Emergence, Exclusion, and the Proper Subset of Powers Strategy

Karen Bennett

Rutgers University

Abstract

Wilson characterizes weak and strong emergence partly based on their differing solutions to the exclusion problem. The weak emergentist should claim that emergent phenomena and their bases can both cause the same effect without overdetermining it, because they literally share causal powers. I compare this strategy with a different but related strategy also available to the weak emergentist, and argue that the virtues of the former cost more than it appears.

Keywords: Causal powers, Dependence, Emergence, Exclusion problem, Mental causation, Nonreductive physicalism, Overdetermination.

1. Introduction

Jessica Wilson's *Metaphysical Emergence* (2021) is an excellent and important book that brings together roughly twenty years of work on the ways in which one set of phenomena could be dependent on, and yet to some degree autonomous from, another set of phenomena. Wilson identifies the core shared ideas in the sea of mushy and contradictory usages of the term 'emergence', and articulates notions of 'weak' and 'strong' emergence that (in the philosophy of mind case) correspond to nonreductive physicalism and dualism respectively. She distinguishes these positions, in part, by how they approach the well-known exclusion problem for mental causation. Wilson's discussion of emergence and exclusion will be my focus in this commentary. What exactly does solving the exclusion problem require, and how exactly does her version of weak emergentism pull it off?

Before getting started in earnest, however, I would like to briefly call attention to a particular virtue of Wilson's book: its engagement with, and reliance upon, classic older work in the metaphysics of mind. She engages with a lot of material by people like Terence Horgan, Jaegwon Kim, Andrew Melnyk, Sydney Shoemaker, and Stephen Yablo. This is both appropriate and important, because a lot of excellent work in this area has been somewhat neglected of late. Both Wilson and I began our careers thinking about the mind-body problem, and are therefore well aware that the question of how some things give rise to other things

Argumenta 10, 1 (2024): 237—246
ISSN 2465-2334

DOI 10.14275/2465-2334/202418.ben
First published: 30 November 2024

is not exactly a new topic in metaphysics, as those in the contemporary grounding literature sometimes seem to suggest.

2. Weak and Strong Emergentism, Characterized by How They Handle the Exclusion Problem

Although terms like 'emergence' and 'emergentism' are used in many slightly different ways, Wilson argues that the most basic commitment of philosophical positions worthy of these labels is that emergent properties and states of affairs involve 'autonomy with dependence'. They are synchronically and non-causally dependent on their base, and yet *somehow or other* are autonomous from it: they have different causal powers, figure in different laws, or something along those lines.

That 'somehow or other' is, of course, crucial. Wilson distinguishes two primary forms of emergentism as meaning quite different things by the claim that emergent phenomena have 'different causal powers'. *Weakly* emergent features—if there are any—have fewer causal powers than the bases from which they arise, and *strongly* emergent features—ditto—have more causal powers then their bases. Wilson draws this distinction in the course of exploring available emergentist answers to the exclusion problem. It's a rather neat methodological trick: she simultaneously explains how these two kinds of emergence have different available responses to the exclusion problem, and uses their responses to the exclusion problem to shed light on the difference between them (Chapter 2).

Here's a simple version[1] of the exclusion problem, formulated as a set of five inconsistent claims:

Distinctness: Mental properties (and perhaps events) are distinct from physical properties (events).

Efficacy: mental events cause things, including physical things, and at least sometimes do so in virtue of their mental properties.

Completeness:[2] every physical effect has a sufficient physical cause.

Exclusion: all events that have multiple sufficient causes (that are not themselves causally related)[3] are overdetermined.

Nonoverdetermination: the effects of mental causes are not routinely and systematically overdetermined.

So, the physical effects of mental causes both are and are not systematically overdetermined. No bueno.

[1] The main way in which this version is simplified is that I merely gesture at how it can be run in either or both a property (type) or event version (token). Further, this is not how Wilson presents it. While the differences do not matter to anything of substance, footnotes 4 and 6 are worth reading.

[2] Most people, including Wilson, call this 'closure'. I prefer the label 'completeness', because the term 'closure' suggests that physical effects have *only* physical causes. That is an excessively strong premise that blocks the weak emergentist solution from the start.

[3] The parenthetical clause is there because the proper formulation of Exclusion ought not say that the outcome of a single, non-branching causal chain is overdetermined. If $c_1 \rightarrow c_2 \rightarrow c_3 \rightarrow e$, then e has multiple distinct sufficient causes but is not overdetermined by anyone's lights. An alternate way to circumvent this issue is to instead stipulate that the multiple sufficient causes be direct/unmediated.

One can of course dissolve the exclusion problem by denying that there are any mental phenomena, or claiming that they are epiphenomenal, or insisting that they are to be identified with the physical after all. But, as Wilson points out, these are not *emergentist* responses. They do not respect the core commitments that a) the mental is in some sense emergent (and thus exists), and b) emergent phenomena are in some sense causally autonomous (so mental events/properties are neither epiphenomenal nor identical to their physical bases).

So how should emergentists respond to the exclusion problem? Wilson claims that there are two and only two properly emergentist moves that can be made. The first is to deny Completeness, and claim that mental phenomena have genuinely novel causal powers that are neither determined by nor dependent on their physical bases. This strategy is non-physicalist, and is the distinctively strong emergentist position. The second solution is to deny Exclusion, and say that mental phenomena are causally efficacious and yet their effects are not overdetermined, or at least not overdetermined in the two-kids-simultaneously-throwing-two-rocks-at-a-window variety.[4] This is the weak emergentist or nonreductive physicalist (henceforth 'WE/NP') strategy.

The key WE/NP move is to appeal to an intimate relation short of identity, such as—to borrow Wilson's list (55-57)—functional realization, constitutive mechanism, mereological realization, the determinate-determinable relation, or 'superdupervenience'. (Though Wilson herself would wince (2014, 2018), we might replace some or all of those relations with grounding.)

I have long been fond of the WE/NP response to the exclusion problem, which I once called 'compatibilism'[5] (Bennett 2003). It will be the focus of the rest of the paper.

3. A "Deeper Unity of Strategy"? The Proper Subset Condition and the Counterfactual Condition

Wilson suggests that the fact that different WE/NPs appeal to different intimate non-identity relations is relatively unimportant as far as the exclusion problem is concerned, because

> underlying the seeming diversity in these and many other accounts of nonreductive physicalism hides a deeper unity of strategy (57).

[4] It is just a terminological matter whether we describe this move as saying that the effects of mental causes are not overdetermined at all, or as saying that they are not overdetermined in the bad 'double-rock' way. Discussions and defenses of the strategy take both forms in the literature. (See, e.g., Bennett 2003 and Sider 2003.) Wilson herself frames the strategy in the latter way, as "allowing that [the effects of mental causes] are overdetermined […] but maintain[ing] that the overdetermination here is of an unproblematic *non*-double-rock-throw variety" (44). Characterized like that, the move denies Nonoverdetermination rather than Exclusion: the effects of distinct causes are always overdetermined, but it turns out that overdetermination is more widespread and less troublesome than usually thought.

I prefer the characterization in the main text, which reserves the word 'overdetermination' for the double-rock-style cases. It is also the better characterization for Wilson herself. See note 6.

[5] I called it that because it says that the non-overdeterministic causal efficacy of the mental is compatible with the conjunction of Completeness and Distinctness.

I agree that there is a deeper unity of strategy here. Indeed, I have argued that there *must* be a deeper story, in the sense that the WE/NP ought not simply name an intimate non-identity relation, and announce that events related in that way do not overdetermine their effects. That is not good enough. What is required is a story about how and why that relation has that kind of impact:

> the burden is on the compatibilist here. She needs to be able to *argue* that the effects of mental causes are not overdetermined, and to explain *why* they are not (2003: 474).

That is, in essence, what Wilson is after when she claims a "deeper unity of strategy". She is saying that all of the tight relations postulated by the WE/NP lend themselves to a particular sort of explanation: what I hereby dub the "Proper Subset Strategy".

While I clearly agree about the need for *some* kind of deeper explanation, I am not convinced that the Proper Subset Strategy is the right one. An alternative is available whose relative merits must also be investigated. After sketching both Wilson's story and this alternative, I will explore the relation between them, and argue that the apparent virtues of the Proper Subset Strategy cost more than it seems.

4. Wilson's Proposed Underlying Idea: The Proper Subset Strategy

Wilson claims that whenever one phenomenon *E* is weakly emergent from a base phenomenon *B*, *E*'s causal powers will be a non-empty proper subset of *B*'s. In particular, when mental and physical phenomena stand in any of the close relations posited by the WE/NP, it will be the case that mental phenomena have a non-empty proper subset of the causal powers of the physical phenomena from which they weakly emerge (58-66). Thus the various particular mechanisms for securing weak emergence "are unified in each [endorsing the Proper Subset Strategy] as a means of avoiding problematic overdetermination" (66).

The Proper Subset Strategy certainly sounds good. Indeed, it sounds like it decisively solves the exclusion problem. The picture is that mental and physical causes do not overdetermine their effects because there is a *literal shared core* of causal juice: to say that mental phenomenon *M* and its physical base *P* overdetermine their effects would be wrong in the same way that it would be wrong to say that our two favorite hooligans, Billy and Suzy, overdetermine the breaking of the window by holding hands and jointly throwing one single mutually-owned rock. It is wrong in the same way that it would be to say that you and I double-pay the bridge toll by together tossing in one $5 bill from our shared piggybank, or that it would be to say that there are two winners of the local 5K, the Johnson family and the García family, because Inez García-Johnson won it. In none of these cases is there any genuine doubling. The window's breaking has just one proximate cause; the 5K has just one winner; the bridge toll has been paid only once. Exclusion begone![6]

[6] Now it can be seen that it is less than optimal for Wilson to characterize the WE/NP solution to the exclusion problem as saying that the effects of mental causes *are* overdetermined, but not in the bad double-rock way—that is, as denying Nonoverdetermination

Unfortunately, this is all a bit of legerdemain. But before I explain why, I need to put the alternative on the table.

5. An Alternative Underlying Idea: The Counterfactual Strategy

Talk of overlapping sets of causal powers is not the only way to explain how various intimate relations between the causes defuse the threat of overdetermination. In a 2003 paper, I offered a different explanation. I provided a necessary condition on overdetermination (genuine, 'double-rock' overdetermination), and argued that it is not met by pairs of causes related in any of the ways WE/NPs think that mental and physical phenomena are.[7]

The necessary condition is simply that two causes overdetermine an effect only if had either happened without the other, the effect would still have occurred.[8] That is, causes c_1 and c_2 overdetermine e only if both of the following counterfactuals are nonvacuously true:

$(c_1 \,\&\, \sim c_2) \,\Box\!\!\rightarrow e$

$(c_2 \,\&\, \sim c_1) \,\Box\!\!\rightarrow e$

This is a very intuitive test for overdetermination. We implicitly rely on it whenever we distinguish between overdetermination and joint causation. Indeed, note that those who would appeal to modal fragility to claim that all apparent overdetermination is really joint causation implicitly rely on these counterfactuals.[9]

Yet if the test is legitimate, the WE/NR is again in good shape. At least one of these counterfactuals will be vacuous or false when (2003) and only when (2008) the mental and physical causes stand in one of the WE/NR's favored relations. Though the details get too complicated to revisit here, the basic idea is that on any such relation, the physical base necessitates the weakly emergent mental phenomena, rendering one of the counterfactuals vacuous.

6. The Relation Between the Two Strategies

Two ways of explaining why the existence of certain tight relations falsifies Exclusion are now on the table. Each strategy offers a necessary condition on overdetermination—one, that certain counterfactuals be nonvacuously true; the other, that the two potential causes not be such that one's set of causal powers is a proper subset of the other's—and claims that weakly emergent phenomena and their

rather than as denying Exclusion. (See note 4). Given the Proper Subset of Powers strategy, she should not think that the effects of mental causes are overdetermined *at all*. For an effect to be overdetermined, it must have at least two distinct causes. But the only sense in which Wilson's WE/NP thinks there are two distinct causes is that there are two distinct phenomena that literally share the efficacious part.

[7] Really, in any of the ways *any* physicalist thinks they are: identity works too.

[8] This is not supposed to be an analysis of overdetermination in noncausal terms, just a condition on which causes count as overdeterminers.

[9] Billy and Suzy throw separate rocks, apparently overdetermining the breaking of the window. The fan of the fragility treatment of such cases (Lewis 1986, 2000) would say, "look I know it seems like the window would still have broken if only Billy threw his rock, or only Suzy threw hers. But that's not actually true, because the precise time and manner of the breaking are essential to it. If only one of them had thrown, it would not have been the very same break. So you're wrong about those counterfactuals. The particular window-breaking that actually happened required both Billy and Suzy to throw their rocks".

bases do not meet the condition, and thus do not overdetermine their effects. Here is a bit more about the relation between these two conditions.

First, the failure of the causal powers to nest in a subset relation does not entail that the overdetermination counterfactuals are nonvacuously true. There are at least two reasons for this. One is that someone who denies that there are any such things as causal powers, or that (foreshadowing!) they are the kinds of countable things that can form sets, will deny that any pairs of events are such that their causal powers nest in the relevant way. But such a person is not committed to thinking that all overdetermination counterfactuals, formulated with whatever pair of events you like, are nonvacuously true. Another reason is the case in which c_1 and c_2 share a lot of causal powers, but not all of them; the two sets overlap but neither is a subset of the other. It could still be the case that one or both of the overdetermination counterfactuals is false or vacuous, for example if the non-shared causal powers are irrelevant to the particular effect in question.

What about the other direction? Does the nonvacuous truth of the overdetermination counterfactuals entail that the causal powers fail to nest in a subset relation? Equivalently, does the subset-nesting of the causal powers entail that at least one of the corresponding overdetermination counterfactuals is false or vacuous? It is tempting to say yes, but matters are somewhat tricky.

Suppose that c_1's causal powers are a proper subset of c_2's, and that c_1 and c_2 are both actual causes of e. It is likely nonvacuously true that if c_1 had happened without the 'larger' c_2, the effect would still have happened. The interesting question is whether e would still have happened if c_2 had happened without the 'contained' c_1. The difficulty in assessing the counterfactual is that the mere claim that c_1's causal powers are a proper subset of c_2's says nothing about the modal status of that inclusion, nor about whether either event has any or all of those token causal powers essentially. The whole shebang could be contingent. And that makes it difficult to mount a decisive case for the falseness or vacuity of the overdetermination counterfactual $(c_2 \,\&\, \sim c_1) \,\Box\!\!\rightarrow e$. The options are that a) c_2 cannot happen without c_1, in which case the counterfactual is vacuous, b) c_2 *can* happen without c_1, and indeed with c_1 and all its causal powers deleted completely, in which case the same counterfactual is probably false, and c) c_2 can happen without c_1 in particular, but only if c_1's causal powers are replaced by numerically different but qualitatively similar ones (in the way that an object might survive the replacement but not complete loss of a part). In *that* case, the counterfactual is probably nonvacuously true, despite the 'subsetting'. And this is the most likely case in the situation at hand: where c_2 weakly emerges from c_1, via any of the standard WE/NP relations. Maybe this mental state could happen without this *particular* physical state that underwrites it, but it cannot happen without *any* physical basis.

Now, I do not want to rest a lot of weight on this. I myself have argued that these kind of 'replacement' interpretations of counterfactuals are problematic (2003: 482), and David Lewis seems to agree (2000: 190). My only point here is that the path from causal-power-subsethood to the falseness or vacuity of the overdetermination counterfactuals is neither obvious nor straightforward. Given the entailment failure in the other direction, it is probably best to think of the two strategies as independent. Two events that vacuify or falsify the counterfactuals need not meet the Proper Subset Condition, and it may well be that two events that meet the Proper Subset Condition can fail to vacuify or falsify the counterfactuals.

7. The Proper Subset of Strategy Is Not More Powerful than the Counterfactual Strategy

I have sometimes thought that the Proper Subset Strategy is a more powerful (groan) implementation of the Counterfactual Strategy. (Both appeared in print at roughly the same time: e.g. Wilson, 1999, 2002; Shoemaker 2001, 2003; Bennett, 2003.) I have come to think that this is wrong. The previous section shows that it isn't clearly right to think of the Proper Subset Strategy as an implementation of the Counterfactual Strategy. And although there is a clear case to be made for the claim that it is more powerful, in two specific senses, this advantage is an illusion.

The first sense in which the Proper Subset Strategy seems to be more powerful than the Counterfactual Strategy is that it appears to provide a deeper, more convincing explanation of why there is no overdetermination. Recall the examples of the bridge toll, the 5k, and the hand-holding hooligans: the weak emergentist gets to similarly claim literally shared causal power. In contrast, the Counterfactual Strategy just says something kind of wishy-washy about the truth-values of certain counterfactuals, while remaining silent about *why* those counterfactuals have the truth-values they do.

The second way in which the Proper Subset Strategy seems to be more powerful than the Counterfactual Strategy is that it not only shows that the weakly emergent entities and their bases can both be causally efficacious without overdetermining their effects, but also shows that weakly emergent phenomena are causally efficacious in the first place. If such phenomena have a nonempty proper subset of the causal powers of their bases, then *a fortiori* they have causal powers.[10] The Counterfactual Strategy, in contrast, does not do this. It simply *assumes* that the mental is causally efficacious, and shows that this (together with Distinctness and Completeness) does not entail that the effects of mental causes are systematically overdetermined.

Unfortunately, these two seeming advantages are just that: mere seemings. There is little substance to either point, which I will address in reverse order.

First, a solution to the exclusion problem that establishes the causal efficacy of the mental, or the weakly emergent more generally, is actually not superior to one that does not—at least, not *qua* solution to the exclusion problem. The exclusion problem is an attempt to undermine the causal efficacy of the mental (the emergent), not because of any intrinsic defect, but rather because there is no causal work for it to do.[11] An adequate response to the exclusion problem is simply one that undercuts this reasoning. My point here is just the elementary one that objecting to an argument that $\sim p$ does not require showing that p is *true*. Thus the fact that the Proper Subset Strategy secures the causal efficacy of the mental does not add anything *qua response to the exclusion problem*.

[10] Wilson admits that nothing she says gives the weakly emergent phenomena *novel* efficacy (58, 67-69), but she is right to accept this consequence. It's what makes weak emergence different from strong emergence. No nonreductive physicalist, for example, should grant causal powers to the mental that aren't possessed by its physical base.

[11] Contrast, for example, Princess Elisabeth-style complaints about substance dualism, where the problem is that the mental is not spatially located, has no mass, has no chemical structure, and so forth.

Of course, this does not mean that it is no advantage at all to the Proper Subset Strategy. It could solve the exclusion problem *and* secure the causal efficacy of the mental. But I am still skeptical; I do not think the strategy actually does secure that. All the work is done by Wilson's claim that weakly emergent entities have a *non-empty* proper subset of the causal powers of their bases. This is the only reason we are guaranteed that weakly emergent entities have causal powers. But Wilson never argues that any particular thing or kind of thing has a non-empty set of causal powers; that is just part of her definition of weak emergence. So those who are inclined to be worried about the causal efficacy of the kinds of phenomena she takes to be weakly emergent—like the mental—will simply deny that they are weakly emergent in her sense.

Second, I also doubt that the Proper Subset Strategy truly provides a deeper, more convincing explanation of why there is no overdetermination—no "causal competition" as Yablo puts it (1992). It looks like it does, yes, but, well, that is the nature of prestidigitation.

The problem is that the deeper explanation requires being quite literal about something that it is not so easy to take literally. The way the Proper Subset Strategy so cleanly escapes overdetermination is by *identifying* each and every causal power of the weakly emergent phenomenon with a causal power of the base phenomenon. As Wilson has emphasized since she began defending the view (1999, 2002), it is crucial that each individual causal power of the emergent thing be possessed by both.

To bring this out clearly, consider two similar but hopeless positions that result from removing the 'subset' part from the Proper Subset Strategy. One position simply says that weakly emergent phenomena have fewer causal powers than their bases. This is no help with exclusion at all; a rock presumably has fewer causal powers than a similarly sized iPhone—for example, only the latter can call an Uber—but throwing both can certainly overdetermine the breaking of a window. The second hopeless position says not only that weakly emergent phenomena have *fewer* causal powers than their bases, but also that their causal powers are *qualitatively indiscernible* from those of their bases. But this again is no help with the exclusion problem. Events with non-identical but qualitatively indiscernible causal powers can absolutely overdetermine things. Consider a scenario in which Billy and Suzy stand 5 feet from each other and throw two indiscernible rocks in indiscernible ways at the window, hitting *almost* the same spot with the same force, at the same angle, at the same time. Their rock-throwings share almost all their causal powers at the type level. (That is, the vast majority of the causal powers belonging to Billy's throw are qualitatively indiscernible from those belonging to Suzy's throw.) But the causal powers of the two events are not numerically identical, and their breaking the window is, again, an uncontroversial case of overdetermination.[12]

In short, the success of the Proper Subset Strategy entirely depends on the idea that the causal powers of the emergent phenomena are numerically identical to the causal powers of the base. And this in turn requires that token causal powers

[12] At this point, one might move to the idea that the causal powers of the base *constitute* or *realize* the distinct but qualitatively indiscernible causal powers of the weakly emergent phenomena. This is basically Derk Pereboom's view (2002, 2011). Whatever its merits, it does not avail itself of the Wilson-Shoemaker idea that there is a shared core of causal power.

are the sort of thing that can not only be *counted* but also *individuated*. Indeed, it is very, very hard not to imagine them as pebbles in a bucket—and Wilson's diagrams on page 70 suggest that she cannot resist this picture either. But this is a serious and rather discombobulating ontological commitment. I will not argue here that causal powers are not like that, but I suspect others will share my reticence. Even Wilson takes pains to insist that her causal powers are nothing dubious or creepy:

> Talk of powers is simply shorthand for talk of what causal contributions possession of a given feature makes […] to an entity's bringing about an effect, when in certain circumstances […] no controversial theses pertaining to the nature of powers, causation, properties, or laws are here presupposed (32-33; also 45).

But the question is, can she really make good on this neutrality? More precisely, can she assuage my ontological qualms while retaining the nice claim that strictly speaking, there is really only one cause of an effect caused both by a weakly emergent phenomenon and its base? That is the challenge I lay before her.

Let me be crystal clear: I have not argued that she cannot meet this challenge. I have simply *raised* the challenge. My real point here is that one cannot have the Proper Subset Strategy on the cheap; the cost-benefit analysis must be made. We can shoulder the ontological commitment to trackable, countable causal powers and accept the benefits, or we can be squeamish and reject the whole picture. What we cannot do is help ourselves to the lovely solution to the exclusion problem while acting as though it costs no more than simply believing in causation. When I accuse the Proper Subset Strategy of sleight of hand, that is what I really mean: not that it cannot fulfill its promise at all, but rather that it hides the expensive machinery required to do so. Regardless, I have appreciated the opportunity to drill deeper into it than I previously have, and discover its secrets.[13]

References

Bennett, K. 2003, "Why the Exclusion Problem Seems Intractable, and How, Just Maybe, To Tract It", Noûs, 37, 471–497.

Bennett, K. 2008, "Exclusion again". In J. Howhy & J. Kallestrup (eds.). Being Reduced, 280–305, Oxford: Oxford University Press.

Lewis, D. 1986, "Postscripts to 'Causation'". In his Philosophical Papers, Volume II, 172–213.

Lewis, D. 2000, "Causation as Influence", The Journal of Philosophy, 97, 182–197.

Pereboom, D. 2002, "Robust Nonreductive Materialism", The Journal of Philosophy, 99, 499–531.

Pereboom, D. 2011, Consciousness and the Prospects for Physicalism, Oxford: Oxford University Press.

Shoemaker, S. 2001, "Realization and Mental Causation", in Identity, Cause, and Mind: Philosophical Essays, Oxford: Oxford University Press, 427-451.

[13] Thanks to Jessica Wilson and the audience at the 2023 Eastern APA Author-Meets-Critics for discussion. I was tempted to somehow work the phrase "metaphysical emergency" into the paper, but I resisted. You're welcome.

Shoemaker, S. 2003, "Realization, Microrealization, and Coincidence", Philosophy and Phenomenological Research, 67, 1-23.

Sider, T. 2003, "What's so Bad About Overdetermination?", Philosophy and Phenomenological Research, 67, 719-726.

Wilson, J. 1999, "How Superduper Does a Physicalist Supervenience Need to Be?", The Philosophical Quarterly, 49, 33-52.

Wilson, J. 2002, "Causal Powers, Forces, and Superdupervenience", Grazer Philosophische Studien, 63, 53-78.

Wilson, J. 2014, "No Work for a Theory of Grounding", Inquiry: An Interdisciplinary Journal of Philosophy, 57, 535-579.

Wilson, J. 2018, "Grounding-Based Formulations of Physicalism", Topoi, 373, 495-512.

Wilson, J. 2021, The Metaphysics of Emergence, Oxford: Oxford University Press.

Yablo, S. 1992, "Mental Causation", The Philosophical Review, 101, 245-280.