

Could Mental Causation Be Invisible?

David Robb <[darobb@davidson.edu](mailto:darobb@ davidson.edu)>

Draft of May 2012; comments are welcome

I. Emergentism

My topic is an emergentist model of mental causation recently proposed by Jonathan Lowe (2003). There are several forms of emergentism,¹ some making epistemic or explanatory claims, but I'll take emergentism to be a purely metaphysical view consisting of three theses, the first of which is:

(1) Mental properties are not physical.

Lowe's emergentism goes beyond this to event and substance dualism, but in what follows, I'm not going to make much of the distinctions among these various forms of dualism. The official emergentist position will be property dualism, but it will sometimes be easier to speak, as Lowe does, of events or processes rather than properties.

The second thesis is:

(2) Mental properties depend on physical properties.

Emergent properties, while non-physical, cannot "float free" of the physical.

Emergentists differ on the strength of this dependence. Lowe takes it to be *causal*.

Adding more detail to (2), then, an emergentist says that when (and only when) physical systems achieve a certain level of complexity, they cause to be instantiated non-physical properties of those same physical systems.

The final emergentist thesis is:

(3) Mental properties exert a novel (non-redundant), downward causal influence on the physical systems from which they emerge.

¹ Overviews include McLaughlin 1992, Kim 1999, and O'Connor and Wong 2012.

Emergent properties make a difference, a causal difference, to what happens physically, and in particular to behavior. They do causal work that the physical on its own cannot do. This third doctrine distinguishes emergentism from weaker forms of property dualism on which mental properties are realized in, and so inherit causal powers from, their physical base properties. Emergent properties aren't like this. There are new causal powers in the world—extra causal “oomph”—when an emergent comes on the scene.

Whatever else you might think about emergentism, it does seem to have the scientific virtue of making a bold empirical claim about the long-term results of human physiology. As we study the structure and dynamics of the human body, emergentism predicts we will eventually find that physical properties do not have the causal resources to produce all of what we observe.

But with emergentism's prediction comes risk, for it might not turn out this way. Suppose we discover that everything physical—and in particular, whatever happens in and to the human body—can be fully accounted for by physical causes. In that case, we would have discovered that the physical world (at least our part of it) is “causally closed” or “causally complete”. Following Lowe and others, we can formulate the resulting Closure principle like this:

Closure: Everything physical that has a cause at time t has a sufficient physical cause at t .²

If Closure were empirically confirmed, this would mean trouble for emergentism. Indeed, we'd have the state of affairs envisioned by those who have advanced a “causal argument” for physicalism, an argument using Closure as a premise.³ Closure ensures that any mental cause of a physical effect will be accompanied by a sufficient physical cause of that effect. If the mental cause is included in (reducible to) that physical cause, then emergentism is false, as thesis (1) says mental properties aren't physical. On the

² The principle is stronger than is needed in this context, as a local version, restricted to the human body, would do just as well. But I'll follow Lowe and others in using the universal version (and in using the name “Closure”, though “Completeness” is another common, and slightly more accurate, label).

³ Papineau 2001 is a recent example. The standard causal argument requires a “no overdetermination” or “non-redundant” premise, but the version in the text can dispense with this, since emergent properties (as defined here) exert non-redundant efficacy.

other hand, if the mental cause is *not* included in the physical cause, then any efficacy of the mental cause will be redundant, as the physical cause is (by hypothesis) sufficient for the effect. But this too is incompatible with emergentism, as thesis (3) says the efficacy of mental properties is non-redundant. And to abandon mental causes altogether—the epiphenomenalist option—is also incompatible with thesis (3). Closure seems to leave the emergentist with no options, leaving some form of physicalism as the best alternative.

That said, it's not my aim here to develop or defend this argument for physicalism, nor will I weigh in on whether science has vindicated (or will vindicate) Closure. Instead, my more modest aim is to argue that—contrary to some recent claims by Lowe—the case for or against emergentism is what it appears to be, namely *empirical*. Along the way, I also want to use the issue to discuss some broader topics in the metaphysics of mental causation, including the proper formulation of Closure and the nature of causal powers.

II. Lowe's model of mental causation

Turn then to Lowe's model of mental causation.⁴ Lowe is an emergentist, but he has recently argued that emergent mental causation could be empirically undetectable, and in this sense "invisible". This seems to me unwelcome news, for it means that what we thought was a step forward on the mind-body problem—namely, that a significant portion of it can be handed off to the scientists—may not be progress after all. Indeed, Lowe suggests that if his "invisibility" model is correct, the dispute between physicalists and emergentists may have to remain largely within philosophy (where Lowe thinks emergentists have the stronger case).


First the model itself. I'm going to compress Lowe's presentation in various ways, but I hope I'm not oversimplifying it. Start with the physical effect, *B*. You can think of *B* as an instance of *behavior*, say my raising my arm, but given other aspects of the model, it's probably better to think of *B* as the *beginning* of a physical process that leads to (or

⁴ In a number of publications, Lowe has explored several models of dualist interactionism (in, e.g., Lowe 1992; 1996, ch. 3). My remarks here are limited to the one defended in Lowe 2003 (and 2008, ch. 3).

partly constitutes) that behavior. Ignoring such niceties, I'll refer to *B* as the *behavioral effect*. *B* is physical, and occurs at time t_2 .

Now suppose *B* is initiated by a mental cause, *M*, at t_1 , the time immediately preceding t_2 .⁵ Given Closure, *B* also has at t_1 a sufficient physical cause, *P*. Lowe has doubts about the proper formulation, not to mention the truth, of Closure (2000; 2008, ch. 2). But here he grants the principle as formulated; indeed, this is central to his case for invisibility. So now the materials for the causal argument are in place: *M* must either be included within *P* or causally redundant, and neither option is compatible with emergentism.

But in an ingenious block of this argument, Lowe says the mental cause could be emergent, but in a way that respects Closure. Suppose that at t_1 , *P* causes emergent property *M* to be instantiated at that very time. Moreover, *M* then aids non-redundantly in the production of *B*. That is, if (counter-nomically) *P* were present without *M*, *P* would not be enough to cause *B*. *M*'s added oomph is needed. But this does not impugn *P*'s sufficiency for *B*, since *P* itself is sufficient for *M*, its needed helper. The model can be represented like this, where each arrow stands for causal sufficiency:



On this model, is *M*'s efficacy empirically detectable? Except for a hedge at one place,⁶ Lowe says No: *M*'s efficacy is invisible.⁷ Crucial is that Closure is not violated, so that anyone empirically tracking the causal history of *B* will find a sufficient physical cause at

⁵ Here time is discrete, a simplifying (but not essential) aspect of the model.

⁶ "While this is not to say that there is no kind of empirical evidence that could have a bearing on this debate, the implication would seem to be that purely metaphysical arguments must have a larger role to play than is commonly supposed in determining which of [interactive dualism or one of its physicalist rivals] is superior" (Lowe 2008, 58).

⁷ Note that the question here and below is not whether *M* itself is visible (Lowe 2003, 151). Presumably there are any number of methods, both first-person and third-person, to detect the presence of a mental property. The question is whether *M*'s efficacy—and in particular, its efficacy with respect to *B*—is visible to empirical investigators.

t_1 , and so will see no reason to postulate a non-physical cause—this in spite of the fact that a non-physical cause is required for B .

Note that there's not going to be a way to test for M 's non-redundant efficacy by removing the mental state (by, say, distracting the subject or rendering him unconscious) while holding P fixed. From our God's-eye view, we see that the counter-nomic conditional mentioned earlier is true, but our hypothetical scientists can't see this, at least not in the way envisaged. Any removal of M would also involve tampering with P , thus altering B 's physical causal history. It will continue to appear as if P is doing all of the work.

By the way, it's important for Lowe's purposes that P 's causation of M be simultaneous. To see this, suppose for a moment that t_1 and t_2 are not adjacent times, but that there is some time between them, and that P causes M to be instantiated, not at t_1 , but at this intermediate time.⁸ In that case, we'd have a detectable violation of Closure, for at the intermediate time, there would be a cause of B (namely, M) but no sufficient physical cause of B . Might we salvage Closure by postulating a sufficient physical cause of B , call it P^* , at this intermediate time? But such a P^* would be either sufficient on its own for B , or sufficient for B by producing the needed helper M . In the former case, M 's efficacy becomes redundant, contrary to emergentism. In the latter case, we are back to Lowe's model with just some re-labeling. The upshot is that an emergentist trying to respect Closure should require M to be simultaneously caused by its physical base P .

III. Powers

Having sketched Lowe's model, I now want to start to look more critically at his invisibility claim.

First, it's because Closure is allegedly respected that M is invisible. But doesn't there seem to be a clear sense in which the model *violates* Closure? After all, something non-physical, namely M , is required for the production of a physical effect, B . The point of Closure is that the physical world contains the full resources for the production of any

⁸ I take it such is the picture endorsed by, for example, O'Connor and Wong 2005, who thereby reject Closure.

physical effect, yet on this model, the physical world *doesn't* contain the full resources for the production of *B*, since the non-physical *M* is required. So something has gone wrong somewhere.

I suspect the problem is that Lowe is reading Closure's "sufficient physical cause" as *nomio* sufficiency: *P* and the relevant laws of nature ensure that *B* occurs. In this sense, *P* is indeed sufficient for *B*. But since Closure is violated on the model, it looks as if *nomio* sufficiency isn't the appropriate way to understand causal sufficiency, at least not as it appears in Closure.

How else could causal sufficiency be understood? Suppose you think of *M*, *P*, and *B*, not just as properties, but more specifically as *powers* and causation as the manifestation of powers (see, e.g., Molnar 2003, Anjum and Mumford 2010, Heil ms). On Lowe's model, *M* is the manifestation of *P*, and *B* is the manifestation of *P* "partnered" with *M*. Switching to powers-talk might be optional in this context, but it provides a useful way to discuss Lowe's model and, in particular, why Closure is violated.

On the powers account of causation, power *y*'s being a *manifestation of powers* $x_1...x_n$ (equivalently, $x_1...x_n$'s *producing* *y*) is taken as an undefined primitive,⁹ but one of its important features can be mentioned here: If *y* is a manifestation of $x_1...x_n$, then $x_1...x_n$ are *complete* in the sense that no other power needs to be partnered with them to produce *y*. It is in this sense that $x_1...x_n$ are causally sufficient for *y*, and it's this productive sense of causal sufficiency, not the *nomio* sense, that, it seems to me, should appear in Closure. (Another reason causal sufficiency here is not *nomio* necessitation is that it's usually possible that the complete set be present, but that some other power be on the scene to prevent the manifestation (see Anjum and Mumford 2010). One may be tempted to say in that case that the original set isn't complete after all: the *absence* of the preventing power must be present as well. But the absence of a power isn't a power, any more than the absence of a chair is a chair.)

⁹ This conceptual point should not be taken to rule out an informative metaphysical (or scientific) theory of the nature of manifestation. For example, it could turn out that manifestation essentially involves the transference of some quantity, such as energy. I'm neutral on such matters here.

Return then to Closure, which currently reads:

Closure: Everything physical that has a cause at time t has a sufficient physical cause at t .

We could stay with this, so long as “sufficient physical cause” is understood in terms of causal production (manifestation), as above. But just to make this feature explicit, I propose reformulating Closure as:

Closure (reformulated): Every physical manifestation is produced by physical powers.

A few notes on this new version:

First, it allows the principle to be a bit more compact. Since manifestations are by definition caused, “every physical manifestation that has a cause” would be redundant. And we no longer need to refer to times, since producing powers always immediately precede their manifestations. Any would-be production that was temporally mediated would not be complete, as the intermediate power(s) would have to be added to produce the manifestation.¹⁰

Second, this modified principle, like its predecessor, is not so strong as to preclude a role for non-physical powers in producing physical manifestations. For all Closure tells us, a given physical manifestation might be produced by physical powers *and* by non-physical powers (or by some combination of the two). Closure requires merely that every physical manifestation is produced at least by physical powers, whether or not it happens to be produced (that is, over-produced) by other sorts of powers as well.

Third, this rephrasing has the virtue of explaining why Lowe’s model is, as it should be, a violation of Closure, for B on Lowe’s model is *not* produced by P , since P , although complete for M , is not complete for B . Rather, B is produced by the partnering

¹⁰ This means that if time is continuous, manifestations are always simultaneous with their producing powers (cf. Huemer and Kovitz 2003, Heil ms.). This makes it especially apt to see the manifestation B , not as behavior, but as the beginning of a process that leads to (or partly constitutes) behavior.

of P and M . Furthermore, the invisibility claim starts to look suspect in a way that it might not have on the nomic understanding of causal sufficiency. After all, if P doesn't produce B , shouldn't this be detectable? With respect to producing B , P should look as odd as a gear with some missing teeth that nevertheless smoothly turns another gear. This is in fact what I will propose, though getting there will require a closer look at the way powers combine to produce their manifestations.

IV. Principles of combination

Thesis (2) of emergentism, when elaborated, says that mental properties emerge only in physical systems of a certain degree of complexity. This means M emerges from a complex combination of physical powers, so that P is really a structural property composed of properties of the system's physical parts (and, perhaps, properties of the immediate environment). Call these the "micro-powers" composing P .

Now I assume that these micro-powers can be studied individually in non-mental contexts—that is, when they do not compose P or any other psychological base, so that the meddling of non-physical powers is not in question. Such a study will let us construct an increasingly detailed *causal profile* for each micro-power, one that specifies the contribution the power makes when a partner in purely physical production.¹¹ The general strategy is then unsurprising: predict, based on these empirically discovered causal profiles, how the micro-powers will partner in the mental context. If B is what's predicted by their profiles, we have evidence that they are the powers producing this behavior. If not—if B is anomalous with respect to these profiles—we have evidence for the contribution of a non-physical power, and M 's efficacy is no longer hidden.

But exposing M 's role cannot be this simple. The worry here is not inductive skepticism about causal profiles. (If skepticism were in play, Lowe's invisibility claim would be more secure, but less interesting, perhaps trivial.) The worry, rather, is that the causal profiles gained from non-mental contexts aren't useful for this project without some *principle of combination* saying how micro-powers with these profiles combine

¹¹ Anjum and Mumford (2010) represent the causal profile of a property with a vector. This is useful for visualizing a power's contribution to a manifestation, but the most accurate representation of a causal profile may need to be considerably more complex, perhaps what Hawthorne (2001, 376) calls a "Shoemaker sentence".

when partnered in a new (in this case, mental) context. I doubt such a principle can be known *a priori*.

But can't it be known empirically? Take a simple case for illustration. Suppose we determine that a solid object shaped like "I" produces an "I" shape when pressed into soft clay, while one shaped like "<" produces a "<" shape. These are (partial) causal profiles for each of these shapes, themselves powers. A principle of combination could tell us that when these two shapes are arranged in the form of a "K", they will produce a "K" shape in soft clay. I will not attempt to precisely formulate the relevant principle here, but it will roughly be an "additive" principle:

Additivity: Powers in combination produce the sum of the manifestations they produce independently.¹²

Thus, if Additivity is true, two bowling balls put on a scale will produce the sum of their independent readings. Two powers expressed as vectors will act as a power expressed as the sum of those two vectors. And so on. Additivity sometimes appears in the literature as the principle that emergence violates. For example, what Mill calls the "principle of the Composition of Causes" (1850, III.vi.1) is additive. Broad's mechanist appeals to a "single and simple principle of composition", which, at least in its paradigmatic form, is additive (1925, 45). And more recently, O'Connor and Wong say an emergent property "will confer causal capacities on the object that go beyond the summation of capacities directly conferred by the object's microstructure" (2005, 665).

In any case, suppose Additivity is confirmed in non-mental contexts. Armed with the causal profiles of our micro-powers, can't we then predict, when they compose *P*, what their additive result should be? If something besides *B* is predicted, then when *B* *does* occur, we have evidence of a non-physical power at work.¹³ Note that the

¹² This principle (and those that follow) could be explicitly restricted to physical powers. Perhaps this would be more natural, given that our evidence is entirely physical. In any case, I don't think such a restriction would affect the main points I want to make, though it would require that they be rephrased.

¹³ Why couldn't this simply be evidence for a previously undetected *physical* power? Because, as noted earlier, skeptical worries about physical properties are not in play here. Lowe's invisibility thesis is interesting because it's claimed to hold when empirical information about all of the relevant physical properties (and their causal profiles) is already in.

epistemic bar here is low. Empirical evidence for some conclusion needn't be infallible, and inductive skepticism is set to one side. We just need good empirical reasons for thinking that a non-physical property is efficacious, at it looks as if this fits the bill.

Additivity is not the only candidate principle of combination. Another is:

Continuity: Similar powers in combination with similar partners produce proportionally similar manifestations.

Suppose that by studying physical powers in non-mental contexts, we can place them along one or more dimensions of similarity, so that the closer two powers are along a dimension, the more similar they are causally. From these similarity orderings, use Continuity to predict how our micro-powers will combine when they compose P , expecting that their joint manifestation will be proportionally similar to the manifestations that their close relatives produce in a similar (but non-mental) combination. Again, if the manifestation B does not meet this expectation, this is evidence for the presence of another power on the scene, which would need to be non-physical.

I take it that Continuity is at work for some of those emergentists who, unlike Lowe, think emergent causation must be detectable. Mental properties emerge only when physical systems have reached a certain threshold of complexity. Once this threshold is reached, radically *discontinuous* effects are observed, effects that cannot be accounted for by the continuous operation of physical causal powers. Such a violation of Continuity, that is, would be evidence for an emergent.

V. The micro-latency hypothesis and intelligibility

The upshot so far is that the efficacy of M in Lowe's model will not be invisible so long as the manifestation B is not what we would predict from the micro-powers composing P and the governing principle of combination, all of which can be discovered empirically. (It's merely in-principle detectability I'm concerned with. I am not here defending either of the principles in the previous section.) There is, however, more one might say here in defense of invisibility.

Even when some confirmed principle of combination is apparently violated in mental contexts, emergence is not the only viable hypothesis. Physicalists and emergentists alike agree that mentality is special, found only in a small part of the universe (perhaps only this planet), and only in a small sliver of time. Moreover, there are well-known idiosyncrasies of mental contexts: only in them do we find phenomenal consciousness, self-awareness, and other unique capacities. It would not be too surprising, then, if a principle of combination confirmed in non-mental contexts were violated in mental contexts, not because of non-physical powers on the scene, but because mental contexts—unique on everyone’s lights—bring out novel manifestations of physical causal powers. One does not need to be an inductive skeptic to consider this a live possibility.

Call this the micro-latency hypothesis (cf. Shoemaker 2002). This hypothesis is not emergentist in our sense, for there’s nothing in physicalism prohibiting physical powers from having novel manifestations in unique contexts. Imagine a chemical that interacts benignly with all known natural substances, but explosively with an artificial substance, one that can be produced only under unusual circumstances, say in a laboratory. Such a reaction may be surprising, but it would be odd to invoke non-physical powers here. The powers at work are all physical: they just partner in surprising ways. The micro-latency hypothesis proposes that this is what’s going on when *P* apparently produces the surprising result *B*. Even though *M* is at work, empirical investigators will have no reason to postulate *M* rather than what has previously been latent in the ordinary micro-powers composing *P*. *M*’s efficacy will continue to be invisible.

(In defense of this hypothesis, I should also note that it needn’t say the micro-powers composing *P* are transformed into different powers—*pace* O’Connor (2000, 113-4) who, unlike Lowe, is a “visible” emergentist. As O’Connor reads the micro-latency hypothesis, it requires that the micro-powers composing *P* are locally responsive to their macro-level circumstances. He then objects that it’s mysterious how a power’s nature could be affected non-locally, according to the broader structure of which it is a part. How does a micro-power “know” that it’s now included in such a larger, mental structure? But it seems to me that the hypothesis doesn’t have to be understood this way. The micro-powers composing *P* don’t know that they are in a mental context. They keep the natures they had outside this context: they continue to “do what they’re

doing”. It’s just that they combine in novel ways in these new circumstances. Imagine two bowling balls, each giving a reading of 15 on a scale. With both on the scale, we get a reading, not of 30, but 40. The latency hypothesis in this case needn’t say that the bowling balls know they are in a “paired” structure, and, thus informed, acquire new weights. Rather, they continue to weigh 15 pounds each—it’s just that their weights do not combine additively. The micro-latency hypothesis the same sort of thing is going on when P produces the anomalous result B . The micro-powers composing P don’t *change* because they are in a mental context. Rather, in this context, they compose in novel ways.)

If the micro-latency hypothesis is always available as a (non-skeptical) alternative to emergentism, Lowe’s invisibility claim would seem to be vindicated. But continuing to press, could it turn out that there is a principle of combination whose violation could *not* be explained by this hypothesis? Suppose we were to empirically confirm a *meta-principle* of combination, one that governs the way powers combine no matter what first-order principle (Additivity, Continuity, or whatnot) they obey in any particular situation. If powers are observed to combine in various first-order ways in non-mental contexts, we might be able to discern such a meta-principle.

What could a meta-principle look like? Here’s one idea:

Intelligibility: Powers are manifested intelligibly.

As a first pass, think of intelligibility like this: when powers are manifested, we can see that the powers and manifestation fit together, much like pieces of a jigsaw puzzle. Knowing (by prior inspection) how the pieces are shaped, we can see that they should interlock. Similarly, knowing (empirically) the causal profiles of various powers, we can see how they and their manifestations fit together.

Suppose, then, that we have evidence from non-mental contexts for Intelligibility across a variety of first-order principles of combination, and suppose that the micro-powers composing P are *not* manifested intelligibly by B in the mental context. That is, empirical investigators can’t see how P and B fit together. This looks like good evidence for a non-physical power partnering with P to produce B . Now suppose the defender of invisibility tries make the latency move: for all our scientists know, the anomalous

manifestation *B* is a result, not of a non-physical power partnering with *P*, but of *P*'s ordinary micro-powers, which simply combine in new, unexpected ways in mental contexts. Intelligibility is supposed to block this move. Novel first-order principles of combination are fine, but the meta-principle requires that we see *how* the novel manifestation arises. If we can't see this, we again have evidence for an emergent power on the scene.

The Intelligibility principle is very schematic, and will no doubt raise eyebrows. For starters, what exactly does it mean to be intelligible? The example of the jigsaw puzzle is merely an analogy, not a definition. And doesn't Intelligibility add an epistemic (and possibly subjective) element to what should be a metaphysical issue? And anyway, didn't Hume show that casual powers are *never* manifested intelligibly? I can't answer all these questions here, or maybe anywhere. But let me make a few remarks to conclude the paper.

First, to lessen, if only slightly, the worries of Humeans: the principle doesn't require that we can know the causal profile of a property *a priori*. We can agree with Hume that Adam, newly created, couldn't just intuit that fire would burn him, or that bread would nourish him. What it does require is that once we empirically discover the causal profiles of micro-properties, their subsequent manifestations will be intelligible to us. This may involve *a priori* inference, but as Sanford (1994) has argued, it's one that comes late in the game. Sanford's example is useful here. Consider two gears that are meshed. Only empirically do we know the natures of the gears' properties—shape, rigidity, and so on. But given this knowledge, we can see—and if you'd like, *a priori* infer—that one gear's turning clockwise will cause the other to turn counterclockwise. It's this last bit of insight that Sanford calls "intelligibility".

Second, Intelligibility as formulated is actually stronger than we need to answer Lowe. I've been speaking of "micro-powers", but a structural property such as *P* can be broken down into properties at various mereological levels. For simplicity, suppose there are two such levels: the level of basic physics (the micro-level), and above it the level of neuroscience. I take that each level has its characteristic properties (powers). Now Intelligibility as formulated requires that powers manifest intelligibly at each level, but this is stronger than we need. Suppose it turns out that powers *don't* manifest intelligibly at the level of basic physics. Perhaps it just looks brute that, for example, like

charges repel. It's compatible with this that neuroscientific powers *do* manifest intelligibly.¹⁴ (Returning to the puzzle analogy, it would be a bit like discovering that a puzzle is ultimately composed of tiny, colored pixels. Why one pixel should be next to another may look brute. But *given* that pixels are arranged to form certain shapes at higher mereological levels, their subsequent ability to interlock is intelligible.) So if we were to discover that this more limited principle of "Neuro-Intelligibility" holds, we could use it to look at what *P* produces in the mental context. If *B* is an unintelligible result at the neuroscientific level, this again is evidence for the efficacy of an emergent, evidence that cannot be explained away with the micro-latency hypothesis.

The principle of Intelligibility is very ambitious and so may obscure my aim here, which is modest: find some way, at least in principle, to force *M*'s efficacy to reveal itself on Lowe's model of mental causation. Empirical confirmation of Intelligibility (or its weaker cousin) looks like it would be enough. If this is right, then it has a lesson for those bold emergentists who take their view to make empirical predictions: they should hope for confirmation of Intelligibility or some similar meta-principle of combination. Otherwise, the debate between physicalists and emergentists may remain with philosophy, and there the evidence is that the debate is likely to remain unresolved.¹⁵

¹⁴ Mill (1850, III.vi.2) makes what seems to be a similar point, though it's put in terms of laws.

¹⁵ Thanks to an audience at Washington University in St. Louis for helpful comments on an earlier draft of this paper, and thanks to James Van Cleve for his questions and suggestions on a number of issues, including the formulation of Closure and the nature of a causal profile.

References

- Anjum, R. and S. Mumford 2010. A Powerful Theory of Causation. In A. Marmodoro, ed., *The Metaphysics of Powers*. Routledge.
- Broad, C. D. 1925. *The Mind and its Place in Nature*. Routledge & Kegan Paul.
- Hawthorne, J. 2001. Causal Structuralism. *Noûs* 35: 361-78.
- Heil, J. ms. Causing. In J. Jacobs, ed., *Putting Powers to Work*.
- Huemer, M. and B. Kovitz 2003. Causation as Simultaneous and Continuous. *Philosophical Quarterly* 53: 556-65.
- Kim, J. 1999. Making Sense of Emergence. *Philosophical Studies* 95: 3-36.
- Lowe, E. J. 1992. The Problem of Psychophysical Causation. *Australasian Journal of Philosophy* 70: 263-76.
- 1996. *Subjects of Experience*. Cambridge University Press.
- 2000. Causal Closure Principles and Emergentism. *Philosophy* 75: 571-85.
- 2003. Physical Causal Closure and the Invisibility of Mental Causation. In S. Walter and H. Heckmann, eds., *Physicalism and Mental Causation*. Imprint Academic.
- 2008. *Personal Agency*. Oxford University Press.
- McLaughlin, B. 1992. The Rise and Fall of British Emergentism. In A. Beckermann, H. Flohr, and J. Kim, eds., *Emergence or Reduction?* Walter de Gruyter.
- Mill, J. S. 1850. *A System of Logic*. Harper and Brothers.

Molnar, G. 2003. *Powers*. Oxford University Press.

O'Connor, T. 2000. *Persons and Causes*. Oxford University Press.

O'Connor, T. and Wong, H. Y. 2005. The Metaphysics of Emergence. *Noûs* 39: 658-78.

----- . 2012. Emergent Properties. In E. Zalta, ed., *The Stanford Encyclopedia of Philosophy*.

Papineau, D. 2001. The Rise of Physicalism. In C. Gillett and B. Loewer, eds., *Physicalism and Its Discontents*. Cambridge University Press.

Sanford, D. H. 1994. Causation and Intelligibility. *Philosophy* 69: 55-67.

Shoemaker, S. 2002. Kim on Emergence. *Philosophical Studies* 108: 53-63.