

RECENT WORK

Mental Causation

SOPHIE C. GIBB

How could mental entities causally affect, or be affected by, physical entities? Identifying a relationship between mental and physical entities that is both consistent with their causal interaction and independently plausible is one of the perennial problems in the philosophy of mind. In the contemporary mental causation debate, there is not one single problem of mental causation, but several. These include the problem of psychophysical causation generated by the causal closure argument, Davidson's problem of how one can allow psychophysical causal interaction given the anomalism of the mental, Kim's 'pairing problem' for substance dualists, and the problem of the causal relevance of mental content. In this article, I focus on the causal closure argument (because discussion of it has dominated the contemporary mental causation debate) and survey some of the important recent responses to it. I begin by considering the problem of mental causation that this argument presents for dualists and for physicalists (§1). I then survey non-reductive physicalist responses (§2) and dualist responses (§3).

1. The thought that mental causes have physical effects – that our beliefs and desires can give rise to the movement of our bodies – is central to our pre-theoretical notion of human agency. A desire to raise your arm causes you to raise your arm. The desire for the cessation of the pain that you experience when you put your hand too close to a flame causes you to withdraw your hand. However, for there to be psychophysical causation, what must the relation between mental and physical entities be? The problem was once considered to be that of explaining how the mind could bring about bodily behaviour, if mind and body are distinct substances. But few of those in the contemporary mental causation debate would wish to maintain a substance dualism. The popular position is that any so-called mental substance, such as the mind, is identical with the body, or a part of it, such as the brain. However, the causal closure argument presents a problem for every position that maintains that mental and physical properties are distinct, regardless of whether or not mental and physical properties characterize the same substantial particular. Consequently, the problem of psychophysical causation has simply moved from that of explaining how mental substances could be causes

in the physical domain to that of explaining how mental properties could be causally efficacious in the physical domain.

The causal closure argument is also referred to as ‘the argument from causal overdetermination’ or ‘the causal exclusion argument’. Versions of it are defended by Kim (1999, 2005) and Papineau (1993, 2002) among others. It combines the premise that there is psychophysical causation with two further premises to yield the conclusion that mental causes (that have physical effects) are identical with physical causes. It can be formulated as follows:

- (1) Some mental events have physical effects (*Psychophysical Causation*).
- (2) Every physical effect has a sufficient physical cause (*Closure*).¹
- (3) There is no systematic causal overdetermination (*Non-Overdetermination*).
- (4) Therefore, mental events (that have physical effects) are identical with physical events.

To explain the argument: In accordance with *Psychophysical Causation*, say that mental event *m* causes physical event *p2*. Given *Closure*, *p2* must have a sufficient physical cause (*p1*). The mere combination of *Psychophysical Causation* and *Closure* does not entail the identity of *m* and *p1*. According to *Closure*, the whole causal story can be told about any physical effect without leaving the physical domain. However, *Closure* does not rule out the possibility that physical effects have both physical and non-physical causes. To reach the conclusion that *m* is identical with *p1*, a premise is required which excludes *m* from being a cause of *p2* unless it is identical with *p1*. This is the role of *Non-Overdetermination*. To give a standard example of causal overdetermination: Two shots are *independently* fired and 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. *Non-Overdetermination* permits isolated cases of causal overdetermination but rules out events being *systematically* causally overdetermined. Hence, given *Non-Overdetermination*, it cannot be the case that *whenever* *m* causes *p2*, *p1* also causes *p2*, where it is such that if one of the two events *m* and *p1* had not existed, the other would have sufficed, in the circumstances, to cause *p2*. It is precisely this kind of systematic causal overdetermination that the combination of *Psychophysical Causation* and *Closure* seems to give rise to. The problem is removed if *m* is identical with *p1*.

The conclusion that mental causes are identical with physical causes is obviously incompatible with substance dualism.² However, the causal

1 Papineau (1998: 375).

2 The causal closure argument is, of course, unproblematic for epiphenomenalist or non-interactive parallelist dualisms, but neither of these positions has many contemporary advocates. For a modern dualist epiphenomenalism, see Robinson (2004).

closure argument is also problematic for those positions that combine a substance monism with a property dualism. This obviously follows if, for example, events are the causal relata and events are Kimean. According to Kim, a mental event is the instantiation of a mental property by a substance at a time and a physical event is the instantiation of a physical property by a substance at a time. Furthermore, the identity of events requires the identity of the properties that they involve. Hence, for mental causes to be identical with physical causes, mental properties must be identical with physical properties. Even if one's account of the causal relata allows that mental causes can be physical while mental properties are not, the problem facing property dualism does not go away. The question about the causal redundancy of the mental simply arises at the level of properties, as opposed to the level of causes. (For discussion, see Heil and Mele 1993.)

Problematically, most physicalists do not wish to identify mental properties with physical properties. This is largely as a result of Putnam's multiple realizability argument. The popular physicalist stance is instead a non-reductive physicalist one. Mental properties are distinct from, but ontologically dependent on, physical properties – they 'supervene on' or are 'realized by' physical properties. By proposing this dependence relation, non-reductive physicalists aim to provide a property dualism that is compatible with a physicalist stance – all mental properties are 'nothing over and above' physical properties. As non-reductive physicalism is committed to psychophysical causation, and because it is a property dualism, this position is a target of the causal closure argument.

Indeed, Kim – whose attack on non-reductive physicalism extends over two decades (see, for example, Kim 1989, 2005) – argues that non-reductive physicalism not only faces a problem regarding psychophysical causation but also one regarding purely mental causation (Kim 2005). Kim's 'supervenience argument' is as follows: Say that M and M^* are mental properties and M causes M^* . Given non-reductive physicalism, M^* will have a physical realizer, P^* . Is the existence of M^* due to M (its supposed cause) or P^* (its realizer)? This creates a tension. (Note, the tension is not a causal one – P^* is metaphysically sufficient, not causally sufficient for M^* .) Given this tension, the non-reductive physicalist's response must be that M causes M^* by causing P^* . But if M causes P^* , this conflicts with the causal closure argument. Indeed, as M itself will have a physical realizer (P) that it is plausible to identify as a cause of P^* , M and P causally compete as the cause of P^* . Given *Closure*, P wins.

It is often asked whether the supervenience argument generalizes to all macroproperties. This would follow if one understands the world to consist of a hierarchy of objects and properties at distinct levels. Physics's concern is with objects and properties at the basic level. Each of the special sciences deals with a domain of higher-level objects and properties. Each level of being is distinct from, but dependent on, the level below it. If one accepts this

layered view of reality, then the supervenience argument creates a problem not only for mental causation but also for biological causation, chemical causation, etc., for then it seems that ‘causation at *any* level gives way to causation at the next lower level’ (Kim 2005: 52). We are faced with what Block (2003) refers to as the ‘causal drainage problem’. All of the causal powers of entities at higher levels drain away to those at the level of physics. Whether this provides a reason to think that there must be something wrong with the supervenience argument or, instead, throws into question the existence of all higher-level causation is widely disputed. (For recent discussion, see Block 2003; Kim 2005; Walter 2008). For a general discussion of the difficulties facing a layered view of reality, see Heil (2003).

2. Non-reductive physicalists are committed to *Psychophysical Causation* and *Closure*.³ All must therefore respond either by rejecting *Non-Overdetermination* or by demonstrating that the causal closure argument is invalid.

The most popular non-reductive response in the recent debate is the compatibilist one (Pereboom 2002; Bennett 2003, 2008; Shoemaker 2007, 2013; Crane and Árnadóttir 2013). According to it, psychophysical causation always involves two distinct, sufficient causes, but they do not overdetermine the physical effect. The case in which a physical effect has both a sufficient mental and physical cause is importantly different from a genuine case of causal overdetermination – it is not, for example, analogous to the case in which an individual’s death is caused by two bullet strikings, each of which was individually sufficient for the victim’s death. This is because in the latter case the distinct causes are independent of one another, whereas in the former case, given non-reductive physicalism, they are not – mental properties depend on their physical realizers, and it is supposedly their physical realizers that they are in causal competition with. Rather than *Non-Overdetermination*, compatibilists commonly deny a claim left implicit in the causal closure argument presented above, but which Kim makes explicit in his version of the causal closure argument. According to it:

Exclusion: No single event can have more than one sufficient cause occurring at any given time – unless it is a genuine case of causal overdetermination. (Kim 2005: 42).

Kim considers *Exclusion* to be ‘virtually an analytic truth with not much content’ (Kim 2005: 51). Compatibilists disagree, claiming that, at most, this would be true if one’s concern was with *independent* causes.

3 As Crane (2001) comments, to deny *Closure* and, hence, suggest that a purely physical account leaves out some of the real causes of physical things, is not something any physicalist should want to say.

For compatibilism to succeed, one must provide a plausible account of the dependence relation between mental and physical properties and explain exactly why appealing to it allows one to reject *Exclusion*.

It is now standardly recognized that the non-reductive physicalist cannot simply appeal to psychophysical supervenience to capture a property dependence. As Kim (1993) has pointed out, supervenience cannot purport to capture an asymmetrical dependence relationship as it merely reports a pattern of property co-variation. Instead, non-reductive physicalists have, for example, appealed to the determinable–determinate relation (Yablo 1992; Wilson 2009), the mereological relation (Clapp 2001), and metaphysical necessitation (Bennett 2003, 2008). Each is supposed to capture the idea that mental properties are distinct from, but ‘nothing over and above’, their physical realizers.

Because of their intimate relation, the compatibilist’s thought is that far from there being any causal competition between a mental property and its physical realizer, the former actually inherits its causal powers from the latter. The subset account of property realization provides the central example of such an account. (See, for example, Shoemaker 2001, 2007, 2013, and Wilson 1999, 2011). According to it:

Where *X* and *Y* are properties instantiated by the same object, *X* is realized by *Y* just in case the conditional causal powers that *X* bestows are a proper subset of those that *Y* bestows. (Shoemaker 2013)

Given this account, if mental property *M* is realized by physical property *P* on a particular occasion, then every conditional power that *M* bestows will be numerically identical with one that *P* bestows. Hence, rather than the causal powers of *P* excluding those of *M*, they actually include them. Despite this, the instantiation of *M* is distinct from the instantiation of *P*, as *M* and *P* are distinct properties given Shoemaker’s criterion of property identity. According to it, properties are identical if and only if they bestow the same *set* of conditional powers. In other words, *M* and *P* are distinct because *M* only inherits some, not all, of *P*’s conditional powers.⁴

Not all non-reductive physicalists are compatibilists. For example, List and Menzies reject *Exclusion*, but, unlike the compatibilist, they do not maintain that mental causes ‘piggyback’ on physical ones (List and Menzies 2009; Menzies 2013). Their response rests upon a difference-making account of causation. Given this account of causation, *Exclusion* is unsatisfactory because it assumes that one event being causally sufficient for an effect excludes any other event from being causally relevant to that effect. In doing so, *Exclusion* mistakes causal sufficiency for causation. List and Menzies replace *Exclusion* with an exclusion principle formulated in terms of causation

4 For alternative versions of the inheritance solution that do not appeal to the subset account, see Levine (2001) and Pereboom (2002).

understood as difference-making. Crucially, the revised principle allows that there can be not only the familiar upwards exclusion but also downwards exclusion. As List and Menzies observe, cases of downwards exclusion support the causal autonomy of higher-level properties.⁵ This appeal to the notion of causation as dependence rather than production reflects a growing trend amongst non-reductive physicalists. (See Woodward 2008, Raatikainen 2010, and Loewer 2007.)

Finally, not all non-reductive physicalists reject *Exclusion*. Some reject the homogeneity of mental and physical causation – a further premise that is left implicit in the causal closure argument. To deny this premise is to propose that mental and physical causation differ, not merely in what they relate, but also in the kind of causation that they involve (Crane 1995). Jackson and Pettit (1990) arguably adopt this approach, distinguishing between causal efficacy (which is associated with causal production) and causal relevance (which is an explanatory notion). Mental events are causally relevant to behaviour but are not causally efficacious in producing behaviour. It is instead the particular physical realizer of the mental event that is causally efficacious.

One general worry when considering non-reductive physicalist accounts is whether they truly capture the causal efficacy of mental properties or whether they merely capture the causal explanatory relevance of mental predicates. Thus, Heil claims that all non-reductive physicalists are guilty of confusing the idea that mental and physical predicates operate at different levels of description or explanation with the idea that mental and physical properties exist at different ontological levels. The property dependence relations that non-reductive physicalists claim to be capturing are really dependence relations between predicates. It follows that the causal relevance that the non-reductive physicalist appears to secure for the mental is mere causal explanatory relevance, not causal efficacy. According to Heil, claims to the contrary tend to rest on an unsatisfactory account of a property's existence and identity or on a weakened notion of causation. (Heil 2003, 2013. Also see Heil and Robb 2003. For related points, see Gibb 2009 and Kim 2005).

3. Some of those who are dissatisfied with non-reductive physicalism think that if psychophysical causation (and, indeed, mental causation) is to be saved, all hope rests with psychophysical reductionism (Kim 2005: 161). However, in the recent literature the general presumption that an interactive dualist response to the causal closure argument is unworkable has begun to be challenged. (See Koons and Bealer 2010 and Corradini and O'Connor 2010.)

Interactive dualists include both interactive substance dualists and interactive anti-physicalist property dualists. Substance dualism maintains that

5 See Noordhof (2013) whose version of non-reductive physicalism provides an interesting contrast with that of Menzies and List's

mental and physical substances are distinct. But note, to be a substance dualist one need not be committed, as Descartes was, to the existence of a purely non-physical substance (such as an immaterial mind). Thus, Lowe maintains a non-Cartesian substance dualism, according to which persons are simple substances possessing both physical and mental properties: they are things that think and feel but which also have, for example, a spatio-temporal location. He thereby avoids many of the traditional problems facing Cartesian dualism (Lowe 2006, 2008, 2010). However, all substance dualists agree that mental properties are not properties of the body.

Anti-physicalist property dualists, like non-reductive physicalists, maintain that mental properties are non-physical properties of the body. But, unlike non-reductive physicalists, they maintain a ‘strong ontological emergence’. What exactly should be understood by ‘strong ontological emergence’, beyond its negative characteristic of being incompatible with the physicalist claim that mental properties are ‘nothing over and above’ physical properties, is a much disputed issue, which I do not have the space to explore properly here. The standard picture of ontological emergence regards emergence as a supervenience relation that holds in virtue of emergent laws (Broad 1925; McLaughlin 1997). However, one concern is whether this picture can allow that emergent properties have genuinely novel downward causal powers or whether it forces one to accept that they are epiphenomenal. This has led to the proposal of alternative accounts of ontological emergence. Hence, O’Connor (2000) and O’Connor and Wong (2005) understand emergence as a non-supervening, causal relationship, and Humphreys (1997) understands emergence as fusion. Arguably to capture a ‘strong ontological emergence’ worthy of the name, mental properties must exist in as metaphysically robust a sense as the physical properties from which they arise. Moreover, mental properties must have full-blooded, independent causal powers that exist over and above the causal powers of the physical properties from which they arise, and which can, given a commitment to *Psychophysical Causation*, affect lower levels of properties on the macro–micro hierarchy of objects.

Contrary to Bennett (2008), who argues that interactive dualists have to accept *Exclusion*, Kroedel (forthcoming) argues that rejecting *Exclusion* is a viable option for the interactive dualist. However, most interactive dualists respond to the causal closure argument by questioning the causal closure principle.

Lowe observes that there is a general lack of agreement about how this principle should be formulated. He argues that the central problem for proponents of the causal closure argument is that of providing a formulation of the principle that is of the *correct* strength (Lowe 2000, 2008). 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, to provide an argument for it, one must first smuggle in physicalist assumptions, and,

hence, risk making the causal closure argument a circular one. Equally, it must not be so strong that it lacks empirical – or, indeed, metaphysical – support. Lowe contends that the physicalist cannot meet the challenge of providing a formulation of the causal closure principle that is neither too weak nor too strong by these standards.

Lowe rejects *Closure* as too weak. Given that causation is transitive, the requirement that every physical effect has a sufficient physical cause would be met by a chain of causation in which a physical event was caused by a mental event which was itself caused by a physical event. Lowe therefore turns to a much stronger principle:

*Closure**: Every physical event contains only other physical events in its transitive causal closure.

By the ‘transitive causal closure’ of an event (call it E), Lowe means the set of events which includes the immediate causes of E, the immediate causes of those causes...and so on (Lowe 2000: 581). According to Lowe, the combination of *Closure**, *Psychophysical Causation* and *Non-Overdetermination* still does not rule out interactive dualism. This combination, Lowe observes, is perfectly consistent with his dualist model of psychophysical causation which rests upon the idea that the causal role of mental events in the physical domain is to make the fact that a causal tree of neural events converges upon a particular bodily movement non-coincidental. Mental events are able to play this role because of their intentional nature (Lowe 2000, 2008). Of course one might provide an even stronger formulation of the causal closure principle – one which, when inserted into the causal closure argument, would be strong enough to rule this account out. But it is hard to see what empirical evidence one could provide in support of it. One potential problem that stands in the way of an attempt to provide such evidence is that given this account of psychophysical causation there will be no gaps in the chains of neurophysiological events that terminate in bodily movement – mental events do not, according to it, cause physical events by initiating any single physical event or set of physical events in these chains. Hence, the causal role that the mental plays in the physical domain will be completely invisible to the scientist (Lowe 2008: Ch. 3).

There are other dualist models of psychophysical causation that allow one to reach a similar conclusion. Thus Gibb (2013) adopts a powers theory of causation (according to which an effect is the manifestation of a causal power) and advances a model of psychophysical causation according to which the causal role of mental events in the physical domain is to serve as ‘double preventers’ – mental events are, in other words, enabling events which permit physical events to be caused. This model, Gibb argues, is also compatible with the combination of *Closure**, *Psychophysical Causation* and *Non-Overdetermination*. As with Lowe’s account, this causal role will be invisible to science.

Even if *Closure** were strong enough to provide a physicalist conclusion when inserted into the causal closure argument, the question still remains as to whether it is true. Some think that *Closure**, or something like it, is a fact of physics that does not need to be argued for – the burden of evidence is on the interactive dualist to prove it false. Others think it is obviously false. (See, for example, BonJour 2010.) For two recent detailed defences of the causal closure principle, see Papineau (2002) and Melnyk (2003).

A general interactive dualist worry is that when one closely examines the arguments that have been offered by physicalists for *Closure** (or something like it), one often finds that they build in some hidden assumption that interactive dualists would wish to reject, and, hence, risk begging the question. To give a specific example, Gibb (2010) argues that this is the case with Papineau (2002) which provides one of the most thorough defences of the causal closure principle in the recent literature. Papineau's claim is that in the light of evidence from theoretical physics and physiological research, there is probably no non-physical energy. On this basis, he asserts that we can conclude that the causal closure principle has, by any normal inductive standards, been fully established. However, Gibb (2010) argues that Papineau's argument does not establish the probable truth of the causal closure principle, unless certain assumptions are first made about the nature of psychophysical causation. These assumptions are ones that certain interactive dualist models of causation can and do reject. (See the models of psychophysical causation proposed by Broad 1925 and Lowe 2000, 2008.) And new theories of agent causation – which abandon the widely accepted belief in the mental causation debate that all causation is event causation – make the job of defending *Closure** even harder (Lowe 2008, 2013; O'Connor 2000; O'Connor and Wong 2005).

At the same time, philosophers of mind are beginning to challenge the causal closure principle as a result of closer investigation into what physical science itself has to say about the causal structure of the physical domain. It is argued that physical science, far from supporting *Closure**, actually calls it into question. Hence, Hendry (2006) challenges the causal closure principle in the domain of chemistry. The principle has also been challenged in the domain of physics itself. To accommodate the indeterministic aspect of quantum mechanics, probabilistic versions of the causal closure principle have been advanced. (See, for example, Crane 1995: 6 and Papineau 1993: 22.) But this does not engage with the deeper issue – namely, that quantum systems are arguably holistic systems, and that the holistic nature of these quantum systems conflicts with the causal closure principle. For discussions of this particular issue and also other defences of emergence in physics, see Barrett (2006), McGivern and Reuger (2010), Stapp (2005), and Teller (1986).

This is to provide but a snapshot of some, but far from all, of the central stances in the mental causation debate. It should be clear, however, that which position to adopt in this debate is still very much an open question,

and that those in the debate are increasingly turning to contemporary physics and metaphysics for guidance.⁶

University of Durham
Durham DH1 3HN, UK
s.c.gibb@durham.ac.uk

References

- Barrett, J. 2006. A quantum-mechanical argument for mind-body dualism. *Erkenntnis* 65: 97–115.
- Bennett, K. 2003. Why the exclusion problem seems intractable, and how, just maybe, to tract it. *Noûs* 37: 471–97.
- Bennett, K. 2008. Exclusion again. In *Being Reduced: New Essays on Reduction, Explanation, and Causation*, eds. J. Hohwy and J. Kallestrup, 280–306. Oxford: Oxford University Press.
- Block, N. 2003. Do causal powers drain away? *Philosophy and Phenomenological Research* 67: 133–50.
- BonJour, L. 2010. Against materialism. In *The Waning of Materialism*, eds. R. Koons and G. Bealer, 3–24. Oxford: Oxford University Press.
- Broad, C. 1925. *The Mind and Its Place in Nature*. London: Routledge & Kegan Paul.
- Clapp, L. 2001. Disjunctive properties: multiple realizations. *Journal of Philosophy* 98: 111–36.
- Corradini, A. and T. O'Connor, eds. 2010. *Emergence in Science and Philosophy*. London: Routledge.
- Crane, T. 1995. The mental causation debate. *Proceedings of the Aristotelian Society* 69(Suppl.), 211–36.
- Crane, T. 2001. *Elements of Mind: An Introduction to the Philosophy of Mind*. Oxford: Oxford University Press.
- Crane, T. and S. Árnadóttir. 2013. There is no exclusion problem. In *Mental Causation and Ontology*, eds. S. Gibb, E. J. Lowe and R. Ingthorsson, 248–66. Oxford: Oxford University Press.
- Gibb, S. 2009. Review: Shoemaker: Physical Realization. *Mind* 118: 207–11.
- Gibb, S. 2010. Closure principles and the laws of conservation of energy and momentum. *Dialectica* 64: 363–84.
- Gibb, S. 2013. Mental causation and double prevention. In *Mental Causation and Ontology*, eds. S. Gibb, E. J. Lowe and R. Ingthorsson, 193–214. Oxford: Oxford University Press.
- Gibb, S., E.J. Lowe, and R. Ingthorsson, eds. 2013. *Mental Causation and Ontology*. Oxford: Oxford University Press.
- Heil, J. 2003. *From an Ontological Point of View*. Oxford: Oxford University Press.
- Heil, J. 2013. Mental causation. In *Mental Causation and Ontology*, eds. S. Gibb, E. J. Lowe and R. Ingthorsson, 18–34. Oxford: Oxford University Press.
- Heil, J. and A. Mele, eds. 1993. *Mental Causation*. Oxford: Clarendon Press.

6 See further Gibb et al. (2013).

- Heil, J. and D. Robb. 2003. Mental properties. *American Philosophical Quarterly* 40: 175–96.
- Hendry, R. 2006. Is there downward causation in chemistry? In *Philosophy of Chemistry: Synthesis of a New Discipline*, eds. D. Baird, E. Scerri, and L. McIntyre. *Boston Studies in the Philosophy of Science*, 242: 173–89. Dordrecht: Springer.
- Humphreys, P. 1997. Emergence, not supervenience. *Philosophy of Science* 64: 337–45.
- Jackson, F. and P. Pettit. 1990. Program explanation: a general perspective. *Analysis* 50: 107–17.
- Kim, J. 1989. Mechanism, purpose, and explanatory exclusion. *Philosophical Perspectives* 3: 77–108.
- Kim, J. 1993. The nonreductivist's troubles with mental causation. In *Mental Causation*, eds. J. Heil and A. Mele, 189–210. Oxford: Clarendon Press.
- Kim, J. 1999. *Mind in a Physical World*. Cambridge, MA: MIT Press.
- Kim, J. 2005. *Physicalism, or Something Near Enough*. Princeton: Princeton University Press.
- Koons, R. and G. Bealer, eds. 2010. *The Waning of Materialism*. Oxford: Oxford University Press.
- Kroedel, T. forthcoming. Dualist mental causation and the exclusion problem. *Noûs*.
- Levine, J. 2001. *Purple Haze: The Puzzle of Consciousness*. New York: Oxford University Press.
- List, C. and P. Menzies. 2009. Nonreductive physicalism and the limits of the exclusion principle. *Journal of Philosophy* 106: 475–502.
- Loewer, B. 2007. Mental causation, or something near enough. In *Contemporary Debates in Philosophy of Mind*, eds. B. McLaughlin and J. Cohen, 243–65. Oxford: Blackwell Publishing.
- Lowe, E.J. 2000. Causal closure principles and emergentism. *Philosophy* 75: 571–85.
- Lowe, E.J. 2006. Non-Cartesian substance dualism and the problem of mental causation. *Erkenntnis* 65: 5–23.
- Lowe, E.J. 2008. *Personal Agency: The Metaphysics of Mind and Action*. Oxford: Oxford University Press.
- Lowe, E.J. 2010. Substance dualism: a non-Cartesian approach. In *The Waning of Materialism*, eds. R. Koons and G. Bealer, 439–62. Oxford: Oxford University Press.
- Lowe, E.J. 2013. Substance causation, powers, and human agency. In *Mental Causation and Ontology*, eds. S. Gibb, E. J. Lowe and R. Ingthorsson, 153–72. Oxford: Oxford University Press.
- McGivern, P. and A. Reuger. 2010. Emergence in physics. In *Emergence in Science and Philosophy*, eds. A. Corradini and T. O'Connor, 213–32. London: Routledge.
- McLaughlin, B. 1997. Emergence and supervenience. *Intellectica* 2: 25–43.
- Melnyk, A. 2003. *A Physicalist Manifesto: Thoroughly Modern Materialism*. Cambridge: Cambridge University Press.
- Menzies, P. 2013. Mental causation in the physical world. In *Mental Causation and Ontology*, eds. S. Gibb, E. J. Lowe and R. Ingthorsson, 58–87. Oxford: Oxford University Press.
- Noordhof, P. 2013. Mental causation: ontology and patterns of variation. In *Mental Causation and Ontology*, eds. S. Gibb, E. J. Lowe and R. Ingthorsson, 88–125. Oxford: Oxford University Press.
- O'Connor, T. 2000. *Persons and Causes*. Oxford: Oxford University Press.

- O'Connor, T. and H. Wong. 2005. The metaphysics of emergence. *Noûs* 39: 658–78.
- Papineau, D. 1993. *Philosophical Naturalism*. Oxford: Blackwell.
- Papineau, D. 1998. Mind the gap. *Philosophical Perspectives* 12(S12): 373–89.
- Papineau, D. 2002. *Thinking About Consciousness*. Oxford: Oxford University Press.
- Pereboom, D. 2002. Robust nonreductive materialism. *Journal of Philosophy* 99: 499–531.
- Raatikainen, P. 2010. Causation, exclusion, and the special sciences. *Erkenntnis* 73: 349–63.
- Robinson, W. 2004. *Understanding Phenomenal Consciousness*. Cambridge: Cambridge University Press.
- Shoemaker, S. 2001. Realization and mental causation. In *Physicalism and Its Discontents*, eds. C. Gillett and B. Loewer, 74–98. Cambridge: Cambridge University Press.
- Shoemaker, S. 2007. *Physical Realization*. Oxford: Oxford University Press.
- Shoemaker, S. 2013. Physical realization without preemption. In *Mental Causation and Ontology*, eds. S. Gibb, E. J. Lowe and R. Ingthorsson, 35–57. Oxford: Oxford University Press.
- Stapp, H. 2005. Quantum interactive dualism: an alternative to materialism. *Journal of Consciousness Studies* 12: 43–58.
- Teller, P. 1986. Relational holism and quantum mechanics. *British Journal for the Philosophy of Science* 37: 71–81.
- Walter, S. 2008. The supervenience argument, overdetermination, and causal drainage: assessing Kim's master argument. *Philosophical Psychology* 21: 673–96.
- Wilson, J. 1999. How superduper does a physicalist supervenience need to be? *Philosophical Quarterly* 49: 33–52.
- Wilson, J. 2005. Supervenience-based formulations of physicalism. *Noûs* 39: 426–59.
- Wilson, J. 2009. Determination, realization, and mental causation. *Philosophical Studies* 145: 149–69.
- Wilson, J. 2011. Non-reductive realization and the powers-based subset strategy. *The Monist* 94: 121–54.
- Woodward, J. 2008. Mental causation and neural mechanisms. In *Being Reduced: New Essays on Reduction, Explanation, and Causation*, eds. J. Hohwy and J. Kallestrup, 218–62. Oxford: Oxford University Press.
- Yablo, S. 1992. Mental causation. *Philosophical Review* 101: 245–80.