Explaining Causal Closure

by Justin Tiehen

The physical realm is causally closed, according to physicalists like me. But why is it causally closed? That is, what *explains* causal closure? The question has not often been addressed. In what follows I argue that physicalists are committed to one explanation of causal closure to the exclusion of others, and that this has profound consequences for how physicalism is to be defended. In particular, I claim that we physicalists need to give up on using a *causal argument* to defend our view. Whatever the problem with dualism is, it's not distinctively causal in any interesting sense.

1.

Following Jaegwon Kim, we can formulate the causal closure thesis as follows.

(Closure): If a physical event has a cause at a time *t*, it has a physical cause at *t*.[1]

What's nice about this formulation is that it makes transparent the logical relation between causal closure and another physicalist thesis:

(P*): All events are physical events.

For my purposes in this paper, I will treat (P*) as a commitment of physicalism.[2] If all events are physical, then it trivially follows that whenever some physical event has a cause, that cause is physical. That is, (P*) entails (Closure). This entailment corresponds to the first proposal for explaining causal closure. According to the proposal, the reason

[1] Kim (2005: 15).

[2] There are in fact some nonreductive physicalists who deny (P*) and hold that mental events are realized by physical events, not identical with them. We could account for such views by replacing (P*) with (P**): All events are or are realized by physical events. This would not affect anything in my argument, since (P**) still entails (Closure) on any plausible construal of realization. For purely expository purposes, I will work with the simpler (P*) in the discussion that follows.

the physical realm is causally closed is just that there are no non-physical events. (Closure) is true *because* (P*) is.[3]

To better understand this proposal, consider a different example. Physicalists sometimes treat the causal closure of the physical as a profound truth about the world. And perhaps it is. But in general, true causal closure theses are easy to come by. For any non-existent – unicorns, the elephant in the room, the 2003 World Series Champion Kansas City Royals, etc. – there is always a true, corresponding causal closure thesis. My own preferred case in this section is Santa Clause. There is no Santa Clause, and thus the domain of things that are not Santa is causally closed.

> (Santa-Closure): If a non-Santa event has a cause at a time *t*, it has a non-Santa cause at *t*.[4]

Why is (Santa-Closure) true? It's true because there is no Santa. *That's* the explanation. According to the present proposal, the explanation of physical causal closure is exactly like this.[5]

This is not the only possible explanation of physical causal closure one can imagine. Perhaps there is some deep fact about physical events, something almost like a shield surrounding them, which prevents causal impingements from non-physical realms (or, at least, which permits such impingements only when a physical cause is also taking place). There are different ways one could try to cash out this shield metaphor, but one natural way of doing so invokes laws of nature: perhaps (Closure) is true *because* it is a

[3] Here and throughout I operate with an "ontic" conception according to which explanation is an objective, mind-independent relation between things in the world, like events or facts.

[4] A non-Santa event is an event not having Santa Clause as its constituent object. All actual events are thus non-Santa events.

[5] The parallel between (Closure) and (Santa-Closure) might be even closer than it initially appears. According to via negativa characterizations of *the physical*, to be physical is just to be non-mental. In that case, the "negative" components of (Santa-Closure) – the *non*-Santa events – are mirrored by negative components of (Closure), as revealed under analysis – physical events just are *non*-mental events. On the via negativa approach, see for instance Montero (1999) and Papineau (2002: Ch. 1).

law.[6] In putting things this way I am assuming that laws explain their instance – that (Closure)'s lawhood would explain its (mere) truth. The proposal doesn't have to be put this way though. We could instead say that if (Closure) is a law then it is explanatorily basic, and so has no explanation. That is not how I will be framing things, but I have no deep objection to doing so. What matters most to my argument is the difference between this present proposal (however we think of it, exactly) and the first one, which explains (Closure)'s truth by appealing to (P*).

Other potential explanations of (Closure)'s truth seem possible as well, but for my purposes we can stop at these two. Replace the second proposed explanation of causal closure with any third explanation you please, and the arguments I'm about to give should still go through.

2.

I claim that the two proposed explanations of (Closure)'s truth *exclude* one another, in the sense that if you accept one, you cannot reasonably accept the other. The first step in defending an exclusion claim of this sort is to show that the two explanations are independent of one another. That is what I will do in this section.

Imagine God is building a world, and the one thing he wants to guarantee about the place is that (Closure) is true there. Here are three divine building plans he could choose from. First, he could decree that no non-physical events are to take place in his world – that is, that (P*) is to be true. If God did this, he wouldn't also need to go to the trouble of making (Closure) a law, since (P*)'s truth by itself ensures (Closure)'s truth.

[6] This way of putting things assumes that universal generalizations like (Closure) are the sorts of things that could be laws. However, my arguments are all compatible with alternative conceptions of laws, such as the Tooley (1977)/Dretske (1977)/Armstrong (1983) view that laws are not themselves generalizations but rather relations among universals.

Second, God could issue a decree that causal closure is to be a law in his world. If he did this, he wouldn't also need to go to the trouble of making (P*) true, since (Closure)'s lawhood secures its truth even if non-physical happen to occur. Third, God could do both these things. He could both issue a decree that (P*) is true and also issue a decree that (Closure) is a law. If these three divine building plans really are distinct, then the two proposed explanations of (Closure)'s truth must be independent of one another, which is what I'm trying to establish.

Counterfactuals give us another, less metaphorical way to appreciate the independence; looking at them will also help set up the discussion below. Consider a principle of the form: If P explains Q, then (barring explanatory overdetermination) if P had not obtained, Q would not have obtained. Some such principle is widely accepted in the special case of causal explanation, but here I intend it as a principle governing explanation generally. Now consider the following counterfactual: If there had been non-physical events, the physical realm would have been causally closed. That is,

(CF): ~(P*) > (Closure).

I say that whether (CF) is true depends on what explains causal closure. If the physical realm is causally closed solely because there are no non-physical events, it should follow that if there *had been* non-physical events, the physical realm would not have been causally closed. That is, if (P*) explains (Closure), then (CF) is false. On the other hand, if causal closure is a law of nature, it should follow that if there *had been* non-physical events, the physical realm still would have been causally closed. This follows because laws support counterfactuals. The closest counterfactual worlds where non-physical events take place will be worlds where any actual laws compatible with the existence of

such events still hold. (Closure) is so compatible, and so it will hold in such worlds. Thus, if (Closure)'s lawhood is what explains its truth, (CF) is true.

Consider then three different possible worlds where (Closure) is true but has different explanations. At the first, where the explanation of causal closure is solely that there are no non-physical events, (P*) will be true and (CF) false. At the second, where the explanation of causal closure is solely that it is a law, (P*) will be false and (CF) true. At the third, where it is both the case that there are no non-physical events and also that causal closure is a law, both (P*) and (CF) will be true. That these three distinct worlds with their different assignments of truth values to (P*) and (CF) are all possible helps confirm for us that the two proposed explanations of (Closure)'s truth are independent. I will entertain an objection to this independence claim below, but for now let's treat the matter as settled.

3.

Given this independence, I say that the two proposed explanations exclude one another. Reconsider the third divine building plan. Given that the goal is just to ensure (Closure)'s truth, a God who does this by both making (P*) true and also making (Closure) a law isn't very efficient – he has done double the work needed. He has made (Closure)'s truth overdetermined, in a sense. This isn't *causal* overdetermination, of course. It's not that (Closure)'s lawhood causes its truth, for instance.[7] Rather, it's a matter of there being two completely independent facts – that all events are physical, and that causal closure is a law – each of which by itself is fully sufficient for guaranteeing (Closure)'s truth. Each of which *explains* (Closure)'s truth.

[7] Things are less obvious with the first explanation. Maybe it can be understood as a kind of causal explanation that appeals to absences as causes, if one takes the truthmaker of (P*) to be the negative fact that there are no non-physical events. I don't need to take a stand on this for my purposes here.

Call this *explanatory overdetermination*. Causation is one type of metaphysical determination relation, and so causal overdetermination is a form of explanatory overdetermination. There are other determination relations as well though: supervenience, realization, constitution, composition, and so on. Wherever there is metaphysical determination, it would seem, there is the potential for overdetermination (or, at least, for a view entailing that there is overdetermination).[8] To bring out parallels between causal and non-causal overdetermination, I turn now to the causal argument for physicalism, which will be my focus in the remainder of the paper.

There seems to be something objectionable about any view of the mind that posits systematic causal overdetermination in order to account for mental causation. This thought plays a central role in the causal argument, which I canonically represent here:

(P1): If a physical event has a cause at a time *t*, it has a physical cause at *t*.
(P2): All mental events have physical effects.
(P3): The physical effects of mental causes are not all causally overdetermined.
(C): Mental events are identical with physical events.[9]

Notice that the first premise just is (Closure), while the conclusion I will be treating as equivalent to our (P*).[10] For now though, our focus is on the anti-overdetermination

[8] This thought is inspired by Kim's most recent presentations of his exclusion argument. In his (2005: 39-40), Kim suggests there is a prima facie tension between holding that a mental state M both supervenes on a physical state P and is caused by some distinct mental state M'. Given that supervenience and causation are both determination relations, M seems overdetermined – though, of course, not causally overdetermined, since supervenience is not a causal relation. According to Kim the way to resolve this tension is by supposing that M' causes P. This, however, leads to the familiar causal exclusion problem, since P, being a physical state, has fully sufficient physical causes.

[9] This version is taken from Papineau (2002) with minor adjustments. Besides Papineau, proponents of some version or other of the causal argument include Smart (1959), Lewis (1966), Davidson (1970), Tye (1995), Levine (2001), and Melnyk (2003). It is the single most influential argument offered for physicalism today.

[10] (C) says that all mental events are physical, while (P*) says that absolutely all events are physical. These claims are equivalent if we assume the via negativa approach which defines "the physical" as the non-mental. Even if we don't assume this, treating (C) and (P*) as equivalent is harmless since proponents generally take the causal argument to generalize to all events.

premise, (P3). Why should we accept (P3)? What is the alleged problem with causal overdetermination? The issue is controversial, but here is one stab at an answer.

I assume the problem isn't that causal overdetermination is metaphysically impossible. I take it there are bizarre worlds out there where (P3) is false, just as there are bizarre worlds where everyone who dies is killed by a pair of simultaneous gunshots to the heart. Instead, the problem seems *epistemic*, in a way. It is very difficult to imagine what evidence could reasonably convince us that we actually live in one of those worlds where (P3) is false. To bring out the difficulty, imagine that we already possess a fully sufficient causal explanation for all physical effects in terms of purely physical antecedent causes. If so, then what justification could we possibly obtain for positing additional, non-physical mental causes of those effects? Positing such further causes inevitably would seem gratuitous. It would bring us no gain in explanatory power.

A similar dynamic arises in our case of non-causal overdetermination. Holding that (Closure)'s truth is non-causally overdetermined is objectionable, but not because such overdetermination is metaphysically impossible. I assume there are worlds where it is both the case that (P*) is true and also that (Closure) is a law. The trouble is that it is difficult to imagine what evidence could reasonably convince us that we actually live in such a world. To bring out the difficulty, imagine that we already know that (P*) is true and now want to figure out whether (Closure) is a law. What possible evidence could convince us that it is? Clearly, this won't do: going out and observing causal chain after causal chain that complies with (Closure), while never finding a single chain that violates it. Such observed compliance with (Closure) would fail to generate any pressure on us at all to infer (Closure)'s lawhood, since we already possess a fully sufficient explanation

for such compliance: the truth of (P*). Given our knowledge that (P*) is true, we already know in advance that all causal chains will comply with (Closure), regardless of its standing as a law. But if that's right, such observed compliance could do nothing to confirm (Closure)'s lawhood.

The role explanatory overdetermination is playing here can be clarified by considering an alternative scenario in which such overdetermination is missing. Imagine we know that (P*) is false – maybe because we have a knockdown proof for dualism – and yet still we find that all observed causal chains comply with (Closure). In that case, the observed compliance really *would* seem to speak in favor of (Closure)'s lawhood. For, such compliance would cry out for explanation, while if (P*) is false it could not do the job. An alternative explanation would be needed. That (Closure) is a law of nature would seem to be an excellent candidate. In this scenario, observed compliance could help confirm (Closure)'s lawhood. What this demonstrates it that the epistemic problem I'm pointing to here isn't due to (Closure) taken by itself. It arises only through explanatory overdetermination.

Turning back to the scenario where we know (P*) is true, admittedly it is the case that there won't be disconfirming evidence of (Closure)'s lawhood in the form of causal chains violating (Closure). We know that there won't be any such chains, given (P*)'s truth. There is an asymmetry between lawhood and non-lawhood, though. For familiar reasons of parsimony, a theory should posit as few explainers as it needs to get by. Because laws are explainers, there is an initial presumption against the lawhood of any given proposition, a presumption which can be overridden only when there is a gain in explanatory power to be had. If we know (P*) to be true, it's difficult to see how the

presumption against (Closure)'s lawhood ever could be overridden. Given (P*)'s truth, all the explanatory work has already been done, leaving nothing for a (Closure) law to do. It is explanatorily idle. Thus, it is explanatorily excluded.

This at least is my own preferred analysis of what's wrong with explanatory overdetermination, both causal and non-causal.[11] For the sake of the arguments that follow it's not strictly required that you buy the analysis, all that's required is that you agree that non-causal overdetermination is problematic in much the way that causal overdetermination is, whatever way that happens to be. And this seems plausible independently of the analysis offered here: surely the trouble with causal overdetermination lies in the overdetermination part, not the causal part.[12]

4.

Let us grant last section's conclusion that explanatory overdetermination is to be avoided. What follows? First, since all physicalists are committed to accepting (P*), it follows that no physicalist reasonably can accept any other explanation of causal closure, on pain of embracing explanatory overdetermination. Thus, no physicalist reasonably can accept that (Closure) is a law. Now, this conclusion does not undermine physicalism per se, for physicalists as such need not be committed to (Closure)'s lawhood. But, I now will argue, it does seriously undermine the causal argument for physicalism.

---

[11] The analysis entails that explanatory overdetermination in general is problematic just when there are not independent sources of evidence for the different explainers. So, for instance, consider a death causally overdetermined by a pair of simultaneous shots to the heart. According to the analysis, this case of overdetermination is unproblematic if we suppose that there is independent evidence for the causal efficacy of each shot – say, in the form of two separate entrance wounds. And this seems right: this sort of overdetermination doesn't seem especially objectionable. Things are different with the sort of overdetermination at issue in (P3) of the causal argument, though, for there the mental and the physical causes do not leave separate traces. On this very point of difference between the two cases of causal overdetermination, see Kim (2005: 48).

[12] Some, like Sider (2003), have argued that systematic causal overdetermination is not problematic at all. We can set aside the merits of such a view since it is not available to causal argument proponents, given their use of the anti-overdetermination premise (P3).

The crucial issue is how causal argument proponents hope to support their first premise, the (Closure) premise. Speaking as a physicalist who is not a causal argument proponent, my own preferred defense goes like this: first we establish that physicalism is true, (P*), and then from (P*) we deductively infer (Closure). This sort of defense obviously will not do for causal argument proponents. If the causal argument is to be anyone's rational basis for *initially* accepting physicalism, what they need is a defense of (Closure) that does not depend on any prior argument for (P*).

In principle, an a priori defense of (Closure) could be mounted.[13] Many contemporary physicalists will be leery of such a move, however, since it would drain the case for physicalism of any serious empirical content.[14] What's needed is an *empirical* warrant for (Closure). Toward this end, many causal argument proponents have come to defend an *inductive* argument for (Closure), broadly along the lines suggested here by Andrew Melnyk.

> [It is not true] that in order to be persuaded of the causal closure of the physical one must already be persuaded of physicalism. To see this, it is necessary only to review how the closure principle is usually evidenced. First we become persuaded, on the basis of observational evidence and ordinary canons of scientific reasoning, that various physical effects have sufficient physical causes, since the best available explanations of those effects posit physical and only physical causes; surely no assumption of physicalism is needed to take the first step. Then, employing enumerative induction, we treat these well-supported explanations as evidence that *all* physical effects have sufficient physical causes.[15]

[13] In attacking his interactionist dualism, a number of Descartes' critics through the years have made use of seemingly a priori arguments for (Closure). More recently, Davidson's (1970) and (1995) case for (Closure) looks pretty a priori: it turns on apparently conceptual points about causation, laws, and the anomalousness of the mental. In addition, the "pairing problem" discussed by Foster (1991: Ch. 6) and Kim (2005: Ch. 3), among others, potentially serves as the basis of an a priori defense of (Closure).

[14] Within the causal argument, (P2) is supported by commonsense while (P3) seems to be supported by largely a priori considerations. Thus, if there is a role for serious science to play, it needs to be in support of (Closure).

[15] Melnyk (2003: 289-290). Papineau (2002) offers another influential defense of the inductive strategy.

I regard this inductive strategy as causal argument proponents' best hope for coming up with the sort of defense of (Closure) they need, and so in the remainder of this paper will focus my attention on it.

For the inductive strategy to work, (Closure) needs to be the sort of generalization that *can* be confirmed inductively. Notoriously, not just any true proposition can be. It may be that all the employees of the Fort Worth Olive Garden have a third cousin named Amy, but this isn't plausibly the sort of thing we could learn through induction. Which propositions are inductively confirmable, then? Here is one historically influential answer: only the laws of nature are inductively confirmable, non-laws are not.[16] Below I'm willing to suppose that this principle isn't quite right, but for now I propose we operate with is so that I can set out the central argument of the paper.

If this view of the connection between laws and induction is right, then since no physicalist reasonably can accept that (Closure) is a law, it follows that no physicalist reasonably can accept that (Closure) is the sort of proposition susceptible to inductive confirmation. Thus, it follows that no physicalist reasonably can accept her physicalism on the basis of the causal argument taken in conjunction with the inductive defense of (Closure). But, we are assuming, the inductive defense is the only promising way for a causal argument proponent to support (Closure) without first establishing (P*). Thus, it follows that no one reasonably can accept physicalism on the basis of the causal argument.

Call this the *explanatory exclusion objection* to the causal argument. Ironically, according to the objection, the problem with the causal argument is that it implicitly assumes an objectionable form of explanatory overdetermination – the overdetermination

[16] See for instance Goodman (1983).

of the truth of (Closure). My presentation of the explanatory exclusion objection was kept brisk so that the main line of reasoning could be surveyed at a glance. Now that the basic structure of the argument is on the table, it will be helpful to flesh out the details a bit. In the process, I hope to make the objection seem intuitively compelling. I want the causal argument for physicalism to come to seem *obviously* problematic, in a way the explanatory exclusion objection helps bring out. Toward this end I now turn to a series of objections to my argument followed by replies.

5.

Objection. The assumption that only the laws of nature are inductively confirmable is surely wrong.[17] If, more reasonably, we suppose that some non-laws are inducible, then causal argument proponents can make their inductive case for (Closure) without committing themselves to its lawhood, thereby avoiding the exclusion problem.

Reply. The assumption that only laws are inducible was helpful for illuminating the exclusion problem, but the problem can be generated using far weaker assumptions about induction. Let's grant that some non-laws are inducible. To see why this won't help here, let's return to the point made above, that for every non-existent there is always a true, corresponding causal closure thesis. For instance, the 2003 Kansas City Royals did not manage to win the World Series. Therefore, the domain of things that aren't the 2003 Champion Royals is causally closed:

> (KC-Closure) If a non-Champion Royal event has a cause at a time *t*, it has a non-Champion Royal cause at *t*.[18]

[17] For an especially influential argument for this conclusion, see Salmon (1989: 49).
[18] A non-Champion Royal event is an event not having a member of the 2003 World Champion Royals as its constituent object. All actual events are thus non-Champion Royal events.

For very, very many such corresponding causal closure theses – in fact, I would venture nearly all of them – it is completely obvious that they are not susceptible to inductive confirmation. How do we confirm (KC-Closure), for instance? Surely not by induction! It's not that we go out and observe a bunch of non-Champion Royal effects with non-Champion Royal causes and then – *prior* to taking a stand on whether the Royals won the 2003 World Series – directly inductively infer (KC-Closure). Rather, it's obvious that our warrant for (KC-Closure) depends on our antecedent warrant for the proposition that the 2003 World Series Champion Royals don't exist. And that's why no one would try to use a causal argument to show that the Royals didn't win the World Series in 2003, even though (KC-Closure) is far less controversial than (Closure) is.[19] As things go for the World Champion Royals, they also go for many, many other non-existents. No one uses a causal argument against Santa Clause, or unicorns, or the elephant in the room, because in each case our warrant for believing that those entities don't exist is prior to our warrant for believing the corresponding causal closure thesis.

Now, no defender of the causal argument for physicalism should deny any of this. They should grant causal arguments do not work against many non-existents, but insist that things are different when it comes to non-physical events. Why are things different here, however? Is this merely special pleading, or is there a principled reason for distinguishing physical causal closure from these other causal closure theses?[20]

[19] Just to make it explicit, the argument would go as follows: (P1): (KC-Closure); (P2): If the Royals won the World Series in 2003, at least some non-Champion Royal events have Champion Royal causes; (P3): It's not the case that all non-Champion Royal effects of Champion Royal causes are causally overdetermined; (C): The Royals did not win the 2003 World Series.

[20] This question can be framed in terms of a meta-induction: given the non-inducibility of (KC-Closure) and all those other causal closure theses, why shouldn't we inductively infer the non-inducibility of (Closure)?

Here would have been a promising response to this question. The difference, one might have said, is that while these other causal closure theses are true only because the corresponding entities don't exist, (Closure) is true for a different reason: it is true because it is a law of nature. If this response had been acceptable it would have solved the problem facing defenders of the causal argument for physicalism, it would have allowed us to see how (Closure) could be inductively confirmable even though all those other causal closure theses are not, which in turn would have allowed us to understand why the causal argument for physicalism succeeds while the causal arguments against the 2003 World Champion Royals and all the rest do not. The trouble, as we have seen, is that this otherwise promising response is unavailable to physicalists. It would commit them to taking the truth of (Closure) to be explanatorily overdetermined.

What defenders of the causal argument need is a response to the question at hand that avoids this pitfall. Thus, somehow or other they need to divorce the issues of inductive confirmation and explanation. With respect to induction, they need to hold that (Closure) is *unlike* (KC-Closure) and all those other causal closure theses, in that it is inductively confirmable even though the others are not. With respect to explanation, though, they need to hold that (Closure) is *just like* (KC-Closure) and all those other causal closure theses, in that what explains (Closure) is (P*), just as what explains (KC-Closure) is that the 2003 World Series Champion Kansas City Royals don't exist.

I do not see how defenders of the causal argument can have it both ways. Given widely held and extremely compelling views about the intimate connections between

confirmation and explanation,[21] I do not see how there could be a difference in the inductive confirmability of (Closure) versus (KC-Closure) unless there is a corresponding difference in what explains (Closure) versus (KC-Closure). This is the crucial assumption about induction that my explanatory overdetermination objection really turns on – the crucial assumption is not that only the laws are inducible, but rather that confirmation and explanation are closely connected such that inductive differences of the sort in question must always be backed by explanatory differences. As mentioned, that there is a connection of this sort between confirmation and explanation is widely accepted. If defenders of the causal argument for physicalism are forced to reject this principle in order to make their argument, this is a startling result. It gives us good reason to be suspicious of the causal argument.[22]

Let's summarize the results of this section. If they are to avoid positing explanatory overdetermination, physicalists are forced to regard (Closure) as being just like (KC-Closure) and all those other causal closure theses – both with respect to explanation, and – *therefore* – with respect to inductive confirmation. But then, just as (KC-Closure) is plainly not inducible and there is no viable causal argument to be run showing that the Royals didn't win the 2003 World Series, we should expect that (Closure) will not be inducible and that there is no viable causal argument to be run showing that physicalism is true.

---

[21] For classic discussions of the close connection between inductive confirmation and explanation, see for instance Harman (1965) and Dretske (1977). According to Dretske explanation is the converse of confirmation; causal argument proponents had better hope that the relation is not that tight.

[22] In connection, consider the following proposal: perhaps (KC-Closure) is not inducible while (Closure) is because the property of not being a member of the 2003 Champion Royals is gruesome while the property of being physical is natural. My response is that I can see how the naturalness of being physical could help make (Closure) inducible only if it somehow *explains* the truth of (Closure), for instance, by entering into a natural law or some other explanatory relation that explains (Closure). But again, no physicalist can accept such an additional explanation for (Closure) on pain of embracing explanatory overdetermination.

6.

Objection. A successful objection to the causal argument for physicalism should deny the argument's soundness, or, granting soundness, contend that the causal argument is question begging. But the explanatory exclusion objection does neither of these things.

Reply. Right. I grant that the causal argument is sound and not question begging; what I deny is something else. To clarify what I deny, and to set up next section's discussion, it will be helpful to import some terminology from recent epistemological work on the *transmission of warrant*.[23]

In order to obtain knowledge of a conclusion by virtue of going through a given argument for that conclusion, one's belief in the argument's premises must be warranted, and the warrant for those premises must not depend entirely on antecedent warrant for the conclusion itself.[24] So, take the causal argument. If we are to come to know the truth of physicalism through it, then our belief in each of the causal argument's premises must be warranted – otherwise, our subsequent belief in its conclusion will be unwarranted and thus not knowledge – and our warrants for these premises must not depend entirely on prior warrant for physicalism – otherwise, our knowledge of physicalism will be due not to the causal argument but to however we obtained this prior warrant. An argument that does not meet these conditions is said not to be *cogent*. What I deny is the causal argument's cogency, in this special sense.

This denial needs to be understood with care. There is a close connection between arguments that are not cogent and arguments that beg the question, but it would be unfair for me to accuse typical causal argument proponents of begging the question – a

[23] Here I rely especially on Wright (2003) and (2004).
[24] A single premise can have multiple warrants. It is alright for some of the warrants for a given premise to depend on antecedent warrant for the conclusion, but not for all of the warrants for the premise to do so.

charge Melnyk is sensitive to in the passage quoted earlier. For, if the inductive argument for (Closure) succeeded in providing it with a warrant, then the causal argument would be cogent, and no question begged. In light of the explanatory exclusion objection, however, I deny that the inductive argument succeeds in providing a warrant for (Closure). This is the core problem. Now, as a physicalist, I do think that belief in (Closure) is warranted. But, I think this warrant depends entirely on our prior warrant for (P*). It would indeed be question begging for causal argument proponents to appeal to this warrant in making their case, but they are clear sighted about this. That is why they (unsuccessfully) make their inductive argument instead.

So, then, my criticism of the causal argument is not that it is question begging. My criticism is that the only *successful* argument for the (Closure) premise that there is cannot be conjoined with the causal argument without rendering the causal argument question begging. That is what I mean when I say the causal argument is not cogent.

7.

Objection. Contrary to what was argued in §2, the two proposed explanations of (Closure) need not be independent of one another. Suppose (P*) is not just true but also a law. This is arguably a requirement of physicalism anyway.[25] Suppose further that the set of laws is closed under deductive entailment. It then follows that (Closure) is a law – a derived law. This shows that a physicalist can after all reasonably hold both that (P*) is true and that (Closure) is a law – she can do so provided she takes (P*) to be a law. In that case the threat of explanatory exclusion does not arise: (Closure) will be inductively confirmable because it's a (derived) law, while (P*) will also be true.

[25] For discussion, see Chalmers (1996: 363-364), who takes physicalism to require the nomological impossibility of non-physical events, and Melnyk (2003:28), who denies this.

Reply. Let's start with what I reject in this objection. Even if (P*) is nomologically necessary, I have real doubts about whether it could qualify as a law.[26] But granting that it is a law, I deny that the laws are closed under deduction. Instead of adopting a defensive posture trying to justify these (admittedly controversial) views, it will be more illuminating to go on the attack and present two arguments showing that it would be disastrous for defenders of the causal argument to adopt the line suggested by the present objection. Each of these arguments pinpoints a very general challenge that all causal argument proponents must meet; the problem with the present objection is that endorsing it guarantees these challenges go unmet.

First, if the present objection is embraced, it looks as though we will be forced to give up on the intuitive idea driving the causal argument, the idea that non-physical events have a distinctively causal problem, a problem not shared by everything else. To see that this is so, reconsider (CF).

(CF): ~(P*) > (Closure).

If we assume with the objection that (P*) is a law, then (CF) is a counterlegal – a counterfactual whose antecedent is nomologically impossible. Still, we can ask whether (CF) is true or false. If, in accordance with the objection, (Closure)'s lawhood is derived from (P*)'s, then it would seem that (CF) must be false. For in that case, if there had been non-physical events then (P*) would not have been a law (lawhood requires truth), and so (Closure) would not have been a law (since its lawhood is derived from (P*)'s), in which case there would have been nothing to prevent those non-physical events from violating (Closure).

[26] The problem is that I doubt the determinable property of being physical can qualify as a natural kind. Discussion of this point would take us too far afield, however.

If (CF) is false, however, there is no interesting sense in which non-physical events have a special causal problem. For, absolutely everything is such that it wouldn't enter into causal relations *if it didn't exist*. The proper criterion for an entity's having a distinctively causal problem needs to be that the entity would be epiphenomenal (or causally redundant) *even if it existed*. Given the falsity of (CF), which again seems to follow from the present objection, non-physical events do not have a distinctively causal problem in this minimal sense.[27] Perhaps we can say that they have an *existential* problem, in that they don't exist, but it would be perverse to call this a *causal* problem merely on the grounds that non-existents aren't causes. It would be like saying that Santa Clause has a visa problem, since (not existing) he has not obtained a visa to work in the United States.

Here is a striking way to reinforce the point. Any false theory of the mind/body relation can, in effect, be construed as positing non-existents. Take the reductionistic type identity theory, according to which, for instance, pain is identified with firing C-fibers. Let's suppose, for familiar reasons having to do with multiple realizability, that the type identity theory is false. Then the type identity theory says there are reducible mental properties – that is, mental properties identical to physical properties – when in fact there are no such properties. And similarly, the type identity theory says there are

[27] Compare Papineau (2001:11), who, echoing J. J. C. Smart, intuitively glosses the causal argument as follows: "*If* conscious properties were non-material, they *would* thus be epiphenomenal 'danglers', caused by physical occurrences but themselves having no effects on physical activities . . . if there were compelling independent grounds for holding that conscious properties are non-material, then we would have no option but to accept epiphenomenalism about consciousness" (emphasis added). If in accordance with the present objection we suppose that (Closure) is a derived law from (P*), though, Papineau's counterfactual is false: if conscious properties were non-material, they would not be epiphenomenal danglers. Rather, if conscious properties were non-material then the causal closure derived law would no longer hold, and so there would be nothing to prevent such non-material properties from causing physical activities. A great deal of what causal argument proponents say when explaining their view will need to be rejected if the present objection is embraced and (Closure) is taken to be a derived law from (P*).

events that are instantiations of reducible mental properties, when in fact there are no such events. The following anti-reductionistic thesis is thus true:

(A-R): No event is the instantiation of a reducible mental property.

Now, as we have seen, whenever there is a non-existent there is always a corresponding causal closure thesis. Thus, (A-R) entails:

> (A-R-Closure): If a physical event has a cause at time *t*, it has a cause that is not the instantiation of a reducible mental property at *t*.[28]

To get a feel for (A-R-Closure), consider the sorts of causal interactions it covers. Some of these interactions will involve the instantiation of mental properties. Suppose I cry (a physical event), and the cause at *t* of my crying is my pain. If my pain is the instantiation of a reducible mental property – if pain is identical to firing C-fibers, say – then this causal interaction violates (A-R-Closure) (barring overdetermination). However, if my pain is the instantiation of an irreducible mental property – if dualism or nonreductive physicalism is true and pain is identical to no physical property – then the causal interaction complies with (A-R-Closure). Aside from causal interactions involving mental events, (A-R-Closure) also covers causal interactions involving just physical events. Suppose one electron repels another. The effect is physical. The cause, which we can suppose is the first electron's having a negative charge, is an event that is not the instantiation of a reducible mental property – having a negative charge is identical to no mental property. Thus, this interaction too is covered by (A-R-Closure), and complies with it.

We can now set out the causal argument that could be seen coming a mile away.

[28] I formulate (A-R-Closure) in terms of physical effects, rather than in terms of effects that are not the instantiations of reducible mental properties, in order to feed it more smoothly into the causal argument that's coming.

(P1): (A-R-Closure).
(P2): All mental events have physical effects.
(P3): The physical effects of mental causes are not all causally overdetermined.
(C): Mental events are not the instantiations of reducible mental properties.

In effect, (C) is the denial of the type identity view; thus, this is a causal argument against the view. We can summarize the argument's reasoning as follows. Given (A-R-Closure), if the type identity theory is true then mental events are either epiphenomenal danglers or they are causally redundant. But mental events are not epiphenomenal. And they are not causally redundant. Therefore, the type identity theory must be rejected. It must be rejected to save mental causation!

This new causal argument is sound. Or, at least, it should be regarded as sound by anyone who rejects the type identity theory – the majority of philosophers today, including most physicalists. Admittedly, the defense of (A-R-Closure) I have provided assumes the falsity of the type identity theory, and so cannot be used to establish the new causal argument's cogency. We weren't talking about cogency yet, however. We were talking about soundness, and as I said the new causal argument is sound.

Let us only now turn to cogency. Proponents of the new causal argument cannot rely on my above defense of (A-R-Closure), but if they could obtain some other warrant for the thesis, they could establish cogency. Here is my recommendation: make an inductive case for (A-R-Closure). Go out, observe causal chains that comply with (A-R-Closure) and none that don't, and then directly inductively infer (A-R-Closure)'s truth without first inferring (A-R)'s truth. The causal chains to be used as inductive evidence would mostly be just the same chains that proponents of the old causal argument use in making their inductive case for (Closure) – they would consist of things like one electron

repelling another.[29] But, whereas proponents of the old causal argument take this evidence to support (Closure), proponents of the new causal argument take it to support (A-R-Closure).[30]

At an intuitive level, you can think of the inductive case for (A-R-Closure) like this. It looks like we can causally explain quite a lot of physical effects without ever needing to posit as causes the instantiations of reducible mental properties. Now, maybe nature is not uniform. Maybe some of the causes of physical effects are the instantiations of reducible mental properties (although we have no uncontroversial examples of this), even though most causes of physical effects are not. But standard inductive practice gives us good reason to accept (A-R-Closure) at this point until proven otherwise.[31]

The new causal argument doesn't really work, of course. Whatever the problems the type identity theory may have, there is no sense at all in which it has a causal problem. This much should be common ground even among those who reject the theory. A challenge for defenders of the old causal argument for physicalism, though, is to say why their argument succeeds while the new causal argument against the type identity theory so obviously fails. After all, the two arguments are structurally parallel. If, in accordance with the objection to which we are presently responding, we take (Closure)'s lawhood to be derived from (P*)'s, I can see no way to meet this challenge.

---

[29] I recommend that defenders of the new causal argument not use my pain causing my crying as inductive evidence for (A-R-Closure). Doing so would be question begging – it would assume that my pain is not the instantiation of a reducible mental property. By the same token, I recommend that defenders of the old causal argument not use this causal chain either in their inductive case for (Closure), for to do so would be to assume that my pain is a physical event.

[30] They both could be right. Potentially, the two causal arguments could be used together to mount a novel causal defense of nonreductive physicalism.

[31] Obviously, grue-like considerations are playing some role here. But, I think a halfway plausible case can be made that (A-R-Closure) is not obviously unprojectible, in the way that the hypothesis all emeralds are grue is. For, intuitively at least, there really does seem to be a metaphysically significant difference between those physical properties that are identical with mental properties and those that are not. So, why not think that the former could figure in a projectible hypothesis?

Here's why. My objection to (Closure)'s inducibility does not turn on the proposition's *modal strength*. Yes, if (P*) is a law then (Closure) is nomologically necessary. If you then want to call it a "derived law," go ahead. But by the same token, it is highly plausible that if the type identity theory is false, its falsity is at least nomologically necessary. In that case, anyone who is not a type identity theorist should regard (A-R) as nomologically necessary and (A-R-Closure) as thus a similar "derived law." Given this status as a derived law, should we go back and say that (A-R-Closure) must be inducible after all, that the new causal argument against the type identity theory is thus cogent, and that type identity theorists have a terrible problem with mental causation? Surely not! In holding that type identity theorists do not have a causal problem, the modal strength of (A-R-Closure) is beside the point. What is to the point is what explains (A-R-Closure). As long as it is true *only because* of (A-R) – either (A-R)'s truth or its necessity – there is no interesting sense in which type identity theorists have a distinctively causal problem, rather than a mere existential problem.

Similarly, the modal strength of (Closure) is beside the point regarding the question of whether dualists have a causal problem. What is to the point is, once again, what explains (Closure). If (Closure) is true *only because* of (P*) – either its truth or its necessity – then dualists have no more of a causal problem than type identity theorists do. Their causal situations are perfectly parallel. In that case, just as we hold that (A-R-Closure) is not inducible and that the causal argument against the type identity theory is not cogent, we should hold that (Closure) is not inducible and that the causal argument for physicalism is not cogent. The present objection, which takes (Closure) to be a

derived law from (P*), is disastrous for defenders of the causal argument for physicalism in that it seemingly makes these results inevitable.

Let's turn now to the section's second argument. Suppose we grant the present objection: (Closure) is a derived law, and thus it *is* inductively confirmable, making the causal argument cogent. Fine. If the causal argument's cogency can be secured in this way, then so too can the cogency of myriad other (non-causal) arguments for physicalism, in which case physicalists have no pressing need for the causal argument. To illustrate the point, I here present a *temporal* argument for physicalism. I focus on time for a reason. It is generally thought that while dualism may have its problems, it has no special problem with time. The temporal argument for physicalism says otherwise.

The argument:

(P1): Every event that is temporally preceded by a physical event is itself a physical event.
(P2): All mental events take place in time.
(P3): The Big Bang (a wholly physical event) precedes all other events.
(C): All mental events are identical with physical events.

The temporal argument is meant to be roughly analogous to the causal argument, with (P1) serving as the analogue to (Closure). The temporal argument is clearly valid and plausibly sound. (P2) is a metaphysical necessity, given the identity conditions for (Kimian) events. (P3) is supported by our best cosmological theory. (P1) is deductively entailed by (P*), and so is something that all physicalists must accept.

Moving from soundness to cogency, the crucial question is whether (P1) has a warrant that does not depend on any prior warrant for (C). In accordance with the present objection, (P1) qualifies as a derived law since it is entailed by (P*). If this status as a derived law is good enough to secure (Closure)'s inductive confirmability, the same must

be true of (P1). Now, if we assume that causes at least typically precede their effects, much of the inductive evidence for (Closure) will also count as inductive evidence for (P1), although (P1) will of course have a lot of further inductive support as well, in the form of causally unconnected sequences of events. Since there is more available evidence for it, then, perhaps (P1) counts as having *better* inductive support than (Closure). At any rate, with the temporal argument readily available, why does the physicalist need a causal argument?

Surely, though, the temporal does not really work, and just as surely the trouble lies with the inductive confirmation of (P1). Once again, I suggest that the underlying issue is explanation. Even if (P1) does qualify as a derived law, as long as what explains it is (P*) – or (P*)'s lawhood – there is no interesting sense in which dualists can be said to have a distinctively temporal problem. They would have a temporal problem only if (P1) would be true even if (P*) were false.

This is a conclusion defenders of the causal argument should accept. They should say that (Closure) is different from (P1) of the temporal argument, in that while (P1) is true merely because (P*) is, (Closure) is true for a deeper reason – it has a different kind of explanation. And this difference in explanation, they should say, is why dualists have a causal problem but no temporal problem. However, as I have argued, this otherwise attractive line is unavailable to physicalists, because to take (Closure) to have a different explanation from (P1) would be to posit (Closure)'s explanatory overdetermination, which must be avoided. The further point to add here is that if causal argument defenders were to adopt the position staked out in the present objection, this possibility of distinguishing between the explanations of (Closure) and (P1) would be closed off to

them. The present objection thus seems to guarantee the parity of the causal and temporal arguments for physicalism.

Let's summarize the two challenges posed for the causal argument in this section. The first is to say why there is a causal problem *for dualists*, given that the causal situation dualists confront looks exactly like the causal situation other theorists confront, notably including type identity theorists. The second is to say why there is a *causal problem* for dualists, given that the causal situation dualists confront looks exactly like the non-causal situations they confront, notably including their temporal situation. My own view is that these challenges cannot be met. Dualists have no causal problem, just as type identity theorists have no causal problem, and just as dualists have no temporal problem. What's more, I say, all physicalists must agree, or else they will be forced to posit explanatory overdetermination.

8.

Objection. The truth of (Closure) follows directly from our best physical theories – that is, without those theories first entailing (P*) – and thus is directly explained by those theories. Some have thought that (Closure) follows directly from the conservation laws of physics, although maybe the idea that physics explains (Closure) could be developed in other ways.[32] At any rate, if physics directly explains (Closure), then what the explanatory exclusion objection entails is that no physicalist reasonably can accept physics. But surely this conclusion is a reductio of the explanatory exclusion objection.

[32] For discussion, see Papineau (2001). If (Closure) deductively follows from the conservation laws it might qualify as a derived law, although in that case it would not be derived from (P*), and so last section's arguments would not apply.

Reply. I grant this would be a reductio. I deny that (Closure) directly follows from our best physical theories,[33] but if I were proven wrong about this, I would give up the explanatory exclusion objection before I held that physicalists could not reasonably accept physics. In that case, however, I instantly would be provided with a *new* objection to the causal argument. We can set up the new objection and clarify its relation to explanatory exclusion by posing a dilemma to causal argument defenders.

Either positing explanatory overdetermination is objectionable or it is not. If it is objectionable, as I have argued, then the causal argument for physicalism does not work, since it requires that we embrace such overdetermination. If explanatory overdetermination is not objectionable, on the other hand, then we have no reason to accept the anti-overdetermination premise, (P3), of the causal argument. So either way, the causal argument for physicalism does not work.

Connecting this back to the present objection, the crucial point is that even if it could be established that physics directly explains (Closure), this does not blunt the dilemma, but rather impales defenders of the causal argument on its second horn. For, what the present objection suggests is not that (Closure) is not explanatorily overdetermined, but rather that such overdetermination must somehow be okay – it's somehow okay that (Closure) is true both because (P*) is, as physicalism requires, and also for whatever reasons our physical theories give in explaining it. Fine. In order to avoid the reductio, I'm willing to grant that perhaps we will need to say such overdetermination is somehow okay. In that case, though, why should we accept (P3) of the causal argument? In short: Why is the explanatory overdetermination posited here by

[33] For discussion, see Lowe (2003). I think (Closure) might indirectly follow from physics: physics might give us good abductive reason to accept (P*), which in turn entails (Closure). But again, this is useless to defenders of the causal argument.

causal argument proponents supposed to be permissible, while the explanatory overdetermination posited by dualists who reject (P3) objectionable?

What the dilemma brings to light is that physics taken by itself has no chance of saving the causal argument here. For, although our warrant for (Closure) plausibly is due to physics in some way or other, our warrant for (P3) surely is not – it's due to a priori or commonsense reflection. To blunt the dilemma, here is what causal argument defenders really need to do. They need to distinguish between two kinds of explanatory overdetermination, the objectionable kind and the permissible kind. They then need to show that the overdetermination they themselves posit is permissible, while the overdetermination that dualists who reject (P3) posit is objectionable. Obviously, the challenge will be to distinguish between objectionable and permissible overdetermination in a principled way. Perhaps the basis for a principled distinction would be that the overdetermination posited by the causal argument defender is non-causal while the overdetermination posited by the dualist is causal. It will not do to leave things at that, however. Causal argument defenders need to say why the causal/non-causal difference makes a difference here. I cannot envision how such an account would go. Every plausible attempt I can think of to say what's wrong with causal overdetermination generalizes to non-causal explanatory overdetermination, as we saw in the case of the analysis offered back in §3.

I close this section by warding off a potential misconception. In so far as my attack on the causal argument has shifted in this section from the (Closure) premise (P1) to the anti-overdetermination premise (P3), this might seem to be a weakening of my objection, given that (P3) seems so obviously true to many philosophers. But that would

be a misunderstanding. I still think the real problem with the causal argument is (Closure) –explanatory overdetermination is always objectionable, whether causal or non-causal. It is only when we are forced to suppose that the causal argument's brand of explanatory overdetermination is somehow okay, as the present objection has us assume, that I come to doubt the case for (P3). Think of it like this. Ordinarily, I would be happy to grant (P3) to defenders of the causal argument without much discussion. But, once such defenders start drawing subtle distinctions between permissible and impermissible forms of overdetermination, and the causal argument itself comes to turn entirely on these distinctions, then the case for (P3) needs to come in for serious scrutiny.[34]

9.

Objection. The explanatory exclusion objection is dialectically flawed. To illustrate the flaw, suppose we are presently neutral on the question of physicalism. The inductive case for (Closure) is presented, and we are persuaded. Perhaps this inductive case does presuppose (Closure)'s lawhood, or at least some explanation independent of (P*), but we do not yet accept (P*) and so we do not yet have a problem with explanatory exclusion. Next we are persuaded by the cases for both (P2) and (P3) of the causal argument. We haven't yet deductively inferred (P*), and so again, there is no exclusion problem yet. Finally we infer (P*). At this point it's too late to run the explanatory

[34] There is one more point to make here. As observed back in note 8, Kim in his most recent presentation of the causal exclusion argument has suggested that because causation and supervenience are both metaphysical determination relations, they give rise to the risk of explanatory exclusion – a risk that can be avoided only by supposing that a given cause of an effect also causes the effect's supervenience base. It seems plausible that causal argument proponents will need to appeal to something like Kim's principle if they are to block all forms of dualism. If so, then the permissible/objectionable distinction between types of explanatory overdetermination relations will not map perfectly onto the non-causal/causal distinction.

exclusion objection, for we already have come to accept physicalism on the basis of the causal argument.[35]

Reply. Having arrived at the conclusion of (P*) and thus now finding ourselves positing the explanatory overdetermination of (Closure), we should realize in retrospect that either (Closure) must not be inducible after all – we were wrong to think it was – or else the case for (P3), the anti-overdetermination premise, must be wrong. It's last section's dilemma all over again. Arriving at the argument's conclusion undercuts our warrant for either (Closure) or (P3), which in turn undercuts our warrant for the conclusion. In effect, what this shows is that the causal argument is *self-defeating*: accepting its conclusion undercuts the warrant for its premises.

10.

Suppose everything I have argued here is right. The causal argument for physicalism is then in bad shape. I do not think this weakens physicalism, however. In fact, I suspect that abandoning the causal argument could strengthen the physicalist position.[36] Those who deny the existence of Santa Clause get by just fine without a causal argument, as do those who deny that there are unicorns, or an elephant in the room, or that the 2003 Kansas City Royals won the World Series. Those who deny the existence of non-physical events should be no worse off.

REFERENCES

Armstrong, D. (1983). *What Is a Law of Nature?* Cambridge University Press.

[35] Versions of this objection have been presented to me by G., J. and K. [Names removed]

[36] For one thing, it could cast a different light on epiphenomenal dualism, of the sort defended by Jackson (1982), Chalmers (1996), and Kim (2005). From the standpoint of the explanatory exclusion objection, I would suggest that Cartesian interactionist dualism is the most plausible form of dualism there is.

Chalmers, D. (1996). *The Conscious Mind*. Oxford University Press.

Davidson, D. (1970). "Mental Events," in L. Foster and J. W. Swanson (eds.), *Experience and Theory*. Duckworth.

Davidson, D. (1995). "Laws and Cause," *Dialectica*, 49: 263-280.

Dretske, F. (1977). "Laws of Nature," *Philosophy of Science*, 45: 173-181.

Foster, J. (1991). *The Immaterial Self*. Routledge.

Goodman, N. (1983). *Fact, Fiction, and Forecast*. Harvard University Press.

Harman, G. (1965). "The Inference to the Best Explanation," *Philosophical Review*, 74: 88-95.

Jackson, F. (1982). "Epiphenomenal Qualia," *Philosophical Quarterly*, 32: 127-136.

Kim. J. (2005). *Physicalism, Or Something Near Enough*. Princeton University Press.

Levine, J. (2001). *Purple Haze: The Puzzle of Consciousness*. Oxford University Press.

Lewis, D. (1966). "An Argument for the Identity Theory," *Journal of Philosophy*, 63: 17-25.

Lowe, E. J. (2003). "Physical Causal Closure and the Invisibility of Mental Causation," in S. Walter and H. D. Heckmann (eds.), *Physicalism and Mental Causation*. Imprint Academic.

Melnyk, A. (2003). *A Physicalist Manifesto: Thoroughly Modern Materialism*. Cambridge University Press.

Montero, B. (1999). "The Body Problem," *Noûs*, 33: 183-200.

Papineau, D. (2001). "The Rise of Physicalism," in C. Gillett and B. Loewer (eds.), *Physicalism and its Discontents.* Cambridge University Press.

Papineau, D. (2002). *Thinking about Consciousness*. Oxford University Press.

Salmon, W. (1989). "Four Decades of Scientific Explanation," *Minnesota Studies in the Philosophy of Science*, 13: 3-217.

Sider, T. (2003). "What's So Bad About Overdetermination?", *Philosophy and Phenomenological Research*, 67: 719-726.

Smart, J. J. C. (1959). "Sensations and Brain Processes," *Philosophical Review*, 68: 141-156.

Tooley, M. (1977). "The Nature of Laws," *Canadian Journal of Philosophy*, 7: 667-698.

Tye, M. (1995). *Ten Problems of Consciousness*. MIT Press.

Wright, C. (2003). "Some Reflections on the Acquisition of Warrant by Inference," in S. Nuccetelli (ed.) *New Essays on Semantic Externalism, Skepticism, and Self-Knowledge*. MIT Press.

Wright, C. (2004). "Warrant for Nothing (And Foundations for Free?)", *Proceedings of the Aristotelian Society*, Supplementary Volume 78: 167-212.