

Robert K. Garcia

Closing in on Causal Closure

Abstract: *I examine the meaning and merits of a premise in the Exclusion Argument, the causal closure principle that all physical effects have physical causes. I do so by addressing two questions. First, if we grant the other premises, exactly what kind of closure principle is required to make the Exclusion Argument valid? Second, what are the merits of the requisite closure principle? Concerning the first, I argue that the Exclusion Argument requires a strong, 'stringently pure' version of closure. The latter employs two qualifications concerning the physical sufficiency and relative proximity of the physical cause required for every physical effect. The second question is addressed in two steps. I begin by challenging the adequacy of the empirical support offered by David Papineau for closure. Then I assess the merits of 'level' and 'domain' versions of stringently pure closure. I argue that a domain version lacks adequate and non-question-begging support within the context of the Exclusion Argument. And I argue that the level version leads to a puzzling metaphysics of the physical domain. Thus, we have grounds for rejecting the version of closure required for the Exclusion Argument. This means we can resist the Exclusion Argument while avoiding the implausible implications that come with rejecting one of its other premises. That is, because there are grounds to reject causal closure, one can reasonably affirm the non-overdeterminative causal efficacy of conscious mental states while denying that the latter are identical with physical states.*

This article finds its genesis in a remark by E.J. Lowe: 'One might have hoped for more exactitude and agreement amongst physicalists when it comes to the formulation of a principle so central to their position' (Lowe, 2000, p. 574). The principle Lowe bemoans is the so-

Correspondence:

Robert K. Garcia, Department of Philosophy, Texas A&M University, College Station, TX 77843, USA. Email: robertkgarcia@gmail.com

called causal closure of the physical domain. As a provisional statement, the principle says that all physical effects have physical causes. The closure principle is central to debates over the metaphysics of consciousness and especially to the debate over mental causation and physicalism. Jaegwon Kim, for example, states that the principle ‘seems minimally required of any serious form of physicalism’ (Kim, 1995, p. 290). Indeed, David Papineau takes causal closure to be ‘the crucial premise’ in a kind of argument responsible for ‘the rise of physicalism’ in the twentieth century (Papineau, 2000, p. 177).¹ Papineau has in mind the Exclusion Argument, whose basic form is as follows:

- (1) All physical effects have physical causes.
- (2) Mental events have physical effects.
- (3) Physical effects are not systematically overdetermined.
- (4) Thus, mental events must be identical with physical events.

If sound, this argument establishes a conclusion of great significance for our understanding of consciousness, namely, that conscious mental states are identical with physical states. Many find this conclusion hard to accept and wish to resist the Exclusion Argument by denying one or more of the premises. Each premise, however, is such that denying it has implications for the metaphysics of consciousness.

Consider premises (2) and (3). On the one hand, denying (2) commits us to epiphenomenalism — the thesis that mental states have no physical effects. On this view, George Santayana was right: ‘Conscious will is a symptom, not a cause; its roots... are... material’ (Santayana, 1930, p. 121). If epiphenomenalism is true, consciousness has no causal influence on the physical world. Thus, the vivid sensation of a toothache does not cause you to wince, the faint aroma of coffee does not cause you to look for Starbucks, and the weight of deep existential angst does not cause you to put down *Being and Nothingness*. For most, this is too incredible to believe (McLaughlin, 1994). On the other hand, denying (3) commits us to the possibility that mental causation (e.g. the vivid sensation of a toothache causing you to wince, etc.) systematically involves ‘two independent causal lines — one from your brain, another from your soul — converging on the same effect’ (Robb and Heil, 2013). For many, this would be an ‘intolerable coincidence’ (Melnyk, 2003, p. 291). In short, there is

[1] This kind of argument goes by many names: The Causal Argument for Physicalism (Papineau, 2000, p. 180), The Causal Exclusion Argument (Horgan, 1997, p. 169), The Overdetermination Argument (Sturgeon, 2000, p. 121), The Materialist Argument (Cockburn, 2001, p. 64), The Master Argument (Graham, 1998, p. 229), etc.


much to be said on behalf of both (2) and (3). This creates considerable pressure to choose between accepting physicalism (4) and rejecting causal closure (1). In this article I consider the prospects of the latter. If there are grounds to reject (or withhold belief in) causal closure, then one can reasonably affirm the non-overdeterminative causal efficacy of conscious mental states while denying that the latter are identical with physical states.

I will proceed by taking up two questions. First, if we grant the other premises of the Exclusion Argument, exactly what kind of closure principle is required to make the argument valid? Second, what are the merits of the requisite closure principle?

Narrowing Down Causal Closure

With respect to the first question, I will argue that the Exclusion Argument requires a closure principle that employs two qualifications, one concerning the *relative proximity* of the physical cause for every physical effect, and another concerning the *sufficiency* of the physical cause for every physical effect.

To appreciate the need for the first qualification, consider the following example in which 'E's stand for events in a causal chain: at time t_1 , E_1 causes E_2 and E_2 causes E_3 ; at time t_2 , E_3 causes E_4 ; and at time t_3 , E_4 causes E_5 . Where an arrow represents a causal relation, this can be diagrammed as follows:



We can now discern two related distinctions. The first distinction is between a proximal cause and a distal cause. The second distinction compounds the proximal/distal distinction with a temporal/ontological one. Thus, assuming the transitivity of causation, in the above example the following causal claims are true:

- E_4 is a *temporally proximal* cause of E_5 : E_4 occurs at a time immediately prior to the time E_5 occurs.
- E_3 is a *temporally distal* cause of E_5 : E_3 occurs at a time prior, but not immediately prior, to the time E_5 .

- E_2 is an *ontologically proximal* cause of E_3 : E_2 directly (i.e. without an intermediate event) causes E_3 , but need not be temporally prior to E_3 — the events may be simultaneous.
- E_1 is an *ontologically distal* cause of E_3 : E_1 indirectly (i.e. through an intermediate event) causes E_3 , but need not be temporally prior to E_3 — the events may be simultaneous.

To demonstrate the relevance of these distinctions, we will consider two points made by Lowe (2000) in his discussion of the Exclusion Argument. First, Lowe points out that certain versions of closure do not preclude an irreducibly mental event from non-overdeterminatively causing a physical event. For example, consider one of Papineau's versions of closure: 'Every physical effect has a sufficient physical cause' (Papineau, 1998, p. 375). Assuming that causality is transitive, this principle is satisfied when its required sufficient physical cause either is temporally distal, ontologically distal, temporally proximate, or ontologically proximate. So, as Lowe points out, this version of closure does not preclude irreducible mental events from non-overdeterminatively causing physical events. This can be seen in the above diagram if we let E_3 and E_5 be physical events — a brain event and grimacing, respectively — and let E_4 be an experience of sharply felt pain. In this case, a conscious mental event (E_4) non-overdeterminatively causes a physical event (E_5) even though the latter has a sufficient physical cause (E_3).

Of course, other versions of the closure principle feature in exclusion arguments. Consider a version offered by Sophie Gibb: 'At every time at which a physical event has a cause, it has a sufficient physical cause' (Gibb, 2013, p. 2). This version is *not* satisfied if E_3 and E_5 are physical events and E_4 is (say) an experience of pain. However, it *is* satisfied where E_1 and E_3 are brain events, E_2 is an experience of pain, and E_4 is grimacing. And this is Lowe's second point: this version of closure 'neglects the possibility of *simultaneous causation*' and thus does not preclude mental states from having physical effects (Lowe, 2000, pp. 575–6).

To be sure, there are versions of closure that explicitly employ distinctions sufficient to preclude the above kind of mental causation.² More often, however, the distinctions are implicit and shown in the use and interpretation of closure principles. According to Kim, for example, closure is violated if both mental and physical events occur as links in the same causal chain. In such a case, closure would be violated since there would be 'physical occurrences that cannot be

[2] See the discussion in Lowe (2000).

causally explained by invoking physical antecedents and laws alone' (Kim, 1993, p. 188). Thus, the intent seems to be to require that the physical cause *not* cause the effect via a non-physical intermediary. Closure, so intended, claims that every physical effect has as its aetiology a causal chain devoid of any non-physical links. Thus, we can distinguish between closure principles that require each physical effect to have an ontologically proximal, or *direct*, physical cause, and those principles that require each physical effect to have *some* physical cause (whether direct or not). For convenience, we may name the former version 'pure closure' and the latter 'impure closure'. What Lowe's points indicate is that the Exclusion Argument requires a pure closure principle, one which requires a direct physical cause for every physical effect.

But to do its job in the Exclusion Argument, the closure principle needs a further qualification concerning the *sufficiency* of the physical cause. We can discern this qualification by considering an ambiguity in the following version of closure: every physical effect E has a direct physical cause C. There are two ways of construing this version:³

- (A) Every physical effect E has a direct *physical sufficient* cause C, in that C has physical properties $P_1 \dots P_n$, and $P_1 \dots P_n$ together are (at least) an INUS⁴ condition for E.
- (B) Every physical effect E has a direct *physically-sufficient* cause C, in that C has physical properties $P_1 \dots P_n$, and $P_1 \dots P_n$ are together sufficient to bring about E.

An example might help. Suppose C causes E, and C has two properties, M and P, where M is a phenomenal property and P is a physical property. If *C's having P* is necessary but insufficient to bring about E, and *C's having P* and *C's having M* are jointly sufficient to bring about E, then (A) is satisfied but not (B). If *C's having P* alone is sufficient to bring about E, then (B) is satisfied (as well as (A)).

I will characterize the distinction between (A) and (B) as the distinction between requiring a *physical sufficient* cause and requiring a *physically-sufficient* cause, respectively. Earlier we considered one of Papineau's versions of closure: 'All physical effects have sufficient physical causes.' Given his (2000) discussion of 'the completeness of the physical', it is clear that Papineau takes closure to require all

[3] One might also entertain a closure principle on which a physical effect could have as a cause an event whose physical properties were epiphenomenal with respect to the effect. But such a principle can be set aside since it would not help the Exclusion Argument.

[4] This useful acronym is J.L. Mackie's and means 'an *insufficient* but *necessary* part of a condition which is itself *unnecessary* but *sufficient* for the result' (Mackie, 1999, p. 414).

physical effects have *physically-sufficient* causes. And indeed, this distinction expresses a central concern in the debate over mental causation, especially with respect to the causal efficacy of mental content. A closure principle meaning (B) would, in effect, say that all non-physical properties of an event which causes a physical event are epiphenomenal with respect to the latter event. In light of closure, many wonder, with Fred Dretske (1989), whether the mental properties of an event are inert just like the meaning of the sounds made by an opera singer as she shatters a glass with a high-C. Tyler Burge, for example, discusses a version of closure on which ‘physical events can be caused only by virtue of physical properties of other physical events’ (Burge, 1993, p. 99). Similarly, Lynne Rudder Baker understands closure to require that ‘...for any event that has a physical property — whether or not it has mental or other properties — there are sufficient physical conditions for its occurrence and for its having all of its physical properties’ (Baker, 1993, pp. 78–9).

But why should we bother with the distinction between *physical sufficient* causes and *physically-sufficient* causes? For example, why consider a version of closure that only requires (A), *viz.* that the cause has *some* physical property — even if that property alone is insufficient to produce the effect in question? Perhaps an emergentist would answer: it may turn out that while the mental requires *some* physical base and genesis, the mental nevertheless has ‘a life of its own’.⁵ Or, perhaps it will turn out that cases of psycho-physical causation involve physical structuring-causes and mental triggering-causes (to use Dretske’s, 1993, distinction but reverse his application). Or, more generally, it may turn out that certain physical events are caused by a combination of prior mental and physical events that are only *jointly* sufficient (though individually necessary) to bring about the effect in question. Thus, the Exclusion Argument needs a version of closure that requires physical effects to have physically-sufficient causes: for every physical effect E, a direct cause C brings about E entirely in virtue of C’s physical properties. Because this version is more demanding than the pure version, let’s call it a *stringently pure* version of closure. Papineau seems to have this version in mind when he speaks of ‘the completeness of physics’: ‘All physical effects are fully determined by law by a purely physical prior history’ (Papineau, 2000, p. 179).

[5] John Searle (1992, p. 112) thinks this is a naïve suggestion.

Sizing Up Causal Closure

We can now consider our second question, which now becomes: what are the merits of a stringently pure version of closure? I will address this question in two steps. In the first step, I will consider the kind of empirical justification needed for a stringently pure version. Second, I will consider the merits of two possible formulations of a stringently pure version — what I will call *level* and *domain* versions.

As Papineau notes, closure principles are typically based upon principles of energy conservation. However, he also notes that energy conservation is not sufficient to entail (what I'm calling) stringently pure closure. This is because while conservation requires that the basic forces operate conservatively, it does not stipulate that all basic forces are *physical* forces (Papineau, 2000, p. 196). Thus, to derive closure from conservation, an additional argument is necessary. Papineau mentions two such arguments.

Papineau calls the first argument 'the argument from fundamental forces'. The argument claims that 'all apparently special [e.g. mental] forces characteristically *reduce* to a small stock of basic physical forces which conserve energy' (*ibid.*, p. 198). A defender of this argument will begin by pointing to many putatively successful reductions — each being a case where a special type of force has been decomposed into a few fundamental forces. The defender then makes an inductive inference to the general conclusion that all special (e.g. mental) forces will reduce to basic physical ones. However, this inference is premature and potentially question-begging. First, whether and to what extent we have actually had this kind of reductive success is a matter of ongoing and considerable controversy.⁶ Second, even if we have had a relative degree of reductive success, it would be question-begging to assume that putative mental forces are not basic. Indeed, some might argue that since mental events are *prima facie* different than physical events, we have a defeasible reason to think that mental forces are irreducible.

The second argument Papineau mentions is an 'argument from physiology'.⁷ This is also an inductive argument for a general conclusion. It begins with the claim that detailed physiological investigations have failed to discover anything but physical forces. The inference is then drawn that there are no 'special' — e.g. mental — forces at work. Although this argument may be more compelling than

[6] See, for example, Hendry (2006), Robb and Heil (2013), and Stapp (2005).

[7] It has also been called the 'no-causal-gaps' argument (e.g. McLaughlin, 1994).

the first, its conclusion seems premature. To see why, it may be helpful to contrast physics with neuroscience.⁸ Physicists often emphasize how little we know about (say) deep space. Nevertheless, there are physical laws (e.g. conservation) that are sufficiently well established to justify their application to deep space. Similarly, neuroscientists often emphasize how little we know about the brain. However, in this case it is doubtful that there are principles of neuroscience that are sufficiently well established to justify their application to the unknown brain. This suggests that the conclusion of the ‘argument from physiology’ is underdetermined by neuroscience. There may come a point when we know enough about the brain to draw such a conclusion, but until then there is something to be said for suspending judgment. And, for all we know, Colin McGinn (1999) is right — the brain and its processes may *remain* a mystery.

This concludes my brief discussion of the empirical justification for stringently pure closure. Having noted why the empirical case for closure fails to convince, I’ll now distinguish between *domain* and *level* versions of stringently pure closure and argue that neither fares well with respect to the Exclusion Argument.

We can generate two versions of stringently pure closure by introducing a further distinction. The two versions result from distinguishing between what I’ll call *domain* and *level* versions of closure. Some philosophers, for example, define closure at a particular physical level: where ‘ ϕ ’ denotes a particular physical level, the ϕ -*level* is said to be causally closed. In fact, some argue that closure is most plausible when defined at the micro-physical level (where ‘micro’ usually means ‘quantum’).⁹ For them, the micro-physical *level* is causally closed. Lynne Rudder Baker and Scott Sturgeon, for example, have characterized closure in this way. Consider their respective versions:

- ‘Every instantiation of a micro-physical property that has a cause at t has a complete micro-physical cause at t ’ (Baker, 1993, p. 79)
- ‘Every quantum event has a fully disclosive, purely quantum history’ (Sturgeon, 2000, p. 124).

Other philosophers, however, seem to take closure as a more general claim, ranging over the entire physical domain. Spurrett and Papineau, for example, offer the following as a plausible closure principle: ‘All non-mental effects are due to non-mental causes’ (1999, p.

[8] I wish to thank an anonymous referee for suggesting this contrast.

[9] See Merricks (2001), Sturgeon (2000), and Baker (1993). But see also Mixie (1996), who argues that the quantum level is *not* causally closed.

26). In other words, Spurrett and Papineau take the entire physical domain to be closed. Kim also asserts domain-closure:

One way of stating the principle of physical causal closure is this: If you pick any physical event and trace out its causal ancestry or posterity, that will never take you outside the physical domain. That is, no causal chain will ever cross the boundary between the physical and the non-physical. (Kim, 1998, p. 40)

It is important to note that, all other things being equal, domain and level versions are not equivalent. Thus, it is puzzling to find Sturgeon glossing his level version (above) with a statement akin to a domain version: ‘...physics considers itself closed and complete... Non-physical events play no role’ (Sturgeon, 2000; see also his 1999, p. 413).

If Sturgeon takes the versions to be equivalent, he is tacitly relying on certain metaphysical assumptions. For example, for certain ontologies, the domain/level distinction will be a distinction without a difference. For there to be a difference, a ϕ -level version must imply both (i) that there is more to the physical domain than the ϕ -level and (ii) that the remainder does not cause effects at the ϕ -level. That is, for a level version to remain distinct from a domain version, it must assume an ontology in which there *are* irreducible events or objects at non- ϕ levels which do not non-overdeterminatively cause ϕ -level effects. Otherwise, if the ϕ -level is all there is to the physical domain, then to assert ϕ -level closure is to assert domain closure. Thus, in order to ensure that the distinction has a difference, I will assume that level-versions have this ontological import. (I am not, of course, claiming that philosophers who have employed level versions knowingly hold to any such ontology.) For convenience, I will call the non- ϕ level(s) the ψ -level.

In light of their importance, the remainder of this paper will focus on a level-version of closure. Nevertheless, I’d like to mention two features of a domain-version that raise suspicions about the acceptability of its employment in the Exclusion Argument. First, a domain-version of closure is a fully general principle, making a claim about the causal ancestry of any physical event whatsoever, at any moment in the history of the universe, and at any level of physical reality. Given this generality, a domain-version amounts to a metaphysical principle about the world. As such, it is doubtful that a domain-version can be derived from conservation principles alone. Second, because it amounts to a metaphysical claim, the deployment of a domain-version in the Exclusion Argument may seem question-begging to an emergentist. More generally, it seems that adequate

empirical and non-question-begging support for a domain version will be hard to come by. Perhaps it is not surprising, then, that the level-version seem to be heralded as the most plausible way to think about closure.

Thus, since we are considering whether or not there is an acceptable version of closure for the Exclusion Argument, I will consider what I've been calling a stringently pure *level* version of closure. Perhaps it will be useful to reiterate exactly what this cumbersome title amounts to:

- *Stringently Pure Level Closure*: For every ϕ -level physical effect, there is a direct cause which brings about the effect entirely in virtue of the cause's ϕ -level physical properties.

As argued above, a level-version implies that there are ψ -level objects or events which are not ontologically reducible to ϕ -level objects and which do not non-overdeterminatively cause ϕ -level events and objects.¹⁰ This raises a question: do ψ -level objects or events have (irreducible) causal power? If they do not, then they are epiphenomenal with respect to ϕ -level events and objects. If so, then there are grounds for simply eliminating the ψ -level objects and/or events altogether. D.M. Armstrong (1997), for example, has argued that to exist is to have causal power — the so-called 'Eleatic Principle'. Others, such as Trenton Merricks, argue that (alleged) inanimate macro-sized physical objects, such as baseballs, do not really exist since their causal powers are redundant upon the causal powers of their micro-constituents (Merricks, 2001, pp. 56–84). Although these considerations are not the last word, they suggest that a plausible version of stringently pure level closure will allow that there are ψ -level events and objects having irreducible causal power. I'll now consider this view.

Unfortunately, several puzzles emerge if we take stringently pure level closure to allow for ψ -level events to have irreducible causal power. There are three puzzles. The first two turn on a striking feature of closure principles. With rare exceptions (Heil, 1998, p. 23; McLaughlin, 1989, p. 111), extant versions of closure do not require every physical cause to have a *physical effect*. This is because almost all closure principles delimit the *causes* of physical events, but not

[10] Although there may be more than one non- ϕ level, it will be simpler to speak of *the* non- ϕ level since nothing relevant to my argument is lost in the simplification.

their *effects*.¹¹ In what follows, ‘closure’ will denote backwards-closure. Thus, the ψ -level will either be closed or open to causal contributions from the ϕ -level. If the ψ -level is closed, we have the strangest result of all: two levels of ontologically distinct physical events that have never causally interacted, and yet somehow coexist and apparently stand in part–whole relations in space-time. This smacks of a pre-established harmony. This is the first puzzle.

The alternative is better, but still puzzling. Recall that we are assuming that there are ψ -level events and objects having irreducible causal power. Suppose further that the ψ -level is open to causal contributions from the ϕ -level. That is, there are ψ -level events that are caused by ϕ -level events. The puzzle here concerns how this fits with the physical law of energy conservation. If bringing about an effect involves the flow of energy into the effect, then in this scenario energy is leaving the ϕ -level and entering the ψ -level — and *cannot* return. This is especially strange since it seems that the most obvious argument for physical closure is based on the law of energy conservation. Indeed, some take closure to be equivalent to the law (e.g. Mixie, 1996). This is the second puzzle.

All three puzzles are generated by the assumption that stringently pure level closure allows for ψ -level events to have irreducible causal power. The first two puzzles show how this assumption leads to implausible metaphysical implications. The third puzzle is different. Here, the said assumption leads to a metaphysical position that may well be plausible, but is at odds with the Exclusion Argument.

The third puzzle stems from a point made by Merricks. Suppose the ϕ -level is causally closed but the ψ -level has ontologically distinct objects with non-redundant causal powers. Thus, ψ -level objects, in virtue of their physical properties, bring about effects that are not caused by ϕ -level objects working in concert. These physical properties are emergent. However, as Merricks notes:

[N]o self-respecting defender of the Exclusion Argument will allow that a [ψ -level] object has causally non-redundant physical properties... She won't do so because allowing it would push the Exclusion Argument to the brink of begging the question. Allowing it puts her in the awkward position of accepting ‘emergent’ causally non-redundant properties as such, but not ‘emergent’ causally non-redundant mental

[11] The fact that almost all extant closure principles are not ‘bi-directional’ (like Heil’s) raises further questions. First, does evidence for backward-closure also support forward-closure, and vice versa? Second, is it coherent to assert backward-closure but not assert, or even deny, forward-closure? These questions are largely beyond the scope of this article, but I will make a couple of suggestions in what follows.

properties, which would undermine her argument. (Merricks, 2001, p. 140)

It seems, then, that the Exclusion Argument cannot aptly deploy a level version of closure — one allowing for the existence and non-redundant causal power of ψ -level events.

So much for the puzzles. Previously, I indicated why a plausible version of stringently pure level closure will allow for ψ -level objects that have irreducible causal power. We've now seen several puzzling implications of such an allowance. Together, these considerations undermine the plausibility of stringently pure level closure, and, in so doing, pose a challenge for an exclusion argument that deploys such a closure principle.

Conclusion

I've been considering the meaning and merits of the causal closure principle as deployed by the Exclusion Argument. I began by raising two questions. First, if we grant the other premises, exactly what kind of closure principle is required to make the Exclusion Argument valid? Second, what are the merits of the requisite closure principle? Regarding the first question, I argued that the Exclusion Argument requires what I call a stringently pure closure principle: for any physical effect E, a direct cause C brings about E entirely in virtue of C's physical properties. Regarding the second question, I argued that we have grounds to reject a stringently pure closure principle. First, the empirical support offered by Papineau for closure is inadequate. Second, both 'level' and 'domain' versions of stringently pure closure are problematic. The domain version lacks adequate support and is potentially question-begging in the context of the Exclusion Argument. The level version leads to a puzzling metaphysics of the physical domain. Thus, we have grounds for rejecting stringently pure closure. This means we can resist the Exclusion Argument while avoiding the implausible implications that come with rejecting one of the other premises. Rejecting (2) would commit us to epiphenomenalism, which implies that consciousness has no causal influence on the physical world. Rejecting (3) would commit us to the possibility that mental causation involves systematic overdetermination. Thus, because there are grounds to reject causal closure, (1), one can reasonably affirm the non-overdeterminative causal efficacy of conscious mental states while denying that the latter are identical with physical states.¹²

[12] I wish to thank Jaegwon Kim and two anonymous referees for the *Journal of Consciousness Studies* for helpful comments and suggestions for improvement.

References

- Armstrong, D. (1997) *A World of States of Affairs*, Cambridge: Cambridge University Press.
- Baker, L.R. (1993) Metaphysics and mental causation, in Heil, J. & Mele, A. (eds.) *Mental Causation*, Oxford: Clarendon Press.
- Burge, T. (1993) Mind–body causation and explanatory practice, in Heil, J. & Mele, A. (eds.) *Mental Causation*, Oxford: Clarendon Press.
- Cockburn, D. (2001) *An Introduction to the Philosophy of Mind*, New York: Palgrave.
- Dretske, F. (1989) Reasons and causes, *Philosophical Perspectives*, **3**, pp. 1–15.
- Dretske, F. (1993) Mental events as structuring causes of behavior, in Heil, J. & Mele, A. (eds.) *Mental Causation*, Oxford: Clarendon Press.
- Gibb, S. (2013) Introduction, in Gibb, S., Lowe, E.J. & Ingthorsson, R.D. (eds.) *Mental Causation and Ontology*, Oxford: Oxford University Press.
- Graham, G. (1998) *Philosophy of Mind: An Introduction*, 2nd ed., Oxford: Blackwell.
- Heil, J. (1998) *Philosophy of Mind: A Contemporary Introduction*, London and New York: Routledge.
- Hendry, R.F. (2006) Is there downward causation in chemistry?, in Baird, D., Scerri, E. & McIntyre, L. (eds.) *Philosophy of Chemistry: Synthesis of a New Discipline; Boston Studies in the Philosophy and History of Science*, **242**, pp. 173–189.
- Horgan, T. (1997) Kim on mental causation and causal exclusion, *Philosophical Perspectives*, **11**.
- Kim, J. (1993) The non-reductivist's troubles with mental causation, in Heil, J. & Mele, A. (eds.) *Mental Causation*, Oxford: Clarendon Press.
- Kim, J. (1995) Dretske on how reasons explain behavior, in *Supervenience and Mind: Selected Philosophical Essays*, Cambridge: Cambridge University Press.
- Kim, J. (1998) *Mind in a Physical World: An Essay on the Mind–Body Problem and Mental Causation*, Cambridge, MA: MIT Press.
- Lowe, E.J. (2000) Causal closure principles and emergentism, *Philosophy: The Journal of the Royal Institute of Philosophy*, **75** (294), pp. 571–585.
- Mackie, J.L. (1999) Causes and conditions, in Kim, J. & Sosa, E. (eds.) *Metaphysics: An Anthology*, Malden, MA: Blackwell.
- McGinn, C. (1999) *The Mysterious Flame: Conscious Minds in a Material World*, New York: Basic Books.
- McLaughlin, B. (1989) Type epiphenomenalism, type dualism, and the causal priority of the physical, *Philosophical Perspectives*, **3**, pp. 109–135.
- McLaughlin, B. (1994) Epiphenomenalism, in Guttenplan, S. (ed.) *A Companion to the Philosophy of Mind*, Oxford: Blackwell.
- Melnyk, A. (2003) *A Physicalist Manifesto: Thoroughly Modern Materialism*, Cambridge: Cambridge University Press.
- Merricks, T. (2001) *Objects and Persons*, Oxford: Clarendon Press.
- Mixie, J. (1996) The causal closure principle of the physical domain and quantum mechanics, *Dialogos*, **31** (68), pp. 119–125.
- Papineau, D. (1998) Mind the gap, in Tomberlin, J. (ed.) *Philosophical Perspectives*, **12**, Oxford: Blackwell.
- Papineau, D. (2000) The rise of physicalism, in Stone, M. & Wolff, J. (eds.) *The Proper Ambition of Science*, London and New York: Routledge.
- Robb, D. & Heil, J. (2013) Mental causation, in Zalta, E.N. (ed.) *The Stanford Encyclopedia of Philosophy (Spring 2013 Edition)*, [Online], <http://plato.stanford.edu/archives/spr2013/entries/mental-causation/>.

- Santayana, G. (1930) *The Realm of Matter: Book Second of Realms of Being*, London: Constable and Company, Ltd.
- Searle, J. (1992) *The Rediscovery of the Mind*, Cambridge, MA: MIT Press.
- Spurrett, D. & Papineau, D. (1999) A note on the completeness of 'physics', *Analysis*, **59** (1), pp. 25–29.
- Stapp, H. (2005) Quantum interactive dualism: An alternative to materialism, *Journal of Consciousness Studies*, **12** (11), pp. 43–58.
- Sturgeon, S. (1999) Physicalism and overdetermination, *Mind*, **107** (426), pp. 411–432.
- Sturgeon, S. (2000) *Matters of Mind: Consciousness, Reason and Nature*, London and New York: Routledge.

Paper received June 2013; revised October 2013.