

Georgia State University

ScholarWorks @ Georgia State University

Philosophy Theses

Department of Philosophy

Summer 8-1-2013

Causal Compatibilism: A Nonreductive Physicalist Solution to the Exclusion Problem

Morgan Thompson

Follow this and additional works at: https://scholarworks.gsu.edu/philosophy_theses

Recommended Citation

Thompson, Morgan, "Causal Compatibilism: A Nonreductive Physicalist Solution to the Exclusion Problem." Thesis, Georgia State University, 2013.
https://scholarworks.gsu.edu/philosophy_theses/138

This Thesis is brought to you for free and open access by the Department of Philosophy at ScholarWorks @ Georgia State University. It has been accepted for inclusion in Philosophy Theses by an authorized administrator of ScholarWorks @ Georgia State University. For more information, please contact scholarworks@gsu.edu.

CAUSAL COMPATIBILISM: A NONREDUCTIVE PHYSICALIST SOLUTION TO THE EXCLUSION PROBLEM

by

MORGAN THOMPSON

Under the Direction of Eddy Nahmias

ABSTRACT

Jaegwon Kim's Exclusion Problem holds that the nonreductive physicalist position is untenable. If the mental and the physical are distinct and both cause their effects, then it seems that their effects were caused twice over. I argue that the nonreductive physicalist should reject the Exclusion principle—a position called Causal Compatibilism. I appeal to our concepts of causal sufficiency and difference making in order to distinguish cases of mental causation, epiphenomenalism, and overdetermination. I appeal to James Woodward's Interventionist framework to individuate causal difference-makers. Mental causation involves two sufficient causes but only one difference-maker. Given that overdetermination involves two sufficient causes and two difference-makers, the Exclusion principle fails to distinguish between overdetermination and mental causation and so, it is false. I conclude that by rejecting the Exclusion principle, the nonreductive physicalist can get out of the Exclusion Problem.

INDEX WORDS: Mental causation, Exclusion problem, Jaegwon Kim, Interventionism, James Woodward, Causal sufficiency, Causal difference making, Overdetermination

CAUSAL COMPATIBILISM: A NONREDUCTIVE PHYSICALIST SOLUTION TO THE EXCLUSION PROBLEM

by

MORGAN THOMPSON

A Thesis Submitted in Partial Fulfillment of the Requirements for the Degree of

Master of Arts

in the College of Arts and Sciences

Georgia State University

2013

Copyright by
Morgan Kelsey Thompson
2013

CAUSAL COMPATIBILISM: A NONREDUCTIVE PHYSICALIST SOLUTION TO THE EXCLUSION PROBLEM

by

MORGAN THOMPSON

Committee Chair: Eddy Nahmias

Committee: Andrea Scarantino
Dan Weiskopf

Electronic Version Approved:

Office of Graduate Studies

College of Arts and Sciences

Georgia State University

August 2013

ACKNOWLEDGEMENTS

Thank you to all of the colleagues, friends, and family members who have helped me over the past two years. I can't mention each person by name, but I'd like to thank a few people who've helped me the most on this thesis and in philosophy. First, I'd like to thank Eddy Nahmias, who put an immense amount of time and effort into my development as a philosopher. He never hesitated to support me in the numerous projects I've begun during my time at Georgia State. His feedback has helped my thesis progress tremendously throughout the many, many drafts. Andrea Scarantino also deserves recognition for helping me work on the clarity of my writing. His advice about reading and writing philosophy has been extremely helpful. I would also like to thank Dan Weiskopf for pointed and valuable criticism of the arguments in this thesis. Finally, I thank Sebastian Rand for always sharing advice and encouragement.

I would like to thank EJ Coffman, Richard Aquila, Clerk Shaw, David Palmer, and John Nolt at the University of Tennessee, all of who fostered my budding interest in philosophy and continue to encourage me throughout the years. Also, I would like to acknowledge Annette Mendola for our discussions about the problem of the underrepresentation of women in philosophy.

I'd also like to recognize GSU's reductive-physicalist-in-residence, Blake Nespica, for his thoughtful and constructive comments at the Phi Sigma Tau graduate conference. I would like to thank Ed Cox as well for his helpful comments. Thank you to those who provided comments on an early draft of this thesis: Toni Adleberg, Sam Richards, Wes Skolits, Cami Koepke, and Blake Nespica.

Finally, I'd like to thank my peers for all of the intellectual discussions, emotional support, and eagerness to partake in libations. Without all of you, I would have lost sight of the bigger picture. Particularly, I would like to thank Michael Mahoney for holding everything together. He has helped me through these two years and kept me sane. Thank you to Wanda and Kerry, my mom and dad, for always supporting me (even when I switched from psychology to philosophy).

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	iv
LIST OF FIGURES	vi
1 INTRODUCTION	1
1.1 Kim's Exclusion Problem.....	3
1.2 Kim's Causal Inheritance Principle and the Production View of Causation	5
2 THE CAUSAL COMPATIBILIST APPEAL TO INTERVENTIONISM	9
2.1 Causal Pathways and Sufficient Causes	11
2.2 Woodward's Interventionist Account of Difference Making	13
2.3 Performing the Correct Intervention	19
3 USING INTERVENTIONISM TO SOLVE THE EXCLUSION PROBLEM.....	29
3.1 Distinct and Sufficient Causes.....	29
3.2 Overdetermination and Difference Making	31
3.3 Denying the Exclusion Principle.....	35
4 POTENTIAL OBJECTIONS AND CONCLUSIONS	38
4.1 Efficacy Objection	39
4.2 Closure Objection.....	40
4.3 Conclusions	41
REFERENCES	43

LIST OF FIGURES

Figure 1. Kim's Diagram	5
Figure 2. A New Diagram	9
Figure 3. Short Circuit Causal Graph	15
Figure 4. Birth Control Causal Graph	17
Figure 5. Incorrect Birth Control Causal Graph.....	19
Figure 6. Intervening on M	20
Figure 7. Intervening on P.....	21
Figure 8. Intervening on the Common Cause	23
Figure 9. Intervening on M&P.....	24
Figure 10. Firing Squad Overdetermination.....	32

1 INTRODUCTION

Nonreductive physicalists are sometimes accused of holding a set of inconsistent theses. The non-reductive physicalist (henceforth NRP) holds that the mental is “irreducible” (or distinct from) the physical, but that the mental is nothing “over and above” the physical.¹ Most NRPs also hold that the mental *qua* mental is efficacious and that mental causes do not overdetermine their effects. Yet, by adding the *prima facie* reasonable Exclusion principle—that any event with more than one sufficient and distinct cause is overdetermined—it seems that the NRP is committed to an inconsistent set of principles. Here is the putative tension: The more the NRP goes out of her way to establish the distinctness of the mental and the physical, the more the putative mental causes seem to overdetermine their effects; similarly, the more she emphasizes the tight relationship between them, the less efficacious the mental *qua* mental seems. Jaegwon Kim’s Exclusion Argument uses the “tension” between these theses to argue, conditionally, for reductive physicalism (2005, 37).

Kim argues that the NRP can only resolve the tension in her view by accepting one horn of the following dilemma: allowing an overabundance of causes or admitting that mental causes *qua* mental are only *apparent* causes. Kim believes the first horn of the dilemma (the Overdetermination Horn) is a “nonstarter” and “absurd” (1998, 65; 1993a, 281 & 354). He assumes that mental causation simply is not a case of genuine overdetermination. Ultimately, Kim argues that the NRP is left with the second horn of the dilemma (the Epiphenomenalism Horn).

I will introduce the Exclusion Argument as a dilemma for NRP between overdetermination and epiphenomenalism in section 1.1. In section 1.2, I will discuss Kim’s use of the causal inheritance principle and his assumption of the production view of causation.

¹ See, e.g., Wilson 2010a.

In section 2.1, I will introduce the distinction between our causal concepts of sufficiency and difference making, which will be useful for my own solution to the Exclusion problem. In section 2.2, I will then describe the following technical notions from James Woodward's Interventionist account of causation: the notion of possible and impossible interventions and the relationship between causal arrows in the causal graphs and distinct mechanisms in the actual world.

I will use the distinction between causal sufficiency and difference making in section 2.3 to argue for treating the mental and the physical as one variable when building a causal graph to represent the causal difference making relationship in cases of mental causation. In arguing for one particular causal graph and its associated intervention, I will describe and reject other proposed causal graphs and interventions. Of the options I present, only my proposed causal graph is well formed: it is modular, each variable can be intervened upon, and each causal arrow represents only one distinct mechanism.

In section 3.1, I will argue that despite the fact that my solution requires treating the mental and the physical as one variable with respect to causal difference making, the mental and the physical are still nonetheless individually sufficient and distinct. In section 3.2, I will argue that my solution allows the NRP to explain the disanalogy between ordinary cases of overdetermination and mental causation. Here is the basic idea: If the NRP distinguishes cases by both the number of sufficient causes and the number of causal pathways², contra Kim, she can suggest a third option between overdetermination and epiphenomenalism: two sufficient causes, but only one causal pathway. By cashing out this metaphor in terms of causal sufficiency and causal difference making, the NRP can hold that the mental and the physical are each distinct and sufficient causes, but only together are they difference-makers. Then in section 3.3, I will argue that my solution amounts to denying the Exclusion principle.

In section 4.1 and 4.2, I will consider some objections to the effect that the individuation of causal pathways will be vague and that the proposed solution violates the NRP's commitment to Clo-

² This phrase will be explained later in the paper in terms of the mechanisms represented by causal arrows in causal graphs.

sure. Finally, in section 4.3, I will conclude that Kim's presentation of the Exclusion Problem involves a false dichotomy, which unnecessarily restricts the solution space to the Exclusion Problem and misleads him into arguing that NRP is untenable.

1.1 Kim's Exclusion Problem

This section is drawn from Kim's (2005, 33-44; 2008, 111) to present his Exclusion Argument.

The dilemma begins by assuming the NRP's position: Closure (required by physicalism) and Distinctness (required by non-reductionism).

Closure If a physical event³ has a cause at *t*, then it has a physical cause at *t*.

Distinctness Mental events are numerically distinct from physical events.

The NRP cannot reject either Closure or Distinctness on pain of collapsing NRP into dualism or reductive physicalism, respectively.⁴ Merely by adding some plausible principles—Exclusion, Efficacy, and No Overdetermination—we've got the Exclusion Problem:

Exclusion If an event at *t* has more than one distinct, sufficient cause at *t*, then it is genuinely overdetermined.

Efficacy Mental events have causal efficacy—that is, their instantiations can, and do, cause other events, both mental and physical to be instantiated.

No Overdetermination Mental causation does not involve overdetermination in the way that firing squads do.

³ Following Kim, I aim to stay neutral as to whether the argument is targeting properties or events, so I will use whichever sounds the most natural. Some Causal Compatibilists (e.g., Bennett 2003) take distinctness to imply a modal distinctness between mental and physical states. Here I only hope to maintain the minimal thesis that the mental and the physical are not identical. Note that the acceptance of distinctness as it is phrased here does not commit the NRP to the dualist notion that the mental and the physical states are independent substances. One way to flesh out the distinctness of the mental and the physical without committing oneself to the independence implied by substance dualism is to accept that a supervenience relation holds between mental and physical states. When A supervenes on B, there can be no change in A without a change in B (McLaughlin and Bennett 2011). By metaphysical necessity, any change in A will be accompanied by a change in B.

⁴ It may be possible for the NRP to reject Closure by arguing that the laws of physics make no appeal to causality, but this discussion is outside the scope of this paper.

Adapted from Kim's (1998) Supervenience argument, the Exclusion Argument goes as follows⁵:

- (1) Assume mind-body supervenience: Mental properties supervene on physical properties in the sense that if something instantiates any mental property M at t, there is a physical base property P such that the thing has P at t, and necessarily any thing with P at a time has M at that time.
- (2) Efficacy: Suppose that an instance of a mental property M causes another instance of a mental property M* to be instantiated.
- (3) M* has a physical supervenience base (by 1).
- (4) M* is either instantiated in this case because: (a) ex hypothesi, M caused M* to be instantiated, or (b) P*, the physical supervenience base of M*, is instantiated on this occasion.
- (5) M caused M* by causing P*, or mental-to-mental causation implies mental-to-physical causation.
- (6) But M itself has a physical supervenience base P, which is also the cause of P* (by Closure).
- (7) Distinctness: P and M are numerically distinct properties.
- (8) So, P caused P*, and M supervenes on P and M* supervenes on P*.
- (9) Exclusion: If an event at t has more than one distinct, sufficient cause at t, then it is genuinely overdetermined.
- (10) No Overdetermination: But mental causation does not involve the type of genuine overdetermination found in firing squad cases.
- (11) So, M-to-M* and M-to-P* causal relations are only apparent, and this appearance arises from the genuine causal process from P to P*.
- (12) Therefore, the NRP position is untenable; Efficacy, Closure, Distinctness, Exclusion, and No Overdetermination are inconsistent.

Kim's ultimate recommendation is to give up on the *nonreductive* part of the NRP's view

(namely, Distinctness) and accept reductive physicalism, which holds that mental properties are not distinct from physical properties and inherit their causal powers from their supervenience base, given that the mental can be successfully reduced.⁶ The NRP might question Kim's insistence that she must give up either Efficacy or No Overdetermination.

⁵ Here I'm leaving aside explanatory versions of the argument.

⁶ Throughout I will ignore Kim's position on qualia, which is that qualia are not functionalizable, and so they are irreducible.

Here is Kim's diagram illustrating the Exclusion Problem (2005, 45)⁷:




Figure 1. Kim's Diagram. The arrows with diamond-heads represent the relation of supervenience with the diamond facing the supervening event. For each set of nodes related by supervenience, the individual nodes are distinct but *not independent*. The arrows represent causal relationships with the arrowhead facing the effect.

Kim's main concern is whether there is a non-overdetermining arrow from M to P*. Kim's diagram leads one to believe there are only two options for the NRP: overdetermination as depicted in Figure 1 or epiphenomenalism where the arrow from M to P* is only apparent. I will argue that these are not the only options—instead, the NRP can reject the Exclusion principle.

1.2 Kim's Causal Inheritance Principle and the Production View of Causation

Kim (1993a) appeals to the principle of causal inheritance to explain why both Closure and Exclusion result in the epiphenomenalism of M with respect to P*. He defines the causal inheritance principle as follows (1993c, 326; italics removed):

if mental property M is realized in a system at t in virtue of physical realization base P, the causal powers of this instance of M are identical with the causal powers of P.

Kim thinks that anyone who accepts this principle is denying distinctness and thus, a reductionist (1993a, 356).

⁷ Here I'm leaving aside issues about mental-to-mental causation because each version of Kim's Exclusion Argument turns crucially on the status of mental-to-physical causation (or "downward causation").

Suppose instead that the nonreductive physicalist chooses to deny the principle of causal inheritance. Kim thinks that anyone who wants to deny this principle is thereby committed to emergent causal powers. In another paper, Kim describes this principle in terms of Alexander's dictum, to be real means to have causal powers (1993a, 350; his italics removed and my italics added):

on Alexander's dictum, that mental properties bring with them *new causal powers*, powers that no underlying physical-biological properties can deliver. For unless mentality made causal contributions that are genuinely novel, the claim that it is a distinct and irreducible phenomenon over and beyond physical-biological phenomena would be hollow and empty. To be real, Alexander has said, is to have causal powers; to be real, new, and irreducible, therefore, must be to have *new*, irreducible causal powers.

Kim is not the only one to suggest that the NRP usually holds the causal inheritance principle.

Michael Baumgartner illustrates the way that many NRP think about the additional causal powers of the mental (2013: 4; emphasis added):

Many non-reductive physicalists additionally assume something along the lines of the Eleatic Principle (cf. e.g. Armstrong 1997; Oddie 1982) according to which *everything that exists (in space and time) has causal powers*. From the fact that mental properties constitute an *ontological category that is distinct and nonreducible* to the domain of the physical [...] *they then infer that mental properties have their own non-reducible causal powers*. That is, mental properties do not have causal powers merely by virtue of being physically realized but by their own right, i.e. they have genuinely mental causal powers. This is known as the principle of the Causal Autonomy of the Mental (CAM) (cf. e.g. Fodor 1989, Lowe 1993, Antony 2007; Menzies and List 2010). Here is the version of the principle advanced in Menzies and List (2010, 111):
(CAM) For some mental property M and physical property P, where an instance of property M is realized by an instance of property P, the causal powers of the M-instance are not a subset of those of the P-instance.

Notice that in this quote, Baumgartner suggests that many NRP jump from a principle that suggests everything that exists has causal powers (Kim calls this Alexander's dictum) to the CAM principle. However, it is unclear that NRP should accept that each ontologically distinct category has distinct and nonreducible causal powers. As long as the mental still has causal powers, it does not violate the Eleatic Principle, as Baumgartner has stated it. I do not want to take a stand on the truth or falsity of CAM, but I do wish to suggest that the NRP does not *need* to accept it.

Kim (2002, 675) explicitly admits that he assumes a production view of causation, which states that causation involves a physical causal process that transmits energy to or marks the effect (e.g., Salmon 1984, Dowe 2000). Production views are taken to rule out omissions as causes. Throughout his writings, it is quite clear that Kim relies on a production view of causation throughout the various forms of the Exclusion argument. He also appeals to notions of causal work: “Given that P has a physical cause P*, what causal work is left for M to contribute?” (1998, 37). He also suggests that the “causal contribution” of M is “totally mysterious” (2005b, 48). Kim’s conclusion is that since M does no additional causal work to bring about P* (and accepting that mental causation does not involve overdetermination), M does not actually cause P. As Kim sometimes says, the mental is dispensable (1998, 45). Some have argued that because Kim assumes a production view of causation, the Exclusion problem could be resolved by adopting a different view of causation.

Loewer (2007) has argued that Kim’s assumption of a production view of causation is problematic by appealing to Bertrand Russell’s argument that *productive* causal relations do not appear in the laws of physics.⁸ Loewer (*ibid.*, 16) states:

My proposal is not that Lewis’ [counterfactual] influence account perfectly captures our intuitive concept of causation. But I do claim that causation as influence is near enough to our folk conception of mental causation to underwrite the role of causation in folk psychology, rational deliberation, action theory and so on.

He then argues that our best understanding of causation (or at least the folk view of causation, which is supposed to ground agency attributions) is found in dependence views (or the counterfactual influence account described by Loewer above).

However, it is unclear what is actually our best theory of causation (or even our best theory of the folk theory of causation), so it would be hasty to rule out production views.⁹ Instead, I take Karen

⁸ This is a contentious point, see Frisch 2009. For arguments supporting Russell’s view, see Norton 2007, 2009.

⁹ On the actual nature of causation, see Lewis (1973) for support of dependence views and see Dowe (2000) for support of production views. On the folk concept of causation, see Gopnik and Schultz (2004) for some support of dependence views and see Wolff (2008) for support of production views.

Bennett's (2008) line on this: choosing a theory of causation in order to solve the Exclusion problem is just the wrong way to go about it. She says: "Whether it can [i.e., dependence accounts of causation can help alleviate worries like the Exclusion problem] or not depends on whether it is a good enough as an account of causation full stop (2008, 24). Even Kim feels that the Exclusion problem will exist regardless of one's theory of causation (1998, 43). This is not to say that one's theory of causation is irrelevant to the Exclusion problem, but rather that one's theory of causation *alone* will not solve it. Hence, we need to consider other ways to approach the problem.

It is important to note that admitting the move from rejecting the causal inheritance principle to positing emergent properties and holding a production view of causation are not required to get the Exclusion Argument off the ground (contra, e.g., Loewer 2007). Given that neither view is necessary for the Exclusion Argument, any part of the argument that rests on one of these principles is controversial at best. In fact, the NRP is able to deny such grounding if she, for example, holds a dependence view of causation. In doing so, she will not have solved the Exclusion Problem. So, in what follows, I will present an independent solution to the Exclusion Argument. However, when Kim's argument relies on controversial principles about causal inheritance or theories of causation, I will suggest getting rid of such reliance.

2 THE CAUSAL COMPATIBILIST APPEAL TO INTERVENTIONISM

In this section, I will appeal to a distinction between causal sufficiency and difference making that will help set up a NRP solution to the Exclusion problem. I will introduce some more formal notions of difference making in order to describe the different ways that our concepts of causal sufficiency and causal difference making apply to mental causation.

Karen Bennett (2003) and others (e.g., Horgan 2001) use the label ‘Compatibilism’ for the view that the mental and the physical do not causally compete, but both are nonetheless efficacious. Since the Causal Compatibilist accepts both No Overdetermination and Efficacy, and given the five theses are inconsistent, she will have to reject Exclusion. Kim assumes that it is not possible to have more than one cause per causal pathway, which leads him to support the application of the Exclusion principle to all cases that do not involve overdetermination.

Recall Kim’s diagram of the Exclusion Problem (Figure 1). Both (non-causal) supervenience relations and causal relations are represented with arrows (with distinct types of arrowheads), which has the potential to be misleading. Perhaps the diagram should look something more like the following¹⁰:



Figure 2. A New Diagram. M and P are distinct causes and M* and P* are distinct effects. There is only one causal relationship between M to M* (and since P and P* contain M and M*, respectively, the causal relationship also between P to P*).

Both M and P are distinct and sufficient causes of M* and P*, respectively¹¹, but there is only one causal pathway between the two. The NRP needs to show that causal difference making and causal sufficiency

¹⁰ This diagram is originally found in Alyssa Ney’s (2010) description of Sydney Shoemaker’s view. It could also represent Jessica Wilson’s (2010b) proper subset of causal powers view.

come apart in particular cases, so she can rely on the individuation of causal events both in terms of ‘causal pathways’ and ‘sufficient causes.’ I will define causal pathways as difference-makers in a later section. As Figure 2 suggests, there is an important sense in which the term ‘downward causation’ is problematic. Downward causation is mental-to-physical causation and it is represented by the arrow on Kim’s diagram (Figure 1) from M to P* and figures into premise (5) of the Exclusion Argument. In Figure 2, the supervenient properties and their supervenience bases are no longer represented on different “levels” of the diagram. Further, Figure 2 suggests that it is unclear why so-called “downward causation” requires an *extra* causal arrow (or more causal “oomph” according to the production view of causation).¹² So, figure 2 has the added benefit of removing an unnecessary reliance on the production view of causation.

I would like to stress that I am not suggesting Figure 2 is the only way to represent the causal structure of cases of mental causation (without relying on “levels”). Perhaps Figure 2 lends itself too easily to particular nonreductive physicalist positions (i.e., subset views of causal powers) and some NRP might find these views to be problematic for various reasons¹³. Instead, perhaps Figure 2 should represent M and P as completely overlapping, but this might lend itself to a reductive interpretation. The most important point is that these diagrams (including Figure 1 and 2) are misleading and, in fact, may be unable to represent all distinct metaphysical views. It may be better not to use diagrams at all, however, I will continue to use those such as Figure 2 with the caveat that we should not make metaphysical inferences from these diagrams. (Note that these diagrams are importantly different from the causal graphs used in the literature on interventionism and causal Bayes nets (see Woodward 2011).

¹¹ One might suggest this seems like a naturalistic version of parallelism, but this is what the notion of a causal pathway is supposed to avoid. Situations of parallelism will involve two *entirely* independent causal pathways, which cause two different effects.

¹² For an alternative to a levels-ontology based on levels in science, see Barnes (Forthcoming) on fundamentality and see Schaffer (2009) on grounding.

¹³ Some argue that the physical state’s set of causal powers can actually be identified with the “subset” of causal powers these NRP identify as comprising the mental state. See R. Wilson (2001) for discussion.

I will now provide a brief sketch of the causal Compatibilist solution, which I will describe in more detail in the following subsections. First, I propose that the causal Compatibilist should appeal to two different notions of causation: sufficiency and difference making. The Exclusion argument is normally described in terms of causal sufficiency and it says nothing about causal difference making. By appealing to these two notions, the causal Compatibilist can stake out some ground for mental causation between overdetermination and epiphenomenalism. To do this, she will need to give an account of overdetermination, mental causation, and epiphenomenalism in terms of both causal sufficiency and causal difference making. If each case varies in its overall response to these two cases (and the causal Compatibilist argues that it does), then she will have shown that the Exclusion principle is false. Recall Exclusion states that if an effect has more than one distinct and sufficient cause, then it is overdetermined. In order to claim that mental causation is neither a case of overdetermination nor a case of epiphenomenalism, the Causal Compatibilist has argued that mental causation involves two sufficient causes, but only one difference-maker. Overdetermination, on the other hand, involves two sufficient causes and two difference-makers. So, the causal Compatibilist can reject the Exclusion principle because cases of mental causation involve more than one distinct and sufficient cause, but which are not cases of overdetermination. However, much more needs to be said about the differences between causal sufficiency and difference making and about the causal structure of mental causation on this account.

2.1 Causal Pathways and Sufficient Causes

For the purposes of this thesis, a ‘causal pathway’ is a difference-maker. Relative to a given context, varying the occurrence of the event (a ‘cause’ on the interventionist theory of causation) will vary the occurrence of the effect. If the first event occurs, then the effect should also occur. Similarly, if the

first event does not occur, then the effect should not occur.¹⁴ Many scientific experiments test for difference-makers.

Larry Shapiro and Elliott Sober (forthcoming) describe the example of a geneticist who tests whether phenotypes make a difference with respect to genetic inheritance of traits. Suppose a parent is born with both genotype G and phenotype P. The parent has developed phenotype S, which is mutually exclusive with phenotype P, due to environmental factors during the course of its life. Will the parent's offspring have phenotype P or S? Holding fixed the genotype G, the offspring will have phenotype P, not S. For example, just because the parent, who was born rather scrawny, developed a muscle-heavy build due to labor-intensive work, the offspring will not be born with a muscle-heavy build. Phenotypes almost never make a difference with respect to genetic inheritance of traits. On the other hand, altering the genotype G to an alternative genotype H will make a difference to both the genetic expression in and phenotype of the offspring.¹⁵

Alternatively, a 'sufficient cause' is an event that is causally sufficient for the effect, given certain background conditions. Some event's being a cause in this sense only tells us something about the effect when the cause actually occurs. For example, paying my bill online is sufficient for avoiding a late fee. However, this says nothing about situations when I don't pay my bill online. Perhaps I decided to pay my bill by mail, which would still avoid a fee. It is often true that the number of sufficient causes is the number of difference-makers in any causal structure, but Compatibilists hold this is not *always* the case.

Note that nomological sufficiency is not the same as difference making. Woodward (2008) states that nomological sufficiency "is a claim about what would happen if the cause were present but says nothing about what would happen if the cause were absent or different" (246). One of the major differences between these two causal concepts is that knowing an event is sufficient to produce an effect

¹⁴ Difference making is the general idea behind the interventionist view of causation (see Woodward 2003).

¹⁵ Admittedly, the difference-making capacities of genes are more complicated than I've described them here. For discussion, see Woodward (2010).

merely shows that the event *could* be the cause of the effect. However, we might insist that we still don't know if the event was the *actual cause* of the effect.

So far I've simply sketched an intuitive distinction between sufficiency and difference making. In the next subsection, I will provide a more formal account of causal difference making by appealing to Woodward's interventionist account of causation. Interventionism will give the NRP a formal framework for individuating the number of difference-makers in a given causal situation. Doing so will allow me to argue that the NRP can individuate causal cases both in terms of sufficient causes and in terms of difference-makers. If there are cases where these two concepts of causation individuate a different number of causes, then the NRP has good reason to suppose that causal situations should be classified based on the number of sufficient causes and the number of difference-makers. In particular, I will argue that mental causation is a causal situation in which there are two sufficient causes but only one difference-maker. This solution will maintain that mental causation does not involve genuine overdetermination of the kind found in firing squad cases, but also that it does not amount to epiphenomenalism.

2.2 Woodward's Interventionist Account of Difference Making

While the intuitive distinction between sufficiency and difference making hints at a way to distinguish between overdetermination, epiphenomenalism, and mental causation, the causal Compatibilist will need to provide a more robust account of the difference making relation. She will need an account that clearly individuates causal difference-makers such that we can count the number of instances in a given causal structure. I will appeal here to James Woodward's formal interventionist account of causation as presented in *Making Things Happen* (2003), which draws from previous work by Judea Pearl (2000). Woodward's account of difference making highlights the notion of causation often taken for granted in scientific experiments where an independent variable is manipulated and its effects on a dependent variable are observed in relation to a control group. It is important to note that these inter-

ventions need not *actually* be performed because sometimes we cannot actually manipulate the variable or it is unethical to manipulate them.

Woodward (2011: 7) states that the intuitive notion of an intervention is “an idealized, unfounded experimental manipulation which would be appropriate for determining whether one variable is causally related to a second variable.” Woodward (2003: 59) describes in more detail the interventionist notion of difference making:

A necessary and sufficient condition for X to be a (type-level) direct cause [in the difference making sense] of Y with respect to a variable set \mathbf{V} is that there be a possible intervention on X that will change Y or the probability distribution of Y when one holds fixed at some value all other variables Z_i in \mathbf{V} .

The idea is that if X is a difference maker with respect to Y , then (holding fixed other variables) any change in the value of Y will be accounted for by an earlier change in X . This causal claim can be tested by representing the causal structure as a directed causal graph (or sometimes as a series of mathematical equations), with ‘nodes’ for each variable in the system. Each variable can be set to at least two values (e.g., occurring, not occurring). Interventions should be treated as instances where the arrows going into some particular node are “broken” or “wiped out” (2003, 39). Any other variables, Z , that have causal arrows in the directed graph entering the variable Y will be held fixed at some particular value. That is, the causal contribution of variables Z to Y are “controlled for”, or held constant while the values of X are manipulated.

In addition to describing what interventions are in general, the interventionist will need an account of what makes a causal graph, or causal diagram, well formed. Being a well-formed causal graph involves: (1) accurately representing the causal structure that is being examined and (2) allowing only legitimate interventions.

(1) will require representing each causal mechanism separately and appropriately representing each node. Consider Woodward’s short circuit fire example (2003, 44-5). Let the oxygen level Ox be 1 or 0 (present, not present), the short circuit S be 1 or 0 (present or not present), and the fire F be 1 or 0

(present or not present). When Ox is set at 0, then regardless of the value of S , F will be 0. However, the interventionist account only requires that there be some intervention on which S makes a difference. So, set Ox at 1 and by an intervention on S , we will see a change in the value of F . However, the interaction between the variables Ox and S will not be represented in a causal graph:




Figure 3. Short Circuit Causal Graph. Both S and Ox have causal mechanisms, represented by arrows, by which they affect the occurrence of F . To correctly test for the causal efficacy of S , we need to intervene on S , but only while Ox is set at 1. The circle-headed arrow represents an intervention.

Note that the arrows here correspond to distinct mechanisms. Although it may be natural to interpret ‘mechanisms’ here as a physical causal process (e.g., transmitting of a mark), it is left unspecified whether these mechanisms involve causal *production* per se.¹⁶ The interventionist view is non-reductive with respect to causation insofar as it does not seek to explain difference making in terms of a primitive notion of causation (Woodward 2003: 20-21). That is, interventionism does not define a difference making cause strictly in non-causal terms.

Unfortunately, Woodward does not define the use of ‘mechanisms’ here. I don’t take it to be a technical term for interventionists, so perhaps we can simply appeal to an intuitive notion of mechanism. Often they will involve some physical process or activity that creates a change in the outcome var-

¹⁶ Elsewhere Woodward (2011) explicitly states that mechanisms should not be interpreted as physical in this sense.

variable. For example, activation of one brain area by another may be caused through a mechanism that involves a certain pattern of firing (e.g., phasic bursts above a certain threshold) by the neurons that connect the first brain area to the second. Yet, we might also think that mechanisms can involve a lack of activity as well. Perhaps the connection between the first brain area and the second is inhibitory; when the neurons from the first brain area fire to the second brain area, the activation of the second brain area is inhibited. Once the inhibitory firing from the first brain area to the second brain area ceases, the second brain area becomes activated again. It seems that the second brain area's activation is caused by the decreased firing rate of neurons coming from the first area.

Another aspect of judging when a given causal graph is well formed requires describing (2) when interventions are legitimate. According to Woodward, it must be possible to intervene on one variable and not others, and vice versa (2003: 48, 52). Another way of stating this requirement: "if a candidate causal claim is associated with interventions that are impossible for (or lack any clear sense because of) logical, conceptual or perhaps metaphysical reasons, then that causal claim is itself illegitimate or ill-defined" (2008: 9). So, illegitimate interventions will generally involve confused causal claims. It is important to get the direction of causation correct here. These causal claims are not confused *because they involve illegitimate interventions*. Instead, considering illegitimate interventions should be an indicator that one's causal claim is confused. For example, intervening on two conceptually identical claims to test whether, e.g., a bachelor or the unmarried man lives at such-and-such address is simply a confused causal claim. There is a sense in which the two do not really causally compete.

For example, it would require an impossible intervention to hold the variable drinking H₂O at 0 while testing whether or not drinking water makes a difference to staying hydrated. We cannot intervene on drinking water in this way because it would require assessing the causal outcome of drinking water without drinking H₂O. The reason this intervention is illegitimate, according to Woodward, is because the causal claim itself is illegitimate. It is not possible to ask whether one and the same thing

makes a difference to an effect because this requires testing an intervention when the thing simultaneously occurs and does not occur—a metaphysical impossibility. Note that confused causal claims will not always involve confusion about identical objects with different names. Any metaphysical reason that restricts the types of interventions that can legitimately be performed will involve confused causal claims. Later in this paper (Section 2.3) I will suggest that interventions that require breaking supervenience relationships are based on illegitimate causal claims.

Let's look at another (more complicated) example from Woodward (2003). Suppose that taking birth control pills (B) causes an increase in the probability of getting thrombosis (T). However, taking birth control pills also significantly reduces the probability of becoming pregnant (P), which itself causes an increase in the probability of getting thrombosis. This causal structure can be depicted with the following diagram (2003: 50):




Figure 4. Birth Control Causal Graph. There are positive relationships between B and T and P and T, but a negative relationship between B and P. So, B does not have a net effect on T.

As we can see, there is no net effect of taking birth control on the probability of getting thrombosis. However, we can “see” the distinct mechanisms, from B to T, and from B to P and then from P to T. The representation of these mechanisms by causal arrows reflects causal mechanisms we know exist. For example, we know that taking birth control reduces the probability of getting pregnant (under the right

conditions, e.g, for women). We also know that getting pregnant increases one's risk of getting thrombosis. Finally, we know that taking birth control pills increases one's risk of getting thrombosis.

We might wonder in what sense B makes a difference to T. Woodward appeals to a distinction between total causes and contributing causes to explain the sense in which B is not a total cause of T, but is still a contributing cause of T. A total cause is defined as follows (50):

X has a non-null total effect on Y, or there is some intervention on X alone (and no other variables) such that for some value of other variables besides X, this intervention on X will change Y.

In this sense of 'cause', B is not a total cause of T. There is no single intervention on B that will make a difference to the occurrence of T. Intuitively, there seems to still be a sense in which B is a cause of T.

Woodward calls this sense a contributing cause (50):

X makes a non-null contribution to Y along some directed path in the sense that for those variables (if any) that are not on this path, there is some set of values of those variables such that if the variables were fixed by interventions at those values, there is some intervention on X that will change the value of Y.

'Directed path' describes those pathways through the causal diagram with at least one variable between X and Y (42). In the birth control case above, P's being along the way from B to T makes it a directed path. So, in Interventionist terms, B is not a total cause of T (as it can also be caused by being pregnant), but B is a contributing cause to T. As the example about birth control's influence on instances of thrombosis shows, asking causal questions merely about total causes will not be sufficient to adequately describe many causal structures.

Direct causes (but not total causes) can help distinguish distinctness of causal mechanisms. Woodward states: "if it is really true that the mechanism by which B influences T is distinct from the mechanism by which B and P influence T, then it should be possible to interfere with one of these without interfering with the other" (52). So, the following representation would be very misleading:




Figure 5. Incorrect Birth Control Causal Graph. This causal graph shows B as influencing T only along one direct route.

It would collapse two distinct mechanisms into one. Often we might need to perform particular interventions to isolate different mechanisms. Woodward suggests that we look for the causal efficacy of B on T by disrupting the inhibitory relationship between B and P on Figure 4 by studying the instance of thrombosis in either infertile women taking birth control or in women already pregnant who nonetheless continue taking birth control pills. This would involve fixing the value of B at 1 and varying the value of P.

These formal interventionist concepts will give the Causal Compatibilist the tools to describe the correct intervention in the case of mental causation. Performing the correct intervention in cases of mental causation will allow the NRP to count the number of difference-makers in those cases—in effect, answering whether M, P, or both make a difference to the occurrence of both P* and M*. As we will see in the following section, the Compatibilist will need to provide an intervention that (1) does not violate metaphysical relationships. She will also need to make sure that the causal graph (2) allows for interventions on each variable.

2.3 Performing the Correct Intervention

Many of the interventions suggested implicitly in the literature on mental causation are illegitimate given the interventionist tools described above. In this section, I will compare interventions proposed by others with the solution I put forward. I will argue that my solution is does not involve illegitimate interventions and does not violate the NRP's commitment to supervenience whereas other interventions are illegitimate and violate supervenience.

It may be that Kim's traditional representation of the exclusion problem (Figure 1) has misled many philosophers in their attempts to apply interventionist insights to the exclusion problem. Interventionism relies on formal causal graphs (adapted from Bayes nets) where all arrows represent causal relationships. It is indeed an extension of the interventionist framework to evaluate causal relationships for diagrams that include both causal and non-causal relationships (Woodward 2011). This technical definition of causal graphs does not preclude extending the interventionist framework to accommodate both causal and non-causal relationships, but rather that we must be careful in doing so. First, we might think that we can intervene on M directly to discover whether M "causally contributes" to P*'s occurrence. To do this, we would have to sever the tie between M and P. Recall that performing an intervention on one variable requires holding fixed all variables on other directed paths. In this case, we would be treating the directed path from P to P* (and M*) as distinct from the directed path that goes from M to P* (and M*).




Figure 6. Intervening on M. This figure comes from Woodward (2008). The circle-head arrow represents an intervention on M. The diamond-head arrows represent supervenience. The arrow represents a causal relationship. The double lines crossing out the supervenience relation between P and M represent "breaking the chain" required by interventions. In this case, M is represented as not occurring.

That is, we fix M at some particular value (perhaps: occurring), while allowing P to continue its normal causal mechanisms, and observe whether P* occurs. Then we fix M at another value (not occurring) and again observe whether P* occurs. In both cases, we would find that there was no difference in P*, since it occurs via the causal mechanism from P. This might lead some to believe that M does not actually cause P* because P is the true cause of P*. However, this intervention violates the supervenience rela-

relationship between M and P. There is a change in M without a change in M's supervenience base P. The main problem with this view is that the supervenience relation is not causal (Woodward 2008, 2011). Correctly performing interventions requires breaking *causal* relations. As described in the previous section, I take it that those interventions that require breaking supervenience relations are illegitimate because they require an impossible intervention. Perhaps the proponent of this intervention as the correct test for the difference making of the mental would suggest that the supervenience relation is causal, but this position would be contrary to Kim's own position (2005, 48).

One intervention described (but not defended) by Larry Shapiro (2010) is often seen in the literature. Testing the difference making of M by holding M fixed while manipulating P. Note that the intervention currently under discussion does not violate the supervenience relationship between M and P. If M supervenes on P, it is still possible to have a change in P without a change in M. In fact, many NRP believe in multiple realizability, which is the thesis that many P-changes do not result in M-changes.




Figure 7. Intervening on P. In this picture, intervening on P requires “breaking the connection” between P and M.

However, I will argue that this intervention is also illegitimate. There are two ways to interpret the case in which M occurs but P does not: either M has no supervenience base at all or M has a supervenience base P' , which is not identical to the physical state P. The former is problematic because it violates Closure. If M occurs without P (or any other suitable supervenience base) and still causes P^* and M^* , then there is a physical state (P^*) with no physical cause. Crisp and Warfield (2001) argued for the other view:

that M must have a supervenience base, P', which is not identical to the physical state P. This strategy is consistent with many NRPs' commitment to multiple realizability.

Unfortunately, replacing P with another supervenience base, P', is problematic for two reasons. First, it fails to answer the main question at stake in Kim's version of the Exclusion Problem: Given non-reductive physicalism, do mental states non-redundantly causally contribute to their effects? If we evaluate the causal difference making of some particular mental state M without one of its possible supervenience bases P but allow another supervenience base P' to fill its place, then we haven't addressed Kim's question. In fact, Kim's question seems to arise again: Does M make a non-redundant causal contribution beyond the contribution its supervenience base, P', to their effects? To test the difference making of M where P' does not occur, we again might posit another supervenience base P''. However, Kim's question can arise once again, and so on.

The second reason it is problematic to replace P with another supervenience base P' in order to evaluate the difference making of M has to do with evaluating counterfactuals. Bennett (2003) has argued that Crisp and Warfield's positing of physical base P' in supervenience world w changes the topic. Instead, we should think of accessing these counterfactuals as using a "metaphysical hole puncher" where we literally think of the supervenience base as simply gone (2003, 15). The idea behind the metaphysical hole puncher is to rule out backtracking counterfactuals.¹⁷ David Lewis has said that when we imagine some event C is deleted, "we imagine that C is completely and cleanly excised from history, leaving behind no fragment or approximation of itself" (2000, 190). Given that both P and P' are appropriate supervenience bases of M, I take P' to be an approximation of P.

¹⁷ Backtracking counterfactuals fail to hold fixed all other causes in the scenario. Say we would like to access the counterfactual where one shooter (out of two) fails to fire at a victim. When we ask "does the victim still die?", we should refrain from describing a *reason* why one shooter fails to shoot. For example, it would involve backtracking reasoning to suppose the shooter was startled by a dog's bark and so did not fire. Following that line of reasoning, we might suspect the other shooter does not fire either, so the victim does not die if the first shooter does not fire.

Notice too that replacing P with P' is perfectly consistent with the epiphenomenalism of the mental. It could have been P (or some other suitable supervenience base P') doing the work all along while M just "goes along for the ride." So, if we're trying to figure out, as Kim asks, if the mental actually makes an additional causal contribution with respect to P*, then we ought not to evaluate situations in which P' replaces P during an intervention. When we control for the influence of P, we shouldn't add anything new to the causal situation. Note that Kim's question in the Exclusion Problem is given at the type level whereas replacing P with P' can only address Kim's worry at the token level.

It may seem intuitive that an intervention that requires holding fixed M while intervening on P will answer Kim's question about the causal contribution of the mental. Particularly, the intervention should show whether or not M causes P* when P is not present. I've described two interpretations of this situation: either M has another supervenience base on this occasion, P', or it does not. If M does not have a supervenience base but still causes P*, then Closure is violated. If M does have a supervenience base P' and still causes P*, then I've argued that we have failed to answer Kim's question. Kim might ask whether M occurring without P' still causes P*. I've argued that this strategy will be repeatedly open to the Exclusion Problem; instead of solving the problem, this strategy simply changes the question from type-level to token-level. Given that both interpretations of M occurring without P are problematic, I argue that we should continue searching for an appropriate intervention.

Shapiro suggests an alternative intervention. He argues that the correct test is to intervene on the common cause, P₀, of both M and P (*ibid.*).




Figure 8. Intervening on the Common Cause. This intervention involves breaking the causal chain between P 's cause, P_0 , and P . Just like Figure 7, it fails to break the causal connection from M to P^* .

Manipulating P_0 will never create a case in which M occurs and P doesn't, and vice versa. It is impossible to test whether M makes a difference when P does not occur. Shapiro admits that the supervenience relationship cannot be broken. Unfortunately, I believe that Shapiro's proposed intervention still involves the misleading representation of the supervenience relation that results from Kim's diagram (Figure 1). Looking at Figure 8, it is too easy to suggest this picture looks like overdetermination.

In more technical terms, the two causal arrows in this picture suggest that there are two causal mechanisms at work. The NRP need not accept this. Recall that in Section 1.3 I suggested the NRP should reject any principle that relies on a production view of causation or the entailment of positing emergent causal properties given one rejects the causal inheritance principle. To suppose that the NRP must posit a separate causal mechanism (i.e., that the mental must do "extra causal work" compared to the physical) is to rule out many NRP views. If the Exclusion Argument relies on these views, then it begs the question against many NRPs.

Further, the causal diagram proposed by Shapiro's contains nodes that cannot be intervened upon, which violates the principle of impossible interventions. While he admits it is metaphysically impossible to intervene on M or P independently (since this amounts to breaking the supervenience relationship), his causal diagram has two variables that cannot be intervened upon individually.

I argue that the correct intervention should treat M&P as a variable together. That is to say, M and P are difference-makers only when they occur together. There should only be one causal mechanism between M&P and M*&P*. Although it may seem like I've denied Distinctness, I will argue against this claim in section 3.1. For now, consider the following diagram:




Figure 9. Intervening on M&P. Here we intervene on M&P by varying their value together. There is only one “causal mechanism” from M&P to M*&P*. M&P will count as a difference-maker with respect to M*&P*.

We can intervene on M&P by holding it fixed at occurring, which will result in M*'s and P*'s occurring. We can then hold M&P fixed at not occurring, which will result in P*'s not occurring (excluding cases where there is another suitable subvenient base P'). Further, this intervention does not require breaking the supervenience relationship. This means it does not require the strange move of treating supervenience as causal. While there may be ways to amend the interventionist framework to evaluate causal diagrams that include both causal and non-causal relationships, it will be more straightforward (and we will make less errors) when variables can be independently set to any combination of values (Woodward 2011, 25). Although Woodward has not described in detail the correct intervention, I take it that my account here would be consistent with what he has in mind. He says (2011: 26): “I suggest that an intervention on a variable X also should be treated as automatically changing (indeed as also an intervention on) the supervenience base $SB(X)$ of X , with $SB(X)$ changing in whatever way is required by the supervenience relation between X and $SB(X)$ ”. The “automatic changing” Woodward describes is built into the intervention I've just described and although it may not be the only interpretation of

Woodward's quote, unlike the interventions describe earlier in this section, there is no room for mistaking the supervenience relationship for a causal one.

The causal arrow from M&P to M*&P* represents only one distinct mechanism. I take this metaphysical fact to be grounded in the rejection of the causal inheritance principle and the rejection of CAM from section 1.2. Given the tight connection between the mental and the physical for NRP, it is unsurprising that the two cause their effects through the same physical mechanism. However, we cannot simply "read" the number of distinct mechanisms off the causal graph. In fact, if we try to adapt Kim's Figure 2 into the interventionist framework, Woodward (2011) suggests that it may in fact be correct to draw only the causal arrow from P to P*, despite the fact that M also causes M* *through the same mechanism*. As these representations can often mislead us in cases that involve both causal and non-causal relationships, I take it as a virtue that my account does not invoke these counterintuitive diagrams.

Perhaps one might suggest that lumping M&P together is illegitimate, not because there is an impossible intervention, but because M and P work through distinct mechanisms. However, this is simply not the case. On our diagram, M&P do not work through distinct mechanisms—a consequence of accepting the principle of impossible interventions. Perhaps the objector would suggest that M and P *should* work through different mechanisms. It is unclear that the NRP is actually committed to the existence of different mechanisms here. Even Kim suggests that mental causation can't be explained by analogy to multiple mechanisms (like in the case of firing squads) because there is no way to make sense of an independent mental mechanism. He says: "it makes no sense to think that there might be an independent, perhaps telekinetic, causal path from the pain to the limb movement" in addition to the neural pathway (1989, 281).

One might suggest that the notion of 'illegitimacy' used throughout this section fails to signify a deep metaphysical problem with these other proposed interventions. Perhaps the fact that many of the

proposed interventions have been labeled as ‘illegitimate’ shows that the interventionist framework simply isn’t very good at representing non-causal dependency relationships. However, the interventionist framework I’ve presented adequately reflects the causal relationships (which it is meant to do) and does not require reworking the causal graphs to include non-causal dependency relationships. Further, it rests not merely on arbitrary rules for causal modeling, but on metaphysical principles. Impossible interventions (e.g., breaking supervenience relationships) should not be allowed because they involve evaluating problematic counterfactuals. By requiring interventions that holding fixed the subvenient base and manipulate the supervenient property, we would be assessing counterfactuals with false (and metaphysically impossible) antecedents. It seems likely that these counterfactuals are simply the wrong ones to evaluate when we’re asking questions about the causal efficacy of properties in a supervenience relationship.

Here is an example that comes from Woodward (2011). Consider a scientist interested in the efficacy of a new drug in treating some condition. Suppose she comes to believe that the macroscopic properties of the drug supervene on the chemical properties of the drug. She would not then attempt to formulate an experiment by which she could hold fixed the chemical composition of the drug while varying the macroscopic properties of the drug—this is impossible. In fact, it doesn’t even seem to be relevant to the types of questions the scientist is interested in, e.g., will this drug be equally efficacious in an older population? The point is that we do not think supervenience (or at least, non-causal dependency) relationships need to be controlled for when testing for causal efficacy. More formally, we should not treat members of a supervenience relationship as being off the directed path that we are evaluating. If we are interested in the difference making of M with respect to P^* , then we should not treat P as an off-path variable that needs to be controlled for. Similarly, if we are evaluating the difference making of P with respect to P^* , then we should not treat M as an off-path variable. It is Kim’s insistence that we con-

control for the causal contribution of M's supervenience base when evaluating M's causal efficacy that misleads much of the discussion about the correct intervention.

In Sections 2.1, 2.2, and the current section, I have argued that our causal notions of sufficiency and difference making come apart in cases of mental causation, but not in cases of overdetermination. Since M and P are both distinct and sufficient causes of P*, each counts as a sufficient cause and the NRP avoids epiphenomenalism. However, M and P are closely related such that we cannot intervene on one without intervening on the other. Interventions that require breaking the supervenience relationship between M and P are illegitimate. Particularly, I addressed the strategy of appealing to multiple realizability in order to replace P with another supervenience base P' during interventions on P. I've argued that it allows the Exclusion Problem to arise again concerning the causal contributions of M and P' and that correctly evaluating counterfactuals might require us to imagine even approximations of the physical state P (which includes P') to be deleted during the intervention. Given the failure of these other interventions, I proposed that M and P only count as a difference-maker, with respect to P* and M*, together. I take this line to be consistent with Woodward's (2011) idea that when evaluating the difference making of a supervening property, we should not control for the supervenience base. Ultimately, I've argued that mental causation involves two sufficient causes but only one difference-maker.

In the next section, I will explain how the Causal Compatibilist's appeal to causal sufficiency and difference making can solve the Exclusion Problem. Given the definition of overdetermination provided above, the NRP avoids the Overdetermination Horn. That is, genuine overdetermination (as in the case of firing squads) involves two sufficient causes and two difference-makers. She also avoids the Epiphenomenalism Horn given that M is still a sufficient cause of P*. Epiphenomenalism involves only one sufficient cause and one difference-maker whereas I've argued that mental causation involves two sufficient causes and one difference-maker.

3 USING INTERVENTIONISM TO SOLVE THE EXCLUSION PROBLEM

In the last section, I argued that the mental and the physical only count as a difference-maker together. I explicated this slogan by using Woodward's formal account of causal interventions—appealing to the requirement to avoid impossible interventions. I have tried to defend the particular type of intervention I proposed and compared it to other previous suggested interventions. In this section, I will argue that the mental and the physical causes are both distinct and sufficient, despite the fact that they only make a difference together. I will also discuss overdetermination in light of the notion of causal difference making. Finally, I will argue that the NRP should reject the Exclusion principle to get out of Kim's Exclusion Argument.

3.1 Distinct and Sufficient Causes

Recall that the Causal Compatibilist's goal is to show that mental causation is not a case of either overdetermination or epiphenomenalism. In this section, I aim to address the objection that the Causal Compatibilist solution described above denies Distinctness.

A defender of Kim's Exclusion Problem might suggest that my solution relies on treating the mental and the physical as non-distinct, at least on the interventionist account of causation. With respect to difference making, it is not just that the mental and the physical work through one mechanism, but rather that the mental and the physical are also treated as a single variable. Perhaps the Causal Compatibilist can appeal to causal sufficiency here, but it seems there are cases where two sufficient causes for an effect are actually the same object or person. For example, Cicero's giving of the Catiline Orations and Tully's giving of the Catiline Orations are both sufficient causes of Catiline leaving Rome. However, Cicero and Tully are actually the same person, so Cicero's and Tully's respective giving of the

Catiline Orations appear not to be distinct sufficient causes.¹⁸ Appealing to the sufficiency of the mental and the physical is not enough to establish that the two are actually distinct.

I argue that the Causal Compatibilist has an adequate reply to this worry. First, although the interventionist framework requires that there is only one causal mechanism that generates M* & P* represented in the causal graph, there is no ontological commitment behind describing a variable as containing both M & P. For example, variables on the interventionist picture can lump together ontologically distinct things (e.g., determinates and determinables).

Consider an example of treating 'red' as a variable. Usually, scarlet is considered to be a determinate of the determinable red. Any particular instance of red will be a type of red, e.g., scarlet. Suppose that a pigeon has been trained to peck when presented with some particular colored stimulus. A lab assistant observes the pigeon peck after the presentation of a scarlet stimulus. Unbeknownst to the lab assistant, the pigeon has been trained to peck after perceiving any shade of red. Consider the following two causal situations:

- (1) Scarlet stimuli cause the pigeon to peck the target.
- (2) Red stimuli cause the pigeon to peck the target.

Clearly, we know that (1) is true. (2) is also true. However, asserting (1) is overspecified in an important way. It might lead the lab assistant to ask whether the pigeon will also peck after the presence of non-scarlet but red (e.g., orange-red) stimuli. The point here is that the scarlet object is also red, and in this particular case, both make a difference to whether or not the pigeon pecks the target. Although red and scarlet are lumped together under the same variable on the interventionist picture, they are still distinct. Red and orange-red can also elicit the stimulus under a different intervention. It is only after multiple interventions on different stimuli that we can distinguish (1) from (2).

¹⁸ I am indebted to Blake Nespica for pressing this objection.

Here the NRP can also appeal to multiple realizability.¹⁹ Many NRPs believe that mental states are multiply realized by a range of neural states. For example, the mental state M may be realized by P or P'. Although I've argued in section 2.3 that an appeal to multiple realizability cannot justify replacing P with P' when evaluating the causal efficacy of M when P does not occur, the Causal Compatibilist can still appeal to multiple realizability in order to establish the distinctness of the mental and the physical. When evaluating the difference making of the mental and the physical, I've argued that the two should be treated as one variable. However, this does not violate Distinctness. Given multiple realizability, M may occur when P does not occur—namely when some other suitable supervenience base P' is present. For the reasons presented in section 2.3, this fact should not be used to evaluate the difference making of M. However, the fact that M can occur without P can be used to show that these states are still numerically (and even modally) distinct.

In this section I have addressed a concern that in making room for the Causal Compatibilist solution by arguing No Overdetermination and Efficacy are consistent, I may have denied Distinctness. I have tried to show that treating the mental and the physical as one variable while assessing the number of difference-makers does not entail that the mental and the physical are non-distinct. In the next section, I will analyze overdetermination more closely. I will then argue that the Causal Compatibilist can distinguish between cases of overdetermination and mental causation by the number of causal pathways.

3.2 Overdetermination and Difference Making

The Causal Compatibilist is more often accused of committing herself to widespread overdetermination than epiphenomenalism. Now that I've described how the Causal Compatibilist can describe the causal structure of mental causation in respect to both causal sufficiency and causal difference making, it will be important for her to distinguish cases of mental causation from overdetermination. Here

¹⁹ The thesis of multiple realizability is controversial, but evaluating it is outside the scope of this paper. For an overview (including some objections), see Bickle (2013).

she will need to evaluate overdetermination in terms of difference making so that she can individuate the causal pathways involved. I argue that cases of genuine overdetermination involve two causal pathways whereas mental causation, as I've argued above, involves only one. Ultimately, this distinction will cast doubt on the Exclusion principle itself, which specifies causal structures according to only sufficient causes.

As it happens, the current literature is divided on the precise definition of overdetermination. Often it is unclear what overdetermination actually involves.²⁰ Sometimes overdetermination is described simply with an analogy to firing squads. When overdetermination is more formally defined, it is merely defined as cases when an effect has more than one sufficient cause.²¹

Firing squads are a paradigmatic example of overdetermination, so they are a natural place to start examining the causal structure of genuine overdetermination. Take two members of a firing squad, Alice and Bill, who both simultaneously shoot their victim, Christine. Assume that both shots hit Christine's heart at the same time and that both Alice's and Bill's acts of firing their guns are individually sufficient to cause Christine's death. Given the accepted definition of overdetermination, it appears that firing squad cases like this are actual instances of overdetermination. Overdetermination is typically depicted in the philosophy of causation literature with the following diagram:




Figure 10. Firing Squad Overdetermination. The nodes each represent some event such as Alice's firing her gun. Each arrow represents causal mechanism from the first node to the second.

²⁰ See Paul 2010, p. 473.

²¹ Menzies (2008) notes that the Exclusion principle is often phrased in terms of causal sufficiency. It may be that the almost exclusive focus on causal sufficiency in discussion of the Exclusion Problem is reminiscent of the focus on laws of nature and the Deductive-Nomological model of explanation (See Woodward 2008).

Kim believes that whatever causal structure best represents the mental causation cases, it is not on par with firing squad cases (1989, 281; 1998, 53). Kim offers a few unsatisfactory explanations for the disanalogy.²² It is important that firing squads and mental causation are relevantly disanalogous because if they were not, mental causation would involve genuine overdetermination. One might hold that the type of overdetermination involved in mental causation and the type involved in firing squads are disanalogous, but that mental causation still involves overdetermination, at least of a less pernicious sort. However, it is less clear that this sort of move is a reply for Kim. Say we replaced the instances of ‘genuine overdetermination’ in the Exclusion Problem with ‘pernicious overdetermination’ and ‘harmless overdetermination.’ Firing squads would fall under the former label while mental causation would fall under the latter. No Overdetermination would simply state that mental causation, as a case of harmless overdetermination, is not analogous to firing squads, as cases of pernicious overdetermination. For Kim to use this distinction in overdetermination to argue that mental causes should be excluded, the Exclusion principle would need to state: if an event has more than one distinct and sufficient cause, then it is (at least) harmlessly overdetermined. It then becomes unclear why being harmlessly overdetermined is a problem; at least the mental causation case does not involve the pernicious sort of overdetermination found in firing squad cases. On the other hand, if the Exclusion principle appealed to pernicious overdetermination, the Causal Compatibilist would simply reject it as false. Mental causation is disanalogous to firing squad cases, but it still involves more than one distinct and sufficient cause. Either way, it seems that this line on overdetermination does not threaten the Causal Compatibilist arguments described in the previous section. Perhaps a Causal Compatibilist can simply refer to mental causation as involving a harmless kind of overdetermination, but for the purposes of this thesis, I will continue to

²² Kim mentions that if firing squads and mental causation were analogous, then overdetermination would be “widespread” (1998, 44-5). The frequency of some potentially problematic event is irrelevant to what is problematic about the event. Kim also considers the argument that if firing squads and mental causation were analogous, then Closure would fail in a near-by possible world where the physical cause does not occur and the effect is caused by the mental cause alone (*ibid.*). Fully responding to this argument is outside the scope of this paper, but see both Crisp and Warfield (2001) and Bennett (2003) for responses.

consider firing squad cases as cases of genuine overdetermination whereas cases of mental causation will not involve overdetermination.

Consider the symmetric firing squad case described above in terms of both causal sufficiency and difference making. Each shooter counts as a sufficient cause of Christine's death because each act of shooting is causally sufficient for the occurrence of her death. What about with respect to difference making? At first, it seems like neither shooting-events are difference-makers to the event of the Christine's death. However, some casual structures require making use of multiple interventions (Woodward 2003, 52). When we hold fixed Bill's shooting as it happens, then Alice's shooting appears not to make a difference to Christine's death. However, when we suppose Bill did not shoot and then vary whether or not Alice shoots, her shooting *does* make a difference to Christine's death. So, the question is not whether Alice's shooting rather than Bill's shooting is the cause of Christine's death.²³ Instead, we should ask whether Alice's shooting rather than Alice's not shooting *when Bill does not shoot* is the cause of Christine's death.²⁴ In this case, Alice's and Bill's shootings will turn out to be both sufficient causes and difference-makers. Note also that the variables representing Alice's and Bill's shooting their guns can be manipulated independently in contrast to cases of variables in a supervenience (or other dependence relationship).

Given the implications of Kim's claims about firing squad cases, Compatibilists should define genuine overdetermination as involving both two sufficient causes and two causal pathways. Often overdetermination, if it is defined at all, is merely described as involving more than one sufficient cause. Instead, I propose a working definition of overdetermination²⁵:

²³ Here I'm assuming that difference-makers need not always make an actual difference. Much of the literature on overdetermination cases has become hung up on this point: Alice's failure to shoot makes no difference with respect to Christine's death if Bill also shoots.

²⁴ I take this to be the same point as the one behind Bennett's (2003, 2007) test for overdetermination.

²⁵ There is likely much more to overdetermination than merely the number of causes and causal pathways. See, e.g., Bernstein (2010).

(Genuine overdetermination): An effect is genuinely overdetermined when there is more than one sufficient, distinct cause *and* more than one causal pathway for that effect.

Even if mental causes do not overdetermine their effects, I will be assuming that genuine overdetermination sometimes actually occurs.²⁶ That is to say, whatever our best theory of event individuation turns out to be, it should be able to account for at least some cases of genuine overdetermination.

Instead of accepting either the Overdetermination Horn or the Epiphenomenalism Horn, I argue that the Compatibilist should deny Kim's dichotomy. While I agree with Kim that the NRP cannot easily reject that the mental and the physical are distinct or that every physical event has a sufficient physical cause, I think Exclusion stands on shakier ground. Recall the Exclusion principle: If an event at *t* has more than one distinct, sufficient cause at *t*, then it is genuinely overdetermined. Kim states that Exclusion "is virtually an analytic truth without much content" and he provides no further discussion (2005, 51).²⁷

3.3 Denying the Exclusion Principle

In the last few sections, I suggested that the causal structure of overdetermination and mental causation should be differentiated in terms of the number of causal pathways. While mental causation and overdetermination both involve two sufficient causes, the former only involves one causal pathway while the latter involves two. However, the Exclusion principle as Kim formulates it cannot distinguish cases of mental causation and overdetermination. Since the Exclusion principle makes use of only causal sufficiency, it lumps mental causation in with genuine overdetermination. Above I have argued that the two come apart, and for that reason, in this section, I will argue that the Exclusion principle is false.

²⁶ I will be explicitly rejecting views that hold overdetermination does not occur because events are fragile (e.g., temporally, modally) and overdetermination is really joint causation underdescribed (e.g., Collectivism). See Kim (1993b) and Bunzl (1979) respectively. For discussion, see Schaffer (2003). This is not to say that *some* cases of seeming overdetermination might be best described by Collectivist theories, for example, presidential voting scenarios.

²⁷ Barry Loewer agrees that this assumption cannot rightly be called analytic: 2007, p. 13.

Instead, I argue the Causal Compatibilist should deny the Exclusion principle. As Kim formulates the principle, it is simply false. As I've argued, mental causation is a case where an effect has more than one distinct, sufficient cause, but it does not involve overdetermination. Perhaps Kim could attempt to revise the Exclusion principle in light of our two causal concepts—sufficiency and difference making. Suppose Kim defined the revised Exclusion principle as follows: If an event has more than one distinct and sufficient cause, and it has more than one causal pathway, then it is genuinely overdetermined. This version of the Exclusion principle will not be strong enough to exclude the mental cause, since, as I've argued, mental causation involves more than one distinct and sufficient cause but only one causal pathway.

Perhaps Kim would respond that the Exclusion principle does not need to be revised. In fact, he might suggest that the NRP has given him no reasons accept that cases of mental causation, epiphenomenalism, and overdetermination should be individuated with respect to causal difference making at all. Instead, he might reply that causal sufficiency is the relevant causal concept for the Exclusion Problem and possibly for all of science.²⁸

First, such a response is not motivated by the causal contribution Kim requires of the mental, if it is to be efficacious. That is, Kim does not hold that the NRP can simply show that the mental is sufficient for its cause in order to get out of the Exclusion Problem. Kim appeals to notions of causal work: "Given that P has a physical cause P*, what causal work is left for M to contribute?" (1998, 37). He also suggests that the "causal contribution" of M is "totally mysterious" (2005b, 48). Kim's conclusion is that since M does no additional causal work to bring about P* (and accepting that mental causation does not involve overdetermination), M does not actually cause P. As Kim sometimes says, the mental is dispensable (1998, 45). However, Kim also states that the mental needs to make a "causal difference" (1998,

²⁸ Thanks to Paul Henne for bringing my attention to this objection.

31; 2005b, 10). This last metaphor seems to indicate that the presence of the mental should matter to the occurrence of the effect—that the mental should be a difference-maker.

Although it may be unlikely that Kim would accept all the formalism found in Woodward's interventionist theory of causation (and especially since Kim openly admits to holding a production view of causation), it might be the case that Kim is committed to at least some intuitive notion of causal difference making. Recall the discussion from Section 1.2. Kim holds that the mental needs to perform extra causal work above and beyond the physical. If Kim is referring only to our notion of causal sufficiency, it would be irrelevant whether or not the mental performed any extra causal work. All we would need to know is that the mental is sufficient for the effect and that the physical is sufficient for the effect. There would be no impulse to "control for" the physical while manipulating the mental in order to see whether or not the effect occurs. Further, if Kim would object to this appeal to causal difference making, it would seem that his theory of causation leaves out very common and intuitive notion of causation. In particular, causal difference making is important to scientific methodology (Woodward 2003: 10-12), presumably the type of causation Kim would like to account for on his view.

Instead, it seems that even Kim appeals to our causal concept of difference making to suggest that the mental will need to make a difference to the occurrence of its effects. The objection described above as rejecting any reference to causal difference making in a solution to the Exclusion Problem is ill motivated by Kim's own views.

4 POTENTIAL OBJECTIONS

In section 2, I introduced two notions of causation: difference making and sufficiency. I appealed to Woodward's interventionist framework in order to individuate the number of causal difference makers in a given causal graph. Then I argued that by treating the mental and the physical as the same variable in a causal graph, they both make a difference to their effects together. My justification for this claim is that other possible interventions break the supervenience relationship or require impossible interventions. This solution is consistent with Woodward's (2011) claim that when evaluating whether a supervening property makes a difference to its effects, then we should *not* control for (or hold fixed) the supervenience base. In section 3, I have just argued that if the solution I propose in section 2 is correct, then the NRP can get out of Kim's Exclusion Problem by denying the Exclusion principle. I have also argued that my solution does not deny the distinctness of the mental and the physical. Here I think the NRP should appeal to multiple realizability in order to establish Distinctness. Then I argued that the mental does not overdetermine its effects. I proposed a new definition of overdetermination based on our notions of causal sufficiency and difference making. Ultimately, I have argued that Kim's Exclusion principle is false because it fails to distinguish between cases of mental causation and overdetermination. So, the NRP has no reason to accept it and she has a way out of the Exclusion Problem.

In this section, I will consider a few possible objections to the solution I've just presented. I will conclude that my solution does not violate Efficacy or Closure. The Efficacy objection holds that causal sufficiency and causal difference making never produce different answers to the individuation of causes, but as I've argued above, they do come apart in cases of mental causation. The Closure objection rests on a confusion between causal sufficiency and causal difference making.

4.1 Efficacy Objection

Bernstein (2010) argues that Compatibilists cannot rely on solutions like mine because Efficacy entails that mental causes overdetermine their effects. First, Bernstein argues that causal Compatibilists rely on an intuitive notion of overdetermination, where cases that involve a particular type of distinctness between the causes do not count as overdetermined, rather than an ontological notion. However, according to Bernstein: “It is this antecedent ontological commitment—a commitment to an extra set of causes—that generates overdetermination” (2010, 28). She then examines the relationship between the ontological commitment that generates overdetermination and Efficacy. Recall that Efficacy states: The instantiations of mental properties can and do cause other physical and mental properties to be instantiated. Efficacy, she argues, requires that mental causes have both potential causal processes and follow-throughs (or actual, completed causal processes) connecting them to their effects.²⁹ Her objection ultimately relies on explicating Efficacy in terms of requiring unique causal processes (both potential and actual) from mental causes to their effects.³⁰

In order to deny that mental causation involves overdetermination, the causal Compatibilist needs to deny the ontological commitment to two causes, but unfortunately, according Bernstein, neither way is consistent with Efficacy. The Compatibilist can either deny that the mental has causal potential, or she can deny that the mental has causal follow-through. However, causal potential seems to be required by Efficacy. On the other hand, mental causes will always fail to be efficacious (given Closure). Unfortunately, on Bernstein’s argument, the mental has the potential to cause effects, but it simply never does.

²⁹ By ‘causal process’, Bernstein does not mean to appeal to the philosophically loaded version found in other theories (e.g., Salmon 1984). Instead, she means a type of unspecified causal connection between the cause and effect. The difference between a potential causal process and a follow-through can be shown in a case of asymmetric preemption. Suppose Suzy throws her ball faster than Billy throws his ball, and so Suzy’s ball shatters the window while Billy’s ball only flies through the air where the window once stood. Suzy’s throw is a cause with follow-through while Billy’s throw is only a potential causal process (*ibid.*, 29).

³⁰ Note that Efficacy requires that mental properties cause their effects *qua* mental rather than inheriting their causal efficacy from their subvenient base. See Horgan (1989).

Bernstein likely rejects the Compatibilist's causal pathway solution because she assumes that the numbers of sufficient causes and of "causal processes" are always equal once individuated. Bernstein says: "There are two causal processes *in virtue of* there being two causes: accept the causes and the complete causal processes come attached" (*ibid.*, 47). However, as I argued in Section 2.3, this is not always the case. Mental causation might be an instance in which the number of causes does *not* determine the number of causal pathways. Causal Compatibilists need not accept this relationship between causes and causal pathways.

4.2 Closure Objection

Another objection might be that I've violated Closure in my description of Figure 2. It might appear that the causal mechanism runs from M to M*; the objector argues that, in my diagram, P inherits this causal mechanism. Yet, if P inherits its causal efficacy from M, then it appears that the NRP has denied the *sufficiency* of the physical.

I agree that were this NRP solution to violate Closure, this would be a bad outcome. However, such an objection confuses causal pathways with sufficiency. Recall that I've defined 'causal pathway' in terms of difference making. Nothing I've said suggests that P is not sufficient for P*. Likewise, there is no reason to think that having more than one sufficient cause is problematic. Consider the following causal chain: $A \rightarrow B \rightarrow C$. We might say that A is sufficient for C and then that B is sufficient for C, but that does not entail that A's (or B's) sufficiency is now suspect.

Conceding the point about sufficiency, the objector (as well as Kim) might press the issue of inheriting causal efficacy. The problem is that P's purported inheritance of causal efficacy from M threatens Closure. However, the talk of one state inheriting causal efficacy from the other (where the inheritor is either P or M) is misleading. The causal mechanism between the two sets of events is equally M's as it is P's. Since M and P only count as difference-makers when both are present, we might say that the

causal mechanism belongs to both.³¹ As Shapiro (2010) says treating M as a free rider “involves confusing M’s causing P* with M’s making a causal contribution to P* that is additional to the contribution that P makes” (600). That is to say, Kim believes the mental needs to make an additional “causal contribution” to the effect in order to be efficacious (2005, 47). This assumption is what led Kim to suggest that Exclusion applies immediately after the NRP rejects mental causation as a case of genuine overdetermination. The Compatibilist need not agree with Kim that Exclusion applies here. Instead, she can reject the assumption that the mental has to add anything *new* to the causal situation (Wilson 2010b, Paul 2007, Shapiro 2010).

If I have successfully defended the distinction between sufficient causes and causal pathways, the NRP has a way out of the Exclusion Problem. She can provide a more intuitive definition of overdetermination and then she can differentiate sufficient causes and causal pathways in terms of causal sufficiency and causal difference making. The NRP can defend a causal Compatibilist solution in which she accepts No Overdetermination and Efficacy but rejects Exclusion. Her distinction between causes and causal pathways shows why Kim’s formulation of the Exclusion principle is incomplete—it only makes use of our causal concept of causal sufficiency.

4.3 Conclusions

I have provided the NRP with a response to Kim’s dilemma. Instead of accepting either the Epiphenomenalism Horn or the Overdetermination Horn, I suggest the NRP reject Kim’s false dichotomy. The NRP should reject Kim’s version of Exclusion because it is inadequately formulated. I also argued

³¹ This is what I take to be the problem with Alyssa Ney’s (2009) argument that Foundationalists about causation, who think that all of our difference-making talk is ultimately grounded in facts about the type of causation discovered “in the real world” by way of scientific investigation, ought to land on Kim’s side of the argument. Ney thinks that science indicates that physical causes are productive, and while this is contentious, I think a Foundationalist about causation can even admit that our difference-making talk is grounded in a production view of causation. However, Kim’s conclusion does not follow directly. A causal Compatibilist might suggest the problem with Ney’s argument is that she, like Kim and Bernstein, assumes that the mental needs to causally contribute *something extra*, as opposed to merely causally contributing.

that the NRP should specify her own definition of overdetermination. Both should require individuation of causal structure according to the number of sufficient causes *and* the number of causal pathways. The NRP should argue that mental causes are neither excluded as causes nor overdetermine their effects; mental causation involves two causes but only one causal pathway. I considered some objections, but ultimately one objection rested on the intuition that causal sufficiency and difference making do not come apart and another objection rested on a confusion about the distinction between causal sufficiency and causal difference making. Ultimately, I concluded that NRP have little reason to be convinced by Kim's Exclusion Argument, given his incomplete and unconvincing version of the Exclusion principle.

REFERENCES

- Barnes, Elizabeth. (Forthcoming). "Emergence and Fundamentality." *Mind*.
- Baumgartner, Michael. (2013). "Rendering Interventionism and Non-Reductive Physicalism Compatible." *Dialectica*. 67(1): 1-27.
- Bennett, Karen. (2007). "Exclusion Again." J. Kallestrup and J. Hohwy (Eds.). In *Being Reduced: New Essays on Reduction, Explanation, Causation*. Oxford University Press, 280-305.
- Bennett, Karen. (2003). "Why the Exclusion Problem Seem Intractable, and How, Just Maybe, to Tract It." *Noûs*. 37(3): 471-497.
- Bernstein, Sara. (2010). "Essays on Overdetermination." Dissertation.
- Bickle, John. (2013). "Multiple Realizability." *The Stanford Encyclopedia of Philosophy*. Edward Zalta (Ed.). <<http://plato.stanford.edu/archives/spr2013/entries/multiple-realizability/>>.
- Bunzl, Martin. (1979). "Causal Overdetermination." *The Journal of Philosophy*. 76(3): 134-150.
- Crisp, Thomas, and Ted Warfield. (2001). "Kim's Master Argument." *Nous*. 35(2): 304-316.
- Dowe, Phil. (2000). *Physical Causation*. Cambridge University Press.
- Frisch, Mathias. (2009). "'The Most Sacred Tenet'? Causal Reasoning in Physics." *British Journal for the Philosophy of Science*. 60(3): 459-474.
- Hall, Ned. (2004). "Two concepts of causation." J. Collins, N. Hall, and L.A. Paul (Eds.). *Causation and Counterfactuals*. Cambridge, Mass.: MIT Press), 225-276.
- Horgan, Terence. (1989). "Mental Causation." *Philosophical Perspectives 3: Philosophy of Mind and Action Theory*. 47-76.
- Horgan, Terence. (2001). "Causal Compatibilism and the Exclusion Problem." *Theoria*. 40, 95-116.
- Kim, Jaegwon. (1989). "The Myth of Nonreductive Materialism." In *Supervenience and Mind*. Cambridge: Cambridge University Press, 265-284.
- Kim, Jaegwon. (1993a). "The Nonreductivist's Troubles with Mental Causation." In *Supervenience and*

- Mind*. Cambridge: Cambridge University Press, 336-357.
- Kim, Jaegwon. (1993b). "Events as Property Exemplifications." In *Supervenience and Mind*. Cambridge: Cambridge University Press, 33-52.
- Kim, Jaegwon. (1993c). "Multiple Realization and the Metaphysics of Reduction." In *Supervenience and Mind*. Cambridge: Cambridge University Press, 309-335.
- Kim, Jaegwon. (1998). *Mind in a Physical World*. Cambridge, MA: Bradford.
- Kim, Jaegwon. (2005). *Physicalism, or Something Near Enough*. Princeton, NJ: Princeton University Press.
- Kim, Jaegwon. (2008). "Reduction and Reductive Explanation: Is One Possible Without the Other?" In *Being Reduced: New Essays on Reduction, Explanation, and Causation*. Hohwy, J. and Kallestrup, J. (Eds.). Oxford: Oxford University Press.
- Kim, Jaegwon. (2010). "Thoughts on Sydney Shoemaker's Physical Realization." *Philosophical Studies*. 148(1): 101-112.
- Lewis, D. (1973). *Counterfactuals*. Oxford: Blackwell.
- Lewis, D. (2000). Causation as Influence. *The Journal of Philosophy*. 97: 182-197.
- Loewer, Barry. (2007). Mental Causation or Something Near Enough. In Brian P. McLaughlin & Jonathan D. Cohen (eds.), *Contemporary Debates in Philosophy of Mind*. Blackwell, 243-264.
- Menzies, Peter. (2008). "The Exclusion Problem, the Determination Relation, and Contrastive Causation." In *Being Reduced: New Essays on Reduction, Explanation, and Causation*. Hohwy, J. and Kallestrup, J. (Eds.). Oxford: Oxford University Press.
- McLaughlin, Brian, and Karen Bennett. (2011). "Supervenience." The Stanford Encyclopedia of Philosophy. Edward Zalta (Ed).
<<http://plato.stanford.edu/archives/win2011/entries/supervenience/>>.
- Musallam, S., Corneil, B., Greger, B., Scherberger, H., and Andersen, R. (2004). "Cognitive Control Signals for Neural Prosthetics." *Science*. 305(5681): 258-262.

- Ney, Alyssa. (Forthcoming). "The Causal Contribution of Mental Events." In *The Mental and the Physical: New Perspectives on Type Identity*. S. Gozzano and C. Hill (Eds.) Cambridge University Press.
- Ney, Alyssa. (2010). "Convergence on the Problem of Mental Causation". *Philosophical Issues*. Vol. 20.
- Ney, Alyssa. (2009). "Physical Causation and Difference making". *British Journal for the Philosophy of Science*. Vol. 60.
- Norton, John. (2007). "Do the Causal Principles of Modern Physics Contradict Causal Anti-Fundamentalism?" In P. Machamer and G. Wolters (Eds.). *Thinking about Causes: From Greek Philosophy to Modern Physics*. Pittsburgh: University of Pittsburgh Press, 222-34.
- Norton, John. (2009). "Is There an Independent Principle of Causality in Physics?" *British Journal for the Philosophy of Science*. 60: 475-486.
- Pacherie, Elisabeth. (2008). "The Phenomenology of Action: A Conceptual Framework." *Cognition*. 107: 179-217.
- Paul, L.A. (2010). "New Role for Experimental Work in Metaphysics." *European Review on Philosophy and Psychology*. Special Issue. Joshua Knobe, Tania Lombrozo and Edouard Machery (Eds.). 1(3):461-476.
- Paul, L.A. (2007). "Constitutive Overdetermination." In *Topics in Contemporary Philosophy vol. 4: Causation and Explanation*, Joseph Keim Campbell and Michael O'Rourke (Eds). Cambridge: MIT Press.
- Paul, L.A. (2000). "Aspect Causation." *The Journal of Philosophy*. 97(4): 235-256.
- Pearl, Judea. (2000). *Causality*. New York: Cambridge University Press.
- Salmon, Wesley. (1984). *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton University Press.
- Schaffer, Jonathan. (2003). "Overdetermining Causes." *Philosophical Studies* 114: 23-45.
- Schaffer, Jonathan. (2009). "On What Grounds What." In David Manley, David Chalmers, and Ryan Wasserman (Eds.). *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford University

Press.

Shapiro, Larry. (2010). "Lessons from Causal Exclusion." *Philosophy and Phenomenological Research*. 81: 594-604.

Shapiro, Larry and Sober, Elliott. (Forthcoming). "Epiphenomenalism – the Do's and the Don'ts". In *Studies in Causality: Historical and Contemporary*. G. Wolters and P. Machamer (Eds.). University of Pittsburgh Press.

Shoemaker, Sydney. (2001). "Realization and Mental Causation." In Gillett, Carl and Loewer, Barry (Eds.). *Physicalism and its Discontents*. Cambridge University Press.

Skyrms, Brian. (1984). "EPR: Lessons for Metaphysics." In P. French, T. Uehling, and H. Wettstein. *Midwest Studies in Philosophy IX*. Minneapolis: University of Minnesota Press, 245-255.

Stoljar, Daniel. (2008). "Distinctions in Distinction." In J. Kallestrup and J. Hohwy (Eds.). In *Being Reduced: New Essays on Reduction, Explanation, Causation*. Oxford University Press, 263-279.

Strevens, Michael. (2008). "Comments on Woodward, *Making Things Happen*." *Philosophy and Phenomenological Research*. 77(1): 171-192.

Wilson, Jessica. (2010a). "Non-reductive Physicalism and Degrees of Freedom." *British Journal for the Philosophy of Science*. 61: 279-311.

Wilson, Jessica. (2010b). "Non-reductive Realization and the Powers-based Subset Strategy." *Monist*. 94: 121-154.

Wilson, Jessica, and Bernstein, Sara. (Draft). "Mental causation and Free Will."

Wilson, Robert. (2001). "Two Views of Realization." *Philosophical Studies*. 104: 1-31.

Wolff, P. (2008). "Dynamics and the perception of causal events." In T. Shipley & J. Zacks (Eds.), *Understanding events: How humans see, represent, and act on events*. Oxford University Press, 555-587.

Woodward, James. (2003). *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford Uni-

versity Press.

Woodward, James. (2008). "Mental Causation and Neural Mechanisms." In J. Kallestrup and J. Hohwy (Eds.). In *Being Reduced: New Essays on Reduction, Explanation, Causation*. Oxford University Press.

Woodward, James. (2010). "Causation in Biology: Stability, Specificity, and the Choice of Levels of Explanation." *Biology and Philosophy*. 25: 287-318.

Woodward, James. (2011). "Interventionism and Causal Exclusion."
<<http://philsci-archive.pitt.edu/id/eprint/8651>>.

Yablo, S. (1992). "Mental Causation." *Philosophical Review*. 101(2): 245-280.