The Philosophical Quarterly Vol. 65, No. 261
ISSN 0031-8094

2015
doi: 10.1093/pq/pqv030
Advance Access Publication 26th April 2015

THE CAUSAL CLOSURE PRINCIPLE

By Sophie Gibb

In the mental causation debate, there is a common assumption that interactive dualism is false because of the principle of the causal closure of the physical domain. However, this paper argues that recent advances in metaphysics—more specifically, in the philosophy of causation—reveal a serious, general flaw in contemporary formulations of this principle.

Keywords: causal closure principle, interactive dualism, powers theory of causation, causal sufficiency, double prevention, enabling events.

The idea that mental entities are causally relevant to bodily behaviour is central to our pretheoretical conception of human agency. Hence, for example, my desire to raise my hand seems to be causally responsible for my hand's raising. It is precisely because I had this desire that my hand raised. And, had I not had this desire, in normal circumstances, I wouldn't have raised my hand. Despite the initial plausibility of the claim that mental entities are causally relevant in the physical domain, identifying a relationship between mental and physical entities that is consistent with this causal interaction and, yet, independently plausible is one of the perennial problems in the philosophy of mind. Solving this problem is the main focus of the contemporary mental causation debate.

In this debate, there has been a general assumption that, given that mental entities are causally relevant to physical entities, some version of physicalism—the doctrine that all entities are identical with, or, in some sense 'nothing over and above' physical entities—must be correct. Interactive dualism—regardless of whether we are here concerned with an interactive substance dualism or an interactive anti-physicalist property dualism—is, it is supposed, untenable.[1]

[1] Contrary to the substance dualist, the anti-physicalist property dualist maintains that mental properties are non-physical properties of the body. However, despite maintaining that mental

This is primarily because of the widespread acceptance in the mental causation debate of the causal closure of the physical domain, which I shall here initially formulate as the principle that all physical effects have sufficient physical causes.[2] However, recent advances in metaphysics, particularly in the philosophy of causation, draw this principle into question and, hence, challenge the consensus. The aim of this paper is to demonstrate one central reason why.

I. THE CAUSAL CLOSURE ARGUMENT AND CAUSAL CLOSURE PRINCIPLES

The problem that the causal closure principle raises for interactive dualism can be set out as follows:

1. *Relevance:* Some mental events are causally relevant to physical effects.
2. *Closure*: All physical effects have sufficient physical causes.
3. *Exclusion*: There is no systematic causal overdetermination.

Therefore, mental events (that are causally relevant to physical effects) are identical with physical events.

To explain this argument: In accordance with *Relevance*, say that $\underline{M}$ is a mental event and that it is a sufficient cause of physical event $\underline{E}$. Given *Closure*, $\underline{E}$ must have a sufficient physical cause, '$\underline{P}$'. The mere combination of *Relevance* and *Closure* does not entail that $\underline{M}$ must be a physical event, for *Closure* is consistent with the possibility that physical effects have both sufficient physical causes and sufficient non-physical causes. It is the role of *Exclusion* to rule this possibility out. Turning to *Exclusion*, to give a standard example of causal overdetermination, say that two shots are independently fired and that both bullets reach the victim at the same time. Given that each bullet striking was causally sufficient for the victim's death, the death was causally overdetermined by the strikings. *Exclusion* permits isolated cases of causal overdetermination, but rules out events being *systematically* causally overdetermined. Hence, it cannot be the case that *whenever* $\underline{M}$ causes $\underline{E}$, $\underline{P}$ also causes $\underline{E}$, where it is such that if one of the two events $\underline{M}$ and $\underline{P}$ had not existed, the other would have sufficed, in the circumstances, to cause $\underline{E}$. But it is precisely this systematic causal overdetermination that the combination of *Relevance* and *Closure* seems

properties are properties of the body, it is anti-physicalist in its claim that mental properties are strongly emergent entities which bestow full-blooded independent causal powers on their bearers. For the purpose of this paper, I take a neutral stance between these versions of dualism.

[2] To give some examples of those who consider that physicalism must be correct in virtue of the causal closure of the physical domain, see Kim (2010: ch. 4), Loewer (2001), Melnyk (2003: ch. 6), Papineau (1993: ch.1, 2001) and Smith & Jones (1986: ch. 4). For this formulation of the principle, see Papineau (1998: 375). Note, what I refer to as 'the causal closure principle', Papineau refers to as 'the principle of the completeness of physics'.

to give rise to. The problem is removed if, contrary to interactive dualism, $\underline{M}$ is identical with $\underline{P}$.[3]

Now the above formulation of the causal closure principle is but one of several to be found in the contemporary mental causation debate. To illustrate the range of ways in which this principle has been formulated, I here provide a representative—although by no means exhaustive—list:

(1) All physical effects have sufficient physical causes (Papineau 1998: 375).
(2) All physical effects are due to physical causes (Spurrett and Papineau 1999: 25).
(3) Every physical event has a physical cause which is enough to bring it about, given the laws of physics (Crane 2001: 45).
(4) All physical effects have complete physical causes ('complete' in the sense that those causes on their own suffice by physical law to fix the chances of those effects) (Papineau 1993: 22).
(5) All physical effects are determined or have their chances determined by prior physical causes according to physical law (Crane 1995: 6).
(6) If a physical event has a cause at $\underline{t}$, then it has a physical cause at $\underline{t}$ (Kim 2005: 15).
(7) Every physical event contains only other physical events in its transitive causal closure (Lowe 2000: 581).
(8) Physical events do not have non-physical causes (Smith and Jones 1986: 66).
(9) Any cause of a physical event is itself a physical event—that is, no non-physical event can be a cause of a physical event (Kim 2005: 50).

Clearly, (1)–(9) are not all equivalent. Some of the formulations appeal to the notion of a 'sufficient physical cause'. Some are probabilistic in nature. Some refer to the laws of physics. Furthermore, many of these formulations differ in strength. As E. J. Lowe has argued, one central problem for proponents of the causal closure argument is providing a formulation of the causal closure principle that is of the *correct* strength. (In particular, see Lowe 2000.) On the one hand, it must not be so weak that it renders the causal closure argument invalid. On the other hand, it must not be so strong that it lacks empirical (or, indeed, metaphysical) support. Equally, it must not be so strong that, to

[3] I take causes and effects to be Kimean events. According to Kim, an event is the exemplification of a property by a substance at a time. Hence, a mental event is the exemplification of a mental property by a substance at a time and a physical event is the exemplification of a physical property by a substance at a time. Given Kim's account of events, two events are identical if and only if they involve the same property, substance and time. It follows that both substance and property dualism entail a dualism with regard to mental and physical events and, hence, that both positions directly conflict with the conclusion of the causal closure argument. Note, however, that the assumption that the causal relata are Kimean events is not essential to the causal closure argument. (See, for example, Heil and Mele 1993 for further defence of this claim). Nor is it essential to the argument that this paper presents.

provide an argument for it, one must first smuggle in physicalist assumptions, and, hence, provoke the complaint that it begs the question against interactive dualism. In light of these considerations, Lowe argues that (1) is in fact too weak. If causation is transitive, then (1), in combination with the other two premises of the causal closure argument, does not entail the desired physicalist conclusion. Given the transitivity of causation, a physical event *would* have a sufficient physical cause if it had a sufficient mental cause which in turn had a sufficient physical cause. Hence, the combination of (1) with *Exclusion* would be compatible with a dualist model of psychophysical causal relevance which held that neural events caused bodily movement via mental causal intermediaries (Lowe 2000: 575–6).[4] While (1) is too weak, according to Lowe (8) and (9) are too strong—indeed, they are so strong that they render *Exclusion* redundant within the causal closure argument. Lowe (2000) proposes that any argument that could be provided for a causal closure principle of this strength will inevitably beg the question against interactive dualism.[5]

It might, however, be assumed that a causal closure principle could be formulated which, unlike (1), when combined with the other premises of the causal closure argument, is strong enough to rule out the causal relevance of non-physical events in the physical domain, but which, unlike (8) and (9), is not so strong as to be implausible. Indeed, one of the other causal closure principles that I have listed might be thought to satisfy both of these constraints. However, I consider that contemporary metaphysics provides a new and altogether different reason for thinking that one cannot formulate a causal closure principle that satisfies both of these constraints, and certainly that none of the formulations that I have listed succeed in doing so. My argument is as follows: Despite the many differences in the way that the causal closure principle has been formulated in the mental causation debate, a common underlying assumption unites most (if not all) of its proponents. This is the assumption that every physical event that has a cause has a sufficient cause, or, at least, that every physical event that has a cause has a cause that is sufficient to fix its chances. Depending on the way that the causal closure principle is formulated, this assumption is either directly entailed by the causal closure principle or must be a further hidden premise in the causal closure argument for it to be valid. But this assumption is implausible according to several recent accounts of the causal relation. Given these accounts, current formulations of the causal closure principle therefore must be abandoned. Furthermore, attempts to address this problem by revising the causal closure principle are problematic, as there is a reason to think that the resulting causal closure principles will lack empirical support. This paper will develop and defend these claims.

[4] (2)–(5) also arguably face this problem.

[5] Although, note that the interactive dualist model that I propose is compatible with the acceptance of (8) and (9). (See Section IV.)

II. CAUSAL INSUFFICIENCY

For the moment, I shall put issues regarding the transitivity of causation to one side and adopt formulation (1) of the causal closure principle, as it makes the problem that I wish to raise most transparent. (I shall return to each of the other formulations in later parts of this paper.) (1) makes explicit appeal to the notion of a 'sufficient cause'. A cause is sufficient for its effect in the sense of being enough to bring the effect about. In other words, the existence of the cause guarantees or ensures the existence of the effect. Indeed, the causal closure principle is sometimes simply formulated as the principle that all physical effects are *due to* physical causes or that all physical effects have physical causes which are *enough to bring them about*. [See (2) and (3).] These alternative formulations of the causal closure principle are simply re-expressions of the idea that all physical effects have *sufficient* physical causes. As an aside, note, it is rarely, if ever, the case that the existence of a single physical event ensures the existence of some other physical event. Rather, the thought is that for every physical effect there is a complex cause which is a combination of wholly physical events, and it is this combination of physical events that collectively ensures the existence of the effect, or which, in other words, is collectively causally sufficient for the effect.

The basic problem that I wish to raise for (1) is as follows: Clearly, it can be true that all physical effects have *sufficient physical causes* only if it is in the first place true that all physical effects have *sufficient causes*. However, there are good reasons to think that not all effects—and, more specifically, that not all physical effects—have sufficient causes. That is, in other words, there are good reasons to think that not all (physical) effects have a set of causes whose existence collectively ensures their existence.

To isolate the problem that I'm raising for the claim that all effects have sufficient causes, it is first important to say what the problem is not.

II.1 Quantum mechanics

Any issues with the claim that all effects have sufficient causes are commonly thought to arise as a result of quantum mechanics—given the indeterministic nature of quantum mechanics, causes cannot always be sufficient for their effects. To attempt to avoid any conflict with quantum mechanics, probabilistic versions of the causal closure principle have been advanced. [For example, see (4) and (5).] My objection to the claim that all effects have sufficient causes is unrelated to this matter. I consider that there are metaphysical reasons to question this claim quite aside from those issues arising as a consequence of the indeterministic nature of quantum mechanics. Furthermore, these reasons provide grounds not only for questioning the claim that every effect has a

sufficient cause, but, equally, grounds for questioning probabilistic versions of this claim such as, for example, the claim that every effect has a cause that suffices to fix its chances.

II.2 Background conditions

Some philosophers distinguish between the event that is *the* cause of a particular event and those events which are mere background conditions necessary for the relevant causal relation to take place. Thus, a match being struck is *the* cause of its lighting, while the presence of oxygen and the dryness of the match are mere background conditions that are necessary for the match striking to bring about the match lighting. The match being struck counts as *the* cause of the match lighting because it is the event that deviated from the natural course of events and, hence, which 'triggered' the match lighting. If one takes this distinction to be an objective one, then one might be tempted to conclude that causes are rarely, if ever, sufficient for their effects. The match striking—*the* cause of the match lighting—is not sufficient for the match lighting. A host of other events which are mere background conditions must also be in place for the match striking to bring about the match lighting.

I would reject any such claim. While causation is a wholly objective relation, the distinction between *the* cause and its background conditions is a subjective one. Which event is singled out as *the* cause and which are counted as background conditions is determined by our interests. Hence, regarding the above case, it is easy to conceive of a case in which the presence of oxygen would instead be counted as *the* cause of the match's lighting. For example, imagine a laboratory experiment in which a match is repeatedly struck in a chamber with no oxygen in it. It is only upon the addition of oxygen that the match lights. In such a case, plausibly we would want to count the presence of oxygen as *the* cause of the match lighting. What is singled out as *the* cause does not reflect any special causal role that the event has. The striking of the match, the presence of oxygen, and the dryness of the match are all contributory causes of the match lighting, and it is merely the case that we relegate some of these events to the background for pragmatic purposes.

In suggesting that not all effects have sufficient causes what I'm claiming is that the *complete* cause (that is, the combination of *all* of the contributory causes) is not always sufficient for its effect.

II.3 The powers theory of causation

The reason why I dispute the claim that all effects have sufficient causes is because, in recent metaphysics, several of the accounts of the causal relation that have been offered lead to the abandonment of this idea. To explain why I shall first focus on the powers theory of causation which, as a result of a resurgence

of power ontologies, is increasingly popular in contemporary metaphysics. According to it, dispositions or powers (I use these terms interchangeably) provide the basis for an account of the causal relation. Two general claims are central to any version of the powers theory of causation.

First, properties bestow irreducible powers on their bearers. Given this stance, all intrinsic properties are dispositional, where a property is dispositional if, solely in virtue of being characterized by it, a substance possesses a certain power.[6] Hence, for example, because of its fragility, a porcelain vase is disposed to break if it is dropped on a hard surface. This power to break is built into some property of the vase, and it is because it is characterized by this property that the vase is disposed to break when dropped on a hard surface.

Secondly, causation is the exercise of these powers—or, in other words, causation occurs when these powers manifest themselves. There are several different ways of developing this claim. Here, I shall summarize C. B. Martin and John Heil's, according to which causation is the mutual manifestation of reciprocal disposition partners. [See, for example, Martin (2008) and Heil (2003).] A particular manifestation of a disposition usually depends on other dispositions being present. For example, the vase's breaking when it is dropped on a hard surface depends, not only on the fragility of the vase, but also on the hardness of the surface. Hence, the vase's breaking is a manifestation, not only of the vase's fragility, but also of the surface's hardness. The vase's fragility and the surface's hardness are 'reciprocal disposition partners'. The breaking of the vase is their 'mutual manifestation'. Causation is the mutual manifestation of reciprocal disposition partners. That is, the vase's breaking just is the mutual manifestation of the vase's fragility and the surface's hardness. Note that, returning to my earlier point about background conditions, given this account of causation, there is no distinction between the cause and its background conditions. According to it, the lighting of the match is a *mutual* manifestation of the force of the match, its dryness, the roughness of the surface upon which the match is struck and the oxygen's power to make materials combust.

This is merely a brief sketch of one version of the powers theory of causation, but it is not necessary to go into any further detail here. The crucial point is that, given this account—and, indeed, given any version of the powers theory of causation—the existence of all of the contributory causes of an event (that is, the existence of its complete cause) is not always sufficient for the existence of that event. The central examples that demonstrate this point involve cases of double prevention.

[6] Note this does not commit one to the thesis that properties are exhausted by their dispositionality. It is also consistent with, for example, the thesis that every property is both dispositional *and* qualitative. [For the first approach towards powers, see Shoemaker (1980). For the second, see Martin (2008) and Heil (1998, 2003, 2012).] For the purpose of this paper, I take a neutral stance between these approaches.

Double prevention occurs when an event that would prevent another event from having a certain effect is itself prevented from doing so. To give an example of double prevention, a barrier is placed in front of a porcelain vase, but the barrier is wired up to an explosive device which will blow the barrier up if a button on the device is pressed. Normally, if a rock is thrown at the vase, the barrier would prevent the rock from coming into contact with and, hence, breaking the vase. But, if the device's button is pressed, this destroys the barrier, hence, allowing the rock to hit the vase and break it. The pressing of the device's button is a 'double preventer' event—the barrier would have prevented the rock from breaking the vase, but is itself prevented from doing so by the pressing of the button. Upon careful reflection of paradigmatic causal sequences, it is clear that double prevention is a common phenomenon in the physical world.[7]

Double prevention can be explained by the powers theory of causation. A disposition's manifestation usually depends on the presence of certain dispositions. But it can also depend on the absence of certain dispositions, as one disposition may be disposed to prevent the manifestation of another. Disposition $\underline{A}$ may be disposed to prevent the manifestation of disposition $\underline{B}$ in one of two ways: either the manifestation of $\underline{A}$ results in the loss of $\underline{B}$ or it merely blocks $\underline{B}$'s manifestation. Hence, returning to our example, the pressing of the device's button prevents the mutual manifestation that is the barrier's solidity and the rock's momentum and hardness. It does so because it brings about the destruction of the barrier and, hence, the loss of the barrier's powers. Hence, this is an example of the first type of prevention. Alternatively, if the device's button is not pressed, the barrier would prevent the rock from coming into contact with the vase. In the language of the powers theory of causation, the solidity of the barrier prevents the mutual manifestation that is the vase's fragility and the rock's momentum and hardness. This is an example of the second type of prevention—the vase does not cease to be fragile, but because of the solidity of the barrier the vase is prevented from manifesting its fragility. Regarding *double* prevention, a disposition that is disposed to prevent the manifestation of another disposition, is prevented from doing so by the presence of a further disposition. Thus, in the example that I have just given, the solidity of the barrier is disposed to prevent the rock from breaking the vase, but the barrier's disposition is itself prevented from being manifested by the pressing of the device's button.[8]

Now, say that the button is pressed. Hence, the barrier is destroyed and, consequently, the rock hits the vase. The resulting shattering of the vase is a mutual manifestation of disposition partners which include the rock's

[7] For this example of double prevention, see Gibb (2013). For further examples of double prevention, see Hall (2004), Lewis (2004), and Schaffer (2000).

[8] For a fuller description of this account of the powers theory of causation's analysis of prevention and double prevention, see Gibb (2013).

momentum and the vase's fragility. But what is of crucial significance for the purpose of this paper is that the device's button being pressed cannot be a contributory cause of the breaking of the vase according to the powers theory of causation. More generally, given the powers theory of causation, a double preventer event cannot be a cause of the event that it has prevented from being prevented. In brief, this is because absences cannot be causes according to the powers theory of causation, for an absence cannot bear powers and hence cannot be disposed to act in any way. Given that absences are not causes, there cannot be a chain of unbroken causation from the double preventer event to the event that it has prevented from being prevented. Hence, in our example, the pressing of the button causes the destruction of the barrier, but given the powers theory of causation, the barrier's destruction cannot in turn be a cause of the vase's breaking, for this is really just to say that the *absence* of the barrier is a cause of the vase's breaking. Therefore, given the powers theory of causation, the pressing of the button is not a cause of the vase's breaking, for there is not a chain of unbroken causation from the pressing of the button to the breaking of the vase. More generally, given the powers theory of causation, double prevention is not causation.[9]

One important consequence of this—which seems obvious and yet which has gone largely unnoticed—is that, given the powers theory of causation, events that are prevented from being prevented never have sufficient causes. Consider every single one of the contributory causes of the breaking of the vase in our example—the rock's momentum, the vase's fragility, etc. The total combination of these events—the complete cause of the vase's breaking—is not enough to bring about the vase's breaking. There is a further event—the pressing of the button—that is not a cause of the vase's breaking, but which must take place for it to break. Hence, the vase's breaking—a paradigm example of a physical effect if ever there was one—does not have a sufficient cause. Nor is it the case that the vase's breaking has a cause that is even sufficient to fix its chance of occurrence. The existence of every single one of the contributory causes of the vase's breaking is not sufficient to fix the chance of the vase breaking, for this is, in part, determined by whether the button is pressed. Clearly, these points generalize to every case of double prevention.[10]

[9] I take the claim that double prevention is not causation according to the powers theory of causation to be uncontroversial amongst those who defend this theory of causation. See Mumford & Anjum (2009) and Gibb (2013) for more detailed accounts of why double prevention is not causation according to the powers theory of causation.

[10] Mumford & Anjum (2011) have also recently argued that a proper understanding of causation reveals that it is false that all effects have sufficient causes. (Note their discussion is firmly focused on the philosophy of causation, and they do not relate their claims to the mental causation debate.) However, our arguments are distinct and independent, as are the conclusions that we draw from them. Their argument is based on the notion of antecedent strengthening and their conclusion is that *no* effect has a sufficient cause. Clearly, the notion of antecedent strengthening

II.4 Other accounts of causation that have this consequence

The conclusion that double prevention is not causation is not unique to the powers theory of causation. As others have argued, it is a consequence of those theories of causation that hold that causation is the transfer of some quantity (such as energy or momentum) between cause and effect, of process theories of causation such as Dowe's (2000) and Salmon's (1984), and of Ehring's (1997) account of causation in terms of trope fission and fusion.[11]

I shall not rehearse these arguments here, but shall simply observe that what is common to all of the theories of causation just listed is the requirement that there is some kind of local connection between a cause and its immediate effect (e.g., a transfer of energy, the transmission of a mark, the fusion of two tropes), and that, furthermore, as a consequence of the connection that they propose, absences are ruled out as causes (a nothingness has no energy to transfer, no mark to transmit, no trope for another to fuse with). As a result, according to these theories of causation, double prevention cannot be causation. It follows that, according to each of these theories, not every effect has a sufficient cause (or one that is sufficient to fix its chances).

Of course, not all theories of causation share the conclusion that double prevention is not causation. The theories of causation that I have been considering all conceive of causation as production. That is, according to them, $\underline{C}$ is a cause of $\underline{E}$ just in case $\underline{C}$ in some sense produces $\underline{E}$. This production approach to causation can be contrasted with a dependence approach. If causation is conceived of as dependence, then $\underline{C}$ is a cause of $\underline{E}$ just in case $\underline{E}$ suitably depends on $\underline{C}$. The central example of this approach is the counterfactual theory of causation, which understands the dependence relation to be a counterfactual one. However, one might instead take the dependence relation to be a nomological or probabilistic one. The division between the production approach and the dependence approach is a core one in the contemporary debate about the nature of causation and reflects a fundamental disagreement about the features of causation. If causation is conceived of as dependence rather than production, then one will automatically count double prevention as causation because a double preventer event and the event that it prevents from being prevented will inevitably stand in the relevant dependence relationship.[12] By contrast, standardly, those accounts that understand causation as production will dismiss double prevention as causation for the reasons explained above.

plays no role in my argument and the conclusion of my argument is merely that *some* effects lack sufficient causes. To engage in a discussion of Mumford and Anjum's argument would take me too far from the topic of this paper. For a discussion of this argument and a persuasive objection to it, see Lowe (2012).

[11] Philosophers that have argued that one or more of these theories lead to the rejection of double prevention as causation include Dowe (2000), Psillos (2010), and Schaffer (2000).

[12] For further defence of this claim, see Hall (2004) and Psillos (2010: 8).

[I say 'standardly' because some of those who defend a production approach depart from their central thesis in order to allow causation by omission. See, for example, Fair (1979: 246–8) who advances an energy transference theory of causation, but adopts a counterfactual dependence approach in the case of omissions.]

Thus, the contemporary debate about causation reveals that to accept that all effects have sufficient causes (or ones that suffice to fix their chances) one must take a stance on the nature of the causal relation. The mental causation debate cannot simply ignore those cases in which physical effects do not have sufficient causes, given the widespread occurrence of double prevention in the physical domain. Nor can it ignore those theories of causation that have the consequence that not every effect has a sufficient cause. Indeed, according to Jaegwon Kim, it is precisely causation as production, as opposed to causation as dependence, that those in the mental causation debate should be concerned with. To quote Kim: 'Causation as generation, or effective production and determination, is in many ways a stronger relation than mere counterfactual dependence, and it is causation in this sense that is fundamentally involved in the problem of mental causation' (2005: 18).[13]

III. ENABLING EVENTS

We have seen that, according to several contemporary theories of causation, double prevention is not causation. Consequently, according to these theories, an event that is prevented from being prevented does not have a sufficient cause. For such an event to be brought about, in addition to all of its contributory causes (its complete cause), a further event (the double preventer event) must also exist. This further event is what I shall call an *enabling event*. Hence, putting issues regarding the indeterministic nature of quantum mechanics to one side, it is the existence of the complete cause together with the existence of the enabling event that is sufficient for the existence of an event that is prevented from being prevented.

I reserve the term 'enabling event' for an event which, on a particular occasion, does not cause an event, but which enables it to be caused. Expressed slightly differently, an enabling event is an event that provides the correct structure for a particular causal relation to take place. Hence, given standard production accounts of causation, a double preventer event is an enabling event because it is not a cause of the event that it prevents from being prevented,

[13] Note that Kim is here referring to the problem of mental causation generated by the causal closure principle (see Kim 2005: 15)—his point being that discussions of this problem, and the causal closure principle more specifically, should be embedded in a production account of causation.

but one which enables the event to be caused by preventing an event from preventing it being caused.[14]

Now, although enabling events are not causes of the events that they enable to be caused, I would suggest that the former events are causally relevant to the latter events, and not just in a merely explanatory sense. In causal situations in which enabling events are involved, for the effect to be brought about, in addition to the complete cause, a further event must occur whose role is to enable the causal relation to take place. Clearly, this role is an objective one, not a merely explanatory one. Furthermore, I can see no good reason for thinking that the role of an enabling event is any less important than the role of a cause in accounting for an effect's existence. Indeed, many of the things that can be said of causes can also be said of enabling events. Hence, returning to our example of double prevention: (1) Just as the throwing of the rock (a cause) is required for the breaking of the vase, so is the pressing of the button (an enabling event); (2) The breaking of the vase stands in a relation of counterfactual dependence to both the throwing of the rock and the pressing of the button; (3) If asked to explain why the vase broke, an explanation which only referred to the throwing of the rock and neglected to mention the pressing of the button would be incomplete. (Although, of course, in certain situations our explanations may relegate enabling events to the background, just as, in certain situations, they may relegate some of the contributory causes of an event to the background).

My distinction between causes and enabling events might remind the reader of Fred Dretske's distinction between 'triggering causes' and 'structuring causes', for the role of both my enabling events and Dretske's structuring causes is supposed to be that of providing the correct structure for a particular causal relation to take place. It is therefore worth saying a little about how our accounts differ.

To give one of the examples that Dretske appeals to to explain his distinction, one puts yeast in dough so that the heat of the oven will cause the dough to rise. The presence of the yeast in the dough is, according to Dretske, a 'background condition' for one thing (the heat of the oven) to cause another (the raising of the dough) (Dretske 1992: 39). Whatever event caused this background condition to exist is a 'structuring cause' of the dough's rising. The structuring causes of $\underline{E}$ are, in other words, the causes (the triggering causes) of those background conditions which are required for $\underline{C}$ to cause $\underline{E}$. [In particular, see Dretske (1994: 206) and Dretske (2004: 170).] By contrast, a triggering cause is an event that triggers the causal process that results in $\underline{E}$. Hence, for example, my turning the oven on is a triggering cause of the dough's rising, for it is what caused the oven to become hot which caused the dough to rise.

[14] Whether there are any other enabling events besides double preventer events is a further issue. Certainly, they provide the most obvious example of enabling events.

It should be clear from this brief summary of Dretske's position, that my 'enabling events' are not equivalent to Dretske 'structuring causes'. Indeed, given my stance on background conditions, Dretske's distinction between structuring and triggering causes is not one that I could plausibly accept. Dretske acknowledges that some philosophers consider that the background conditions for $\underline{C}$ to cause $\underline{E}$ are in fact partial causes of $\underline{E}$, and that they, together with $\underline{C}$ (itself just another partial cause), cause $\underline{E}$ (1992: 39). Hence, for example, the presence of yeast in the dough is a partial cause of the dough's rising, alongside the heat of the oven. As made clear earlier, this is the stance that I adopt. Furthermore, I consider that, given this stance, any distinction between the structuring causes of $\underline{E}$ and the triggering causes of $\underline{E}$ collapses. Both are really just (triggering) causes of different partial causes of $\underline{E}$. The distinction between structuring 'causes' and triggering 'causes' might well be relevant at the level of explanation, but it is not at the level of causation.

Having made clear the distinction between causes and enabling events, we are now in a position to reappraise the various formulations of the causal closure principle.

IV. CAUSAL CLOSURE PRINCIPLES RECONSIDERED

Returning to the list of causal closure principles, given standard production accounts of causation, formulations (1)–(5) are false. Each of these formulations entails either that every physical effect has a sufficient cause or that every physical effect has a cause that is sufficient to fix its chances. But, given standard production accounts, this is incorrect. As we have seen, the case of a physical event that is prevented from being prevented demonstrates this point.

Formulations (6)–(9) do not entail that every physical effect has a sufficient cause or a cause that is sufficient to fix its chances. Hence, one might wonder how my argument affects them. Well, take formulation (7). Assuming (7), the causal closure argument is as follows:

i. *Relevance*: Some mental events are causally relevant to physical effects.
ii. *Closure* (7): Every physical event contains only other physical events in its transitive causal closure.
iii. *Exclusion*: There is no systematic causal overdetermination.

Therefore, mental events (that are causally relevant to physical effects) are identical with physical events.

Note that by the 'transitive causal closure' of an event ($\underline{Y}$) Lowe means the set of events 'which includes every event which stands in the ancestral of the "immediate cause" relation' to $\underline{Y}$. That is, the set of events which includes the immediate causes of $\underline{Y}$, the immediate causes of those causes, the immediate

causes of *those* causes. . . and so on' (Lowe 2000: 581). Given (7), where *Y* is a physical event, every event in this set must be physical.

Reflecting upon the distinction between causes and enabling events allows one to recognize that (7) is not strong enough to render the causal closure argument valid. If a mental event need not cause a physical event in order to be causally relevant to it, then clearly the combination of (7) with *Relevance* and *Exclusion* does not entail that mental events (that are causally relevant to physical effects) are identical with physical events.[15] Consequently, to make this argument valid, a fourth premise must be added to it, namely:

iv. *Causal Irrelevance*: If event *X* is not a cause of event *Y*, then *X* is causally irrelevant to *Y*.

It might be assumed that *Causal Irrelevance* is obviously true. This assumption seems entirely reasonable if every effect has a sufficient cause (or if every effect has a cause that is sufficient to fix its chances)—if an effect has a cause that is enough to bring it about (or, in indeterministic cases, one that is sufficient to fix its chances of being brought about), then what further causal role is left for an event that is not its cause to play? However, the plausibility of *Causal Irrelevance* becomes altogether less clear if, for the reasons demanded by standard production accounts of causation, one rejects the claim that every effect has a sufficient cause or one that is sufficient to fix its chances. I have argued that, given standard production accounts, it is important to recognize a distinction between causes and enabling events. Enabling events are not causes of the events that they enable to be caused, but they are causally relevant to them. On this basis, *Causal Irrelevance* is false— *X* does not have to be a cause of *Y* to be causally relevant to *Y*, for *X* could be an event that enables *Y* to be caused.

Formulations (6), (8) and (9) all face exactly the same problem as formulation (7). They are all too weak, as the combination of any one of these causal closure principles with the other two premises of the causal closure argument does not rule out the causal relevance of non-physical events in the physical domain. Hence, the causal closure argument is rendered invalid. To make it valid, a further premise (*Causal Irrelevance*) must be added which limits causally relevant events to those that are causes. However, this additional premise is implausible, given the distinction between causes and enabling events.

Returning to formulation (7), let us flesh out this claim with an example. Call the event that is neuron 1 firing in Fred's brain '*n1*', the event that is neuron 2 firing in his brain '*n2*', and the event that is his hand's moving '*b1*'. Assuming the powers theory of causation, say that the firing of neuron 1 is disposed to make neuron 2 fire, which is disposed to make certain muscles in Fred's body

[15] As with the rest of this paper, I should emphasize that 'causal relevance' is not to be interpreted as mere 'causal explanatory relevance'.

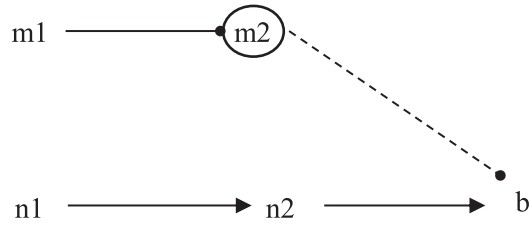

Figure 1.

contract and, thereby, make his hand move. For simplicity, assume that no other dispositions are required for these manifestations. Thus, *n1* causes *n2* and *n2* causes *b1*. Now let us add that Fred's desire to keep his body still (call this mental event '*m2*') is disposed to prevent *n2* from causing *b1*. But that his conflicting, stronger desire to move his hand ('*m1*')—say due to a bad case of pins and needles—prevents *m2*'s manifestation. It might be that Fred retains the desire to keep his body still, but the manifestation of this desire is blocked by his overriding desire to move his hand. Or, it might be that gaining the desire to move his hand causes him to lose the desire to keep his body still. Regardless of which is the case, *m1* prevents *m2* from preventing *n2* causing *b1*.[16] This causal structure is represented diagrammatically in Fig. 1. In Fig. 1, a solid line ending in an arrow depicts a causal relation; a solid line ending in a dot depicts an inhibitory connection; a broken line ending in a dot depicts an inhibitory connection that failed to occur and a circle around a letter signifies the non-existence of the relevant event.

In this example, *m1* prevents *m2* from preventing *n2* causing *b1*. Consequently, *n2* is able to cause *b1*. In such a case, contrary to *Causal Irrelevance*, *m1* is causally relevant to *b1*, despite not being a cause of *b1*, as it enables *b1* to be caused. Now let us say that, contrary to the conclusion of the causal closure argument, *m1* is non-physical. The non-physicality of *m1* does not require one to reject either (7) or *Exclusion*, despite *m1*'s causal relevance to physical event *b1*. It does not lead to the rejection of (7), as *b1* contains only physical events (*n2*, *n1*, . . .) in its transitive causal closure. It does not lead to the rejection of *Exclusion*, as *m1* is not a cause of *b1*, and, hence, does not threaten to causally overdetermine it. Hence, the causal closure argument fails.

This particular dualist model of psychophysical causal relevance is one that I have developed and defended from a metaphysical, empirical and phenomenological point of view in various papers. [See, for example, Gibb (2013, 2015).] I shall therefore not attempt to do so here. However, I shall end this section by responding to one pressing objection to it.[17]

[16] For this example, see Gibb (2013).
[17] I am very grateful to an anonymous referee for raising this objection.

Consider what would be required for *m1* to prevent *m2* from preventing *n2* causing *b1*. If *m1* really does prevent *m2* from preventing *n2* causing *b1*, then must it not be the case that had *m1* not occurred (or been weaker than *m2*) then *m2* would have caused Fred's body to keep still? (Call this physical event '*b2*'.) Unless *m2* would have caused *b2* (in *m1*'s absence), how could *m2* have prevented *n2* causing *b1* in the first place? If *m2* is strong enough to prevent *n2* causing *b1*, this is because it is strong enough to cause *b2* instead.

Clearly, if this is correct, it gives rise to a problem for this dualist model. In such a case, *b2* (a physical event) would have *m2* (a mental event) in its transitive causal closure. The combination of (7) and *Exclusion* therefore entails that *m2* must be identical with a physical event. And, if it can be demonstrated that at least some mental events must be identical with physical events, then this provides us with very good inductive grounds for concluding that probably all mental events are identical with physical events. Hence, the double prevention model collapses. For this dualist model to succeed, *all* mental events must at most be double preventers (and, hence, not causes) of physical events.

I think that this objection is right in one important respect, but wrong in another. Let me begin by explaining what I consider to be wrong with it. In *m1*'s absence, *m2* prevents *n2* from causing *b1*. As observed in Section II, disposition *A* may be disposed to prevent the manifestation of disposition *B* in one of two ways: either the manifestation of *A* results in the loss of *B* or it merely blocks *B*'s manifestation. In line with this, let us say that, in *m1*'s absence, *m2* prevents *n2* from causing *b1* by bringing about the loss of *n2* (the event that is neuron 2 firing in Fred's brain)—that, in *m1*'s absence, *m2* stops neuron 2 from firing.

For *m2* to prevent *n2* from causing *b1*, must it also be the case that *m2* has the power to cause *b2*? I don't see why this must be the case. It might seem to be a further reasonable requirement that for *m2* to prevent *n2* from causing *b1*, *m2* must be causally relevant to *b2*. But, given the distinction between causes and enabling events, it does not follow from this that *m2* has to have the power to cause *b2*. Hence, let us instead say that it is some further neurological event, *n3*, that has the power to cause *b2*, but that *n3*'s causing *b2* would be prevented by *n2*. In *m1*'s absence, *m2* prevents *n2* causing *b1* by bringing about the loss of *n2*. And, by doing precisely this—that is, by bringing about the loss of *n2*— *m2* enables *n3* to cause *b2*. That is, *n2* would have prevented *n3* from causing *b2*, but *m2* prevents it from doing so. Hence, *m2* is causally relevant to *b2*, but not because it causes *b2*, but rather because it enables *b2* to be caused. It enables *b2* to be caused simply by bringing about the loss of *n2*.

But we don't need the claim that *m2* has the power to cause *b2* to arrive at the problem. Given the mere fact that, in *m1*'s absence, *m2* brings about the loss of *n2*, we arrive at the problem. For *m2* to bring about the loss of *n2*—for it to stop neuron 2 from firing— *m2* must presumably have to cause something within the physical domain. Hence, we revert back to the problem that this

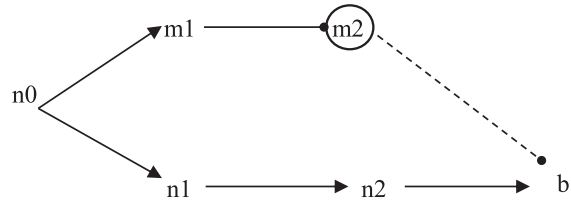

Figure 2.

objection raises. If *m2* does cause some physical event, then the combination of (7) and *Exclusion* entails that *m2* must be identical with a physical event.

However, I provide a detailed response to this problem in Gibb (2013: 202–10), which is further discussed in Gibb (2015). The response can be summarized as follows: The basic problem is that in *m1*'s absence, *m2* would prevent *n2* causing *b1*, and, to do this, *m2* must cause some physical event. If *m2* causes some physical event, then the combination of (7) and *Exclusion* entails that *m2* must be identical with a physical event. What this problem demonstrates is that for the double prevention model of psychophysical causal relevance to work, in any case where there is *m2* and *n2*, *m1 must* be there to prevent *m2* from preventing *n2* causing *b1*. This would be the case if the existence of some event in the chain of neurological events that caused *n2* entailed the existence of *m1*. Hence, for example, if the existence of *n1* entailed the existence of *m1*. And, this entailment relation between *n1* and *m1* would be explained if whatever neurological event that caused *n1* also caused *m1*. Where *n0* is some further neurological event, this proposal is set out in Fig. 2.

Elsewhere, I explore this claim, defending it from both a metaphysical and an empirical point of view (Gibb 2013: 205–10. See also Gibb 2015.) Unfortunately, I do not have the space to provide a further defence of it here. For the purpose of this paper, the crucial point is that if this claim were correct, then it blocks the above objection—that is, the objection that, in *m1*'s absence, *m2* would prevent *n2* from causing *b1* and, to do so, *m2* must cause some physical event. It is true that if *m2* ever actually *did* prevent *n2* from causing *b1*, this would (given *Exclusion*) violate (7). But, given the causal system of events that is presented in Fig. 2, *m2* never actually does prevent *n2* from causing *b1*. *m2* is disposed to prevent *n2* from causing *b1*, but given the presence of *m1*, it is never able to manifest this disposition. (For further discussion of this particular point, see Gibb 2013: 206–7.)

V. REVISED CAUSAL CLOSURE PRINCIPLES

I have argued that given standard production accounts of causations, the formulations of the causal closure principle that have been listed are all either too

strong or too weak. Formulations (1)–(5) are too strong, because they mistakenly assume that every physical effect has a sufficient cause or that every physical effect has a cause that is sufficient to fix its chances. Formulations (6)–(9) are too weak, as the combination of any one of these closure principles with the other two premises of the causal closure argument does not rule out the causal relevance of non-physical events in the physical domain. Hence, given (6)–(9), the causal closure argument is invalid. To make it valid, a further premise must be added (*Causal Irrelevance*). But this additional premise is implausible given the distinction between causes and enabling events. I shall end this paper by briefly considering one way that I envisage proponents of the causal closure argument will attempt to respond to this argument. It is as follows: although this paper might have demonstrated that, given standard production accounts of causation, formulations (1)–(9) are all either too weak or too strong, this does not call for the rejection of the causal closure argument, but instead the revision of its causal closure principle.

I agree that it is possible to formulate a causal closure principle which meets both of the constraints that have been set out. That is, it is possible to formulate a principle which, on the one hand, does not entail that every physical effect has a sufficient cause (or, one that is sufficient to fix its chances), but which, on the other hand, when combined with the other two premises of the causal closure argument, does yield the conclusion that mental events (that are causally relevant in the physical domain) are identical with physical events.

I shall begin by considering a formulation of the causal closure principle provided by Montero which I have not yet discussed. At first glance, it might appear to meet both of these constraints, but it in fact fails to do so. According to it:

(10) Every physical event that has a sufficient cause has a sufficient physical cause (Montero 2003: 174).[18]

(10) certainly meets the first constraint that I have set out. That is, it does not entail that every physical effect has a sufficient cause. This is because it merely claims that every physical event that *has* a *sufficient cause* has a sufficient physical cause. However, upon inspection, it fails to meet the second constraint.

[18] I am very grateful to an anonymous referee for drawing my attention to this particular formulation of the closure principle. For a different version of this kind of formulation, see Yates (2009), who offers the following formulation: 'At every time at which a physical event has a sufficient cause, it has a sufficient physical cause' (p. 115). Note, as Montero observes, quantum effects fall outside the scope of (10) as a quantum effect, such as the radioactive decay of a particle, lacks a sufficient cause. Montero suggests that this potential issue should be addressed by 'interpreting "sufficient cause" as a cause that either fully determines its effect or a cause that fully determines the chances of its possible effects'. My discussion of (10) shall assume this interpretation of a 'sufficient cause'.

That is, it is too weak to render the causal closure argument valid. Moreover, like formulations (6)–(9), the additional premise that is required to make the argument valid is implausible in light of the distinction between causes and enabling events.

(10) is too weak to render the causal closure argument valid precisely because it is silent about those physical events that lack sufficient causes. The combination of (10) with *Relevance* and *Exclusion* does not allow one to rule out the possibility that non-physical events are causally relevant to those physical events that lack sufficient causes. Hence, given this formulation of the causal closure principle, for the causal closure argument to be valid it must be supplemented with the following premise:

Every physical event *that a mental event is causally relevant to* has a sufficient cause.

Clearly, this further premise is a version of the claim that every physical event effect has a sufficient cause, but one which is limited to a subset of physical effects—namely, those physical effects that mental events are causally relevant to. (Which is precisely the subset of physical effects that we are interested in for the purpose of the mental causation debate.) As should be clear from our previous discussion, this additional premise is false if mental events are enabling events in the physical domain. If the causal role of mental events in the physical domain is not to cause physical events but to enable physical events to be caused, then *no* physical event that a mental event is causally relevant to will have a sufficient cause—as explained in Section III, an event which requires the existence of an enabling event to be brought about lacks a sufficient cause. Consequently, (10) fails to meet the challenge that this paper has presented.

However, one can certainly create causal closure principles that do meet this challenge. Modifying (7), one proposal might be the following:

(11) Every physical event contains only other physical events in its causal history. (Where the 'causal history' of event $\underline{Y}$ is defined as including not only every event which stands in the ancestral of the 'immediate cause' relation to $\underline{Y}$, but also every event that enables $\underline{Y}$ to be caused.)

Unlike formulations (1)–(5), (11) does not entail that every physical effect has a sufficient cause (or a cause that is sufficient to fix its chances). However, unlike formulations (6)–(10), (11) leads to the rejection of the claim that mental events could be enabling events in the physical domain, unless they are physical.

But a causal closure principle must have some measure of empirical support. The problem is that a causal closure principle that is specifically designed to overcome the problem that I have raised, such as (11), will be precisely of the sort that it is very hard to muster empirical evidence for. This is because, if the

causal role that a mental event plays in the physical domain is to enable one physical event to cause another physical event by preventing a mental event from preventing it, then this causal role will plausibly be one that empirical science is blind to.

The idea that certain causal roles that mental events might play in the physical domain are invisible to science is one that Lowe advances in his defence of his own interactive dualist model of psychophysical causal relevance. (See, for example, Lowe 2008.) According to Lowe, the causal role of mental events in the physical domain is that of making the fact that a causal tree of neural events converge upon a particular bodily movement non-coincidental. Although there are many differences between Lowe's model of psychophysical causal relevance and the one that has been suggested here, there is one crucial similarity. Both accounts depart from standard dualist models of psychophysical causal relevance in denying that the causal role of mental events in the physical domain is that of initiating any neural event (or set of neural events) in the chains of neurophysiological causation that give rise to bodily movement. Rather than suggesting that a mental event is ever the cause of any neural event, Lowe's proposal is that a mental event is causally responsible for the *fact* that a maze of neural events converge upon a particular bodily movement. According to my suggestion, rather than a mental event ever being the cause of any neural event, mental events enable neural events to give rise to particular bodily movements.

Because he is not claiming that a mental event is causally relevant in the physical domain in virtue of causing any physical event or set of physical events, Lowe argues that the causal role that he provides mental events in the physical domain will be invisible to any scientist who was to examine the situation by empirical means. As this causal role is not that of initiating any single physical event or set of physical events, there will no gaps in the chains of neurophysiological causation for science to discover. Hence:

> Any scientist who was to examine that situation by empirical means, but who was restricted by his means of investigation to observing only purely physical events and causal relationships, would quite naturally come to the conclusion that the physical event [. . .] had a complete and wholly physical causal explanation, in terms of its immediate causes [. . .] and their antecedent physical causes [. . .]. (Lowe 2008: 74)

Consequently, the empirical findings provided by neurophysiologists and neuropsychologists will not allow one to discriminate between a wholly physicalist account of psychophysical causal relevance and the interactive dualist account that Lowe offers.

The same can, I think, be said if mental events play the causal role in the physical domain that I have described. Hence, returning to Fig. 2, an empirical examination of the causal chain of neurological events that results in $\underline{b_1}$ would not reveal the causal role that $\underline{m_1}$ has played in bringing about $\underline{b_1}$, no matter

how closely one looks. This is because there are no gaps in this causal chain of neurological events for empirical science to find—the complete cause of *b1* (that is, the combination of *all* of the contributory causes of *b1*) is wholly physical. *n2*, which is the complete cause of *b1*, has a complete wholly physical cause (*n1*), which in turn also has a complete wholly physical cause (*no*). It would therefore be fair, although incorrect, given this causal structure, for the scientist to conclude that non-physical events play no role in bringing about *b1*.[19]

The central aim of this paper has been to demonstrate that, given standard production accounts of causation, current formulations of the closure principle are unsatisfactory. I have also indicated why I consider that potential new, strengthened versions of the closure principle which aim to address the problem that this paper raises for the causal closure argument will be empirically dubious. My aim has not been to convince the reader that the double prevention model of psychophysical causal relevance is correct. However, the findings of this paper do suggest that it is doubtful that this dualist model can be defeated simply by appealing to the causal closure argument.[20]

REFERENCES

Crane, T. (1995) 'The Mental Causation Debate', *Proceedings of the Aristotelian Society*, 69: 211–53.
—— (2001) *Elements of Mind*. Oxford: OUP.
Dretske, F. (1992) *Explaining Behaviour*. Cambridge: MIT Press.
—— (1994) 'Reply to Slater and Garcia-Carpintero', *Mind and Language*, 9/2: 203–8.
—— (2004) 'Psychological Versus Biological Explanations of Behavior', *Behavior and Philosophy*, 32/1: 167–77.
Dowe, P. (2000) *Physical Causation*. Cambridge: CUP.
Ehring, D. (1997) *Causation and Persistence*. Oxford: OUP.
Fair, D. (1979) 'Causation and the Flow of Energy', *Erkenntnis*, 14/3: 219–50.
Gibb, S. (2013) 'Mental Causation and Double Prevention', in S. Gibb, E. J. Lowe and R. Ingthorsson (eds) *Mental Causation and Ontology*, 193–214. Oxford: OUP.
—— (2015) 'Physical Determinability', *Humana Mente*, 29.
—— (forthcoming) 'Defending Dualism', *Proceedings of the Aristotelian Society*.
Hall, N. (2004) 'Two Concepts of Causation', in J. Collins, N. Hall and L. Paul (eds) *Causation and Counterfactuals*, 181–276. Cambridge, MA: MIT Press.
Heil, J. (1998) *Philosophy of Mind: A Contemporary Introduction*. London: Routledge.
—— (2003) *From an Ontological Point of View*. Oxford: OUP.
—— (2012) *The Universe as We Find It*. Oxford: OUP.
Heil, J. and Mele, A., eds (1993) *Mental Causation*. Oxford: Clarendon Press.
Kim, J. (2005) *Physicalism, or Something Near Enough*. Princeton: Princeton University Press.
—— (2010) *Philosophy of Mind*. Westview Press: Boulder CO.

[19] For a more detailed discussion of these claims, see Gibb (forthcoming).

[20] I am grateful to the John Templeton Foundation for funding the 'Durham Emergence Project', of which this paper is a consequence. I am also grateful to two anonymous referees for their helpful comments, and to participants at workshops held at Durham University and the University of Macerata where I presented earlier drafts of this paper. Most of all I am indebted to E. J. Lowe, conversations with whom helped to shape the ideas that I present in this paper.

Lewis, D. (2004) 'Causation as Influence', in J. Collins, N. Hall and L. Paul (eds) *Causation and Counterfactuals*, 75–106. Cambridge, MA: MIT Press.
Loewer, B. (2001) 'From Physics to Physicalism', in C. Gillet and B. Loewer (eds) *Physicalism and Its Discontents*, 37–56. Cambridge: CUP.
Lowe, E. J. (2000) 'Causal Closure Principles and Emergentism', *Philosophy*, 57/294: 571–85.
—— (2008) *Personal Agency: The Metaphysics of Mind and Action*. Oxford: OUP.
—— (2012) 'Mumford and Anjum on Causal Necessitarianism and Antecedent Strengthening', *Analysis*, 72/4: 731–5.
Martin, C. B. (2008) *The Mind in Nature*. New York: OUP.
Melnyk, A. (2003) *A Physicalist Manifesto: Thoroughly Modern Materialism*. Cambridge: CUP.
Montero, B. (2003) 'Varieties of Causal Closure', in S. Walter and H. Heckmann (eds) *Physicalism and Mental Causation: The Metaphysics of Mind and Action*, 173–87. Exeter: Imprint Academic.
Mumford, S. and Anjum, R. (2009) 'Double Prevention and Powers', *Journal of Critical Realism*, 8/3: 277–93.
—— (2011) *Getting Causes from Powers*. Oxford: OUP.
Papineau, D. (1993) *Philosophical Naturalism*. Oxford: Blackwell.
—— (1998) 'Mind the Gap', *Philosophical Perspectives*, 32/12: 373–88.
—— (2001) 'The Rise of Physicalism', in C. Gillet and B. Loewer (eds) *Physicalism and Its Discontents*, 3–36. Cambridge: CUP.
Psillos, S. (2010) 'Causal Pluralism', in R. Vanderbeeken and B. D'Hooghe (eds) *Worldviews, Science and Us: Studies of Analytical Metaphysics*, 131–51. Singapore: World Scientific Publishers.
Salmon, W. (1984) *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton University Press.
Schaffer, J. (2000) 'Causation by Disconnection', *Philosophy of Science*, 67/2: 285–300.
Shoemaker, S. (1980) 'Causality and Properties', in P. van Inwagen (ed.) *Time and Cause*, 109–35. Dordrecht: D. Reidel Publishing Co.
Smith, P. and Jones, O. (1986) *The Philosophy of Mind*. Cambridge: CUP.
Spurrett, D. and Papineau, D. (1999) 'A Note on the Completeness of "Physics"', *Analysis*, 59/1: 25–9.
Yates, D. (2009) 'Emergence, Downwards Causation and the Completeness of Physics'. *Philosophical Quarterly*, 59/234: 110–31.

Department of Philosophy, University of Durham, UK