingent psychophysical laws that link the two. So the dualist denies that there is anything to substitute for *p* that would make the antecedent metaphysically impossible. He at

¹⁶ This test is supposed to be fully general; I only name the causes 'm' and 'p' in order to streamline the ensuing discussion. Also, the counterfactuals can be tweaked in various ways to account for the fact that a version of the exclusion problem can be run on mental *properties*.

¹⁷ Martin Jones has raised the following counterexample to my claim that the nonvacuous truth of the counterfactuals is necessary for overdetermination. Suppose that there is a small firing squad of just two shooters, Billy and Suzy, with their weapons trained upon the victim. Suppose further that Billy has a different kind of gun, and his bullet will travel somewhat faster than Suzy's (or, alternatively, suppose that they use the same type of gun, but Billy is standing closer to the victim than Suzy is.) Billy is a sensitive chap, however, and wants to avoid being the only person to shoot the victim. So he waits a split second to make sure that Suzy has fired her gun, and only then fires his. If Suzy does not fire when the command is given, Billy fires into the air. But if Suzy does fire, he aims properly and fires at the victim too. Because his bullet travels faster, it makes up the time en route and strikes the victim at the exact same moment as Suzy's does. This certainly looks like a case of overdetermination—after all, the victim gets shot with two bullets! Yet one of the two counterfactuals is (non-backtrackingly) false. Had Billy fired his gun and Suzy not fired hers, the victim would not have died. So it looks like this is a counterexample to the claim that overdetermination requires that both counterfactuals be nonvacuously true.

There are several possible responses to this kind of 'staggered' overdetermination case. One is to insist that the relevant event here is Billy's firing *at the victim*, not his firing full stop. The overdetermination counterfactuals are nonvacuously true for that choice of cause. Another, which is more neutral about the individuation of events, is to grant that the death is not overdetermined by Billy and Suzy's firings, and to claim that it is instead overdetermined by some *intermediate* pair of events for which the counterfactuals are nonvacuously true. Billy and

most thinks that there are choices of p that would make the antecedent *nomologically* impossible. So the dualist cannot claim that any instance of (O2) is vacuous.

Can he claim that it is false, despite thinking that p is causally sufficient for the effect? I don't see how. Let me sketch the story that I've told about how the *physicalist* can say that it is false, and then explain why the dualist cannot say the same thing. First, the physicalist certainly can think that there are nonvacuous instances of (O2). That is, she can think that some physical event p is causally sufficient for effect e , that some mental event m is as well, and that p fails to necessitate m . After all, most of the events and properties that we talk about when we talk about the exclusion problem—things like patterns of neural activity, or properties like *being a C-fiber firing*—don't necessitate anything mental. These ordinary, everyday events and properties tend to be spatio-temporally localized (more-or-less intrinsic) and they only guarantee the existence of the mental ones they 'realize' *given certain background conditions*. For example, it is perfectly possible for C-fiber firings to occur without pain. They could be hooked up rather differently, or not hooked up to anything at all. Context matters. It is not an accident that physicalism is usually characterized by means of a *global* supervenience thesis rather than a local one.

There is, in short, an important mismatch between the sorts of physical properties and events that are typically invoked in instances of the exclusion problem, and those that constitute the supervenience base for the mental. It is only hugely complicated extrinsic physical properties, and physical events with complicated extrinsic essences, that will metaphysically guarantee mental ones. So as long as it is legitimate to plug the more intrinsic, everyday physical events and properties into the counterfactual (O2), it won't come out vacuous. Whether it is legitimate to do so depends on whether such things are ever causally sufficient for anything,

Suzy's firings count as overdetermining the death in a slightly derivative sense, because they cause events that nonderivatively overdetermine it.

which in turn depends upon one's views about the nature of causal sufficiency. Basically, if you think that causal sufficiency is a kind of strict sufficiency, according to which only big complicated sums of everyday events, background conditions, causal intermediaries and the like count as causally sufficient for anything, then the only physical occurrences that will ever be causally sufficient for action will be the complicated nonlocal occurrences. These do guarantee the mental ones, and thus all instances of (O2) come out vacuous. If, on the other hand, you think that causal sufficiency is mere sufficiency in the circumstances, then the more ordinary, localized physical events and properties will count as causally sufficient for action, and not all instances of (O2) will be vacuous.

So even the physicalist can perfectly well say that there are nonvacuous instances of (O2). These are claims like the following:

had these C-fibers fired occurred without the pain, my hand would still have jerked back from the stove.

However, these claims are typically *false*—by physicalist lights, anyway. The idea here is simple. The context within which the physical event or property guarantees the mental one *is the same as the background conditions within which it brings about its effects*. So the C-fibers can perfectly well fire without the pain. They could be wired up differently, or perhaps twitching away in a petrie dish. But in such a situation, they will not at all cause the sorts of things they actually cause—they will not cause me to pull my hand away from the stove, and they will not cause me to jump around swearing like a sailor. For localized choices of p , then, p can indeed happen without m , but if it did, there is no reason to expect the occurrence of e . Those instances of (O2) are false. So says the physicalist, anyway.

Let's pause for a quick rundown: The physicalist says that if your notion of causal sufficiency requires you to plug the complicated extrinsic properties and events in for p , then

(O2) is vacuous. If, on the other hand, your notion of causal sufficiency allows you to plug in the more ordinary sorts of events and properties, instances of (O2) will *not* typically be vacuous, but will instead be *false*. Now, we have already seen that the dualist cannot claim that any instance of (O2) is vacuous. So can he claim that all instances are false? Can he adopt this strategy, and say with me that if the physical cause had occurred without the mental one, it would not have caused the same effects?

No. Doing so would abandoning standard ways of evaluating counterfactuals. For the dualist, the closest world in which the C-fibers fire without pain is not a world in which various surrounding physical facts go differently. It is not a world in which the C-fiber firing takes place in a petri dish, or otherwise without crucial background conditions that actually obtain. It is instead a world in which the psychophysical law linking firing C-fibers in such and such circumstances to pains is violated. It is not a full-blown *zombie* world, mind you—that would clearly involve the kinds of “big, widespread, diverse violations of law” that Lewis says it is of the first importance to avoid (1979, 47). It is instead simply a world in which just that particular physical occurrence fails to give rise to the sort of mental one that usually accompanies it. That is merely a “small, localized, simple violation of law,” that allows us to “maximize the spatio-temporal region throughout which perfect match of particular fact prevails” (47-48). This one tiny little violation of psychophysical law is a lot easier to accomplish—if it can be accomplished at all—than a big sweeping change in circumstances.

Crucially, of course, the nonreductive physicalist does not think it can be accomplished at all. As I have already emphasized, she thinks it is a mistake to think of psychophysical laws as breakable nomological connections between distinct things. The dualist and the nonreductive physicalist disagree about what is possible, about what worlds there are—and this forces them to

disagree about which is the closest world in which the antecedent of (O2) is true. For the relevant choices of p , the closest m & $\sim p$ world that the nonreductive physicalist recognizes is *not* an e world. But for those same choices of p , the closest world that the *dualist* recognizes is still an e world. Nothing physical changes at all; given Completeness, e still occurs.¹⁸

Consequently, the dualist cannot say that (O2) is either false or vacuous, and therefore cannot motivate compatibilism in this way. For the dualist, cases of mental causation *do* meet the necessary condition on overdetermination. He thus has no argument for the claim that mental and physical events and properties are so intimately related that they can both be causally sufficient for the same effect without overdetermining it. He has given us no reason to think Exclusion is false of mental and physical causes. In short, it *matters* that the dualist does not think that the connection between physical facts and mental facts is as tight as the nonreductive physicalist does. A mere nomological connection does not fly.¹⁹

7. Compatibilism II: The Notion of Causation

Let us move on, then, to the other strategy for motivating compatibilism. Recall that I said there were two—one that focuses on the intimate relation between the putatively competing causes, and one that focuses on the notion of causation in play. The latter strategy claims that the plausibility of the Exclusion principle, and thus the force of the exclusion argument as a whole, turns upon a mistaken view about the nature of causation—namely, that it involves some kind of *oomph* over and above mere counterfactual dependence. Now, I shall not say anything about whether this really is a mistaken view about the nature of causation. And I shall not argue that

¹⁸ Of course, the dualist may at the end of the day want to avoid the exclusion problem by denying Completeness. But the question at the moment is whether she can avoid the exclusion problem by means of *compatibilism*.

¹⁹ It turns out that Barry Loewer has given a very similar argument for a slightly different conclusion. In his case, it is for the claim that the dualist cannot say that $\sim m \sqsupset \sim e$ is true, rather than (as for me) for the claim that the dualist cannot say that p & $\sim m \sqsupset e$ is false. See 2001, 51-52.

the dualist somehow cannot adopt this strategy—presumably, he too can wade into the causation literature and emerge with whatever view he likes. I certainly see nothing stopping him from adopting the sort of pure dependence view that is allegedly friendly to compatibilism. Indeed, doing so may well help him count the mental as causally efficacious in the first place. Instead, I will argue that this strategy simply does not work. Rejecting oomphy causation does not in fact provide any reason to think that the exclusion principle is false. The force of the exclusion problem does not turn upon any substantive view about the nature of causation.

The two views about causation I have in mind are those that Ned Hall has called ‘dependence’ and ‘production’ (forthcoming; see also 2002).²⁰ The rough distinction is this. According to the production view, causation is a matter of the transfer of energy, or, to be more sardonic about it, the transfer of ‘causal juice’, ‘oomph’, or ‘biff’. This is the kind of view according to which causes really generate their effects, and according to which absences cannot be causes because there is no connecting process. According to the dependence view, in contrast, causation is *purely* a matter of counterfactual dependence (or probability-raising, or something in that ballpark). On this line, patterns of counterfactual dependence fully constitute causal reality; they do *not* merely indicate underlying oomphy causes.

People sometimes suggest, both in conversation and to some extent in print (e.g. Loewer 2002) that the exclusion problem does not get off the ground on the pure dependence notion of causation.²¹ But while I certainly agree that the production view is often in the background of

²⁰ Hall himself thinks that our causal intuitions are not univocal, and that we actually have two concepts of causation.

²¹ I do not count Jackson and Pettit 1990, because they are explicit that their counterfactual notion of causal relevance is mere “causation lite” (Loewer 2002, 661), not full-blooded causation. They therefore deny the efficacy of the mental rather than Exclusion.

discussions of the problem,²² I do not agree that the problem itself actually *requires* it. I do not agree that rejecting it makes the issue go away.

Let me be clear about what the claim is that I am disputing. I am not disputing that a dependence notion of causation might be handy in establishing the causal efficacy of the mental in the first place. That is, it might well help block the worries about the causal relevance of mental properties that arose around Davidson's anomalous monism (see Lepore and Loewer 1989), and it might also help block Princess Elizabeth's worries about the causal powers of Cartesian souls.²³ Whether it can or not depends on whether it is good enough as an account of causation full stop. Those questions are not currently on the table. What *is* on the table is whether a pure dependence notion of causation can defuse the threat of overdetermination by falsifying the exclusion principle. The claim would be that an effect can counterfactually depend on a multitude of causes without being overdetermined. To think that the presence of one cause excludes others, or renders them overdeterministic, is to think of causation like causal juice of which some effects get a double dose. Reject that in favor of dependency, and the exclusion principle falls away, bringing the exclusion problem with it like a house of cards. This is the second strategy for defending compatibilism, the strategy I claim does not work.

Moving to a pure dependency notion of causation is not sufficient to establish that compatibilism is true, that mental and physical causes do not overdetermine their effects. I actually do not think that think that it is necessary, either—as long as one believes that mental and physical causes are appropriately intimately related, I suspect that one can think that

²² Kim, for example, admits that “Loewer is right... in saying that my thinking about causation and mental causation involves a conception of causation as ‘production’ or ‘generation’” (2002, 675). He goes on to try to defend the production model against Loewer's claim that contemporary physics has no place for such a notion. I think Kim is right to admit this, but wrong to assume that the pure dependence notion alone would dissolve the problem completely

causation is *seriously* oomphy and nonetheless claim that mental and physical causes do not overdetermine their effects—but I will set that aside for now.²⁴ All I will argue is that moving to a pure dependence notion is not by itself enough. The real work must be done by an appeal to the relation between the causes, à la the first strategy.

To see this, note that that most believers in a pure dependence theory *also believe that genuine overdetermination occasionally happens*. They think that classic firing squad cases do happen, and that they are importantly different than cases of mental causation. Consider, for example, the familiar point that simple counterfactual theories, according to which *c* is a cause of *e* just in case *e* would not have happened if *c* had not happened, do not allow overdetermining causes to count as causes at all. Such views are forced to say that all apparent cases of overdetermination are really cases of joint causation. They are forced into what Jonathan Schaffer calls the ‘collectivist’ view of overdetermination rather than the ‘individualist’ view (2003). But—and this is the crucial point—everyone thinks that that is a *problem*, and starts looking for a less simple counterfactual theory. Lewis did, for example (cf. 1986, 2000a).

But if the right version of a dependence theory must allow for genuine overdetermination, then that theory alone cannot dismiss the charge that some particular effect is overdetermined. Sometimes effects have two causes and are overdetermined. Sometimes effects have two causes and are not overdetermined. *No* theory of causation that allows both cases can *all by itself* distinguish between them. Only information about the two causes—and, in particular, how they

²³ While it is very hard to see how nonphysical, nonextended souls could actually *transfer energy* to physical things like neurons, it would not be very hard to argue that there are counterfactual connections between, say, acts of will and the contraction of muscle fibers.

²⁴ The trick would be to claim that mental property instances (or events, etc.) and their physical realizers *only provide one injection of oomph*. Events, properties, and the like are individuated differently than transfers of oomph. Two events, one dose of oomph. Now, I do not actually endorse the metaphysics of this solution. But I do think it is worth mentioning, and that it is worth pointing out that the nonreductive physicalist definitely should not reject it on the grounds that it does not give the mental its own causal power. No self-respecting *physicalist* should think that the mental can bring its own oomph, can have *genuinely new* causal powers. That is the mark of emergentism.

are related—can do so. A mere insistence that causation is not oomphy cannot do the job; it cannot distinguish cases of mental causation from cases in which a person is simultaneously hit with two bullets from two independent shooters. So the appeal to a pure dependence theory of causation cannot itself establish that the exclusion principle is false and compatibilism is true. It cannot show that mental and physical causes do not overdetermine their effects.

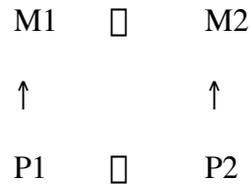
Indeed, I am inclined to suspect that the only way in which the dependence view of causation can help is because anyone who endorses it will be amenable to my counterfactual test for overdetermination, and will consequently be amenable to my own version of the *first* strategy for motivating compatibilism. Be that as it may, the fact is that the only way to properly motivate compatibilism is by appeal to the tight relation between mental and physical causes. And once we go beyond simply asserting that tightly related causes cannot overdetermine their effects, and provide an actual *test* for overdetermination that some pairs of causes pass and others fail, we can see that *compatibilism requires physicalism*. The dualist cannot avail herself of the nonreductivist's solution to the exclusion problem.

8. Conclusion

I have argued that the exclusion problem does not exert the kind of force on a physicalist that it does on a dualist. Dualists really do need to choose between systematic overdetermination, epiphenomenalism, and the incompleteness of physics. Nonreductive physicalists do not. Thus although the argument does provide pressure towards physicalism, it does not provide pressure towards *reductive* physicalism.

A proper understanding of nonreductive physicalism—an understanding that puts the right emphasis on the 'physicalism', and does not get distracted by the 'nonreductive'—makes

the exclusion problem look a lot less threatening. One lesson to be drawn, then, is that these familiar diagrams are dangerous:



They blur the line between emergentism and nonreductive physicalism, and mislead us into thinking that both views are in the same boat vis à vis the exclusion problem—which they straightforwardly are not. The thought that they *are* increasingly seems to me to be based on an almost willful misunderstanding of the nonreductive physicalist’s position. The exclusion argument is an enormous problem for dualist; not for those who say, and mean it, that the mental is nothing over and above the physical.

References

- Bennett, Karen. 2003. Why the exclusion problem seems intractable, and how, just maybe, to tract it. *Noûs* 37: 471-497.
- Boyd, Richard. 1980. Materialism without reductionism: what physicalism does not entail. In Ned Block (ed.), *Readings in the Philosophy of Psychology, Vol. I*. Cambridge, MA: Harvard University Press.
- Burge, Tyler. 1979. Individualism and the mental. *Midwest Studies in Philosophy* 4: 73-121.
- Chalmers, David. 1996. *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford University Press.
- Crane, Tim. 2001. The significance of emergence. In Carl Gillett and Barry Loewer, eds., *Physicalism and its Discontents*. Cambridge: Cambridge University Press, 207-224.
- Crane, Tim, and Mellor, Hugh. 1990. There is no question of physicalism. *Mind* 99: 185-206.
- Crisp, Thomas, and Warfield, Ted. 2001. Kim's master argument. *Noûs* 35: 304-316.
- Hall, Ned. Forthcoming. Two concepts of causation. In J. Collins, N. Hall, and L. A. Paul, eds., *Counterfactuals and Causation*. Cambridge, MA: MIT Press.
- . 2002. Nonlocality on the cheap? A new problem for counterfactual analyses of causation. *Noûs* 36:2, 276-294.
- Hawthorne, John. 2002. Blocking definitions of materialism. *Philosophical Studies* 110: 103-113.
- Hempel, Carl. 1980. Comments on Goodman's *Ways of Worldmaking*. *Synthese* 45:193-9.
- Jackson, Frank. 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Clarendon.
- Jackson, Frank, and Phillip Pettit. 1990b. "Program Explanation: A General Perspective," *Analysis* 50: 107-117.
- Kim, Jaegwon. 1989. The myth of nonreductive physicalism. Reprinted (1993) in *Supervenience and Mind*. Cambridge: Cambridge University Press, 265-284.
- . 1992. Multiple realization and the metaphysics of reduction. Reprinted (1993) in *Supervenience and Mind*. Cambridge: Cambridge University Press, 309-335.
- . 1993. The nonreductivist's troubles with mental causation. Reprinted (1993) in *Supervenience and Mind*. Cambridge: Cambridge University Press, 336-357.
- . 1998. *Mind in a Physical World*. Cambridge, MA: Bradford.
- . 2002. Responses. *Philosophy and Phenomenological Research* 65:3, 671-680.
- . 2003. Blocking causal drainage and other maintenance chores with mental causation. *Philosophy and Phenomenological Research* 67:1, 151-176.
- Lepore, Ernest, and Loewer, Barry. 1987. Mind Matters. *The Journal of Philosophy* 84: 630-642.
- Lewis, David. 1966. An argument for the identity theory. *The Journal of Philosophy* 63: 17-25.
- . 1979. Counterfactual dependence and time's arrow. Reprinted in *Philosophical Papers Vol. II*, NY: Oxford, 32-66.
- . 1983. New work for a theory of universals. *The Australasian Journal of Philosophy* 61: 343-377.
- . 1986. Postscripts to "Causation." In *Philosophical Papers Vol. II*, NY: Oxford, 172-213.
- . 2000a. Causation as influence. *Journal of Philosophy* 97: 182-197.
- . 2000b. Reduction of mind. In S. Guttenplan, ed., *A Companion to the Philosophy of Mind*. Oxford: Basil Blackwell, 412-31.
- Loewer, Barry. 2001. From physics to physicalism. In Carl Gillett and Barry Loewer, eds., *Physicalism and its Discontents*. Cambridge: Cambridge University Press, pp. 37-56.
- . 2002. Comments on Jaegwon Kim's *Mind in a Physical World*. *Philosophy and Phenomenological Research* 65:3, 655-663.

- Melnyk, Andrew. 2003. *A Physicalist Manifesto: Thoroughly Modern Materialism*. Cambridge: Cambridge University Press.
- Merricks, Trenton. 2001. *Objects and Persons*. Oxford: Clarendon Press.
- Papineau, David. 1995. Arguments for supervenience and physical realization. In Savellos and Yalçin, eds., *Supervenience: New Essays*. Cambridge: Cambridge University Press, 226-243.
- . 2001. The rise of physicalism. In Carl Gillett and Barry Loewer, eds., *Physicalism and its Discontents*. Cambridge: Cambridge University Press, 3-36.
- Pereboom, Derk, and Kornblith, Hilary. 1991. The metaphysics of irreducibility. *Philosophical Studies* 63: 125-145.
- Schaffer, Jonathan. 2003. Overdetermining causes. *Philosophical Studies* 114: 23-45.
- Shoemaker, Sydney. 1980. Causality and Properties. In *Time and Cause*, ed. Peter van Inwagen, Dordrecht: D. Reidel: 109–35.
- . 2001. Realization and mental causation. In Carl Gillett and Barry Loewer, eds., *Physicalism and its Discontents*. Cambridge: Cambridge University Press, 74-98.
- Sturgeon, Scott. 1998. Physicalism and overdetermination. *Mind*107: 411-432.
- Stoljar, Daniel. Forthcoming.
- Swoyer, Christopher. 1982. The Nature of Natural Laws. *Australasian Journal of Philosophy* 60: 203–223.
- Unger, Peter. 1979. There are no ordinary things. *Synthese* 4: 117-54
- Van Inwagen, Peter. 1990. *Material Beings*. Cornell University Press, Ithaca.
- Wilson, Jessica. Forthcoming. Supervenience-based characterizations of physicalism. *Nous*.
- Yablo, Stephen. 1992. Mental causation. *The Philosophical Review* 101: 245-280.