

# Exclusion Excluded\*

Brad Weslake, NYU SHANGHAI

[brad.weslake@nyu.edu](mailto:brad.weslake@nyu.edu)

## *Abstract*

The non-reductive physicalist would like to believe that mental properties are not identical to physical properties; that there are complete causal explanations of all events in terms of physical properties; and that there are sometimes explanations of events in terms of mental properties. However, some have argued that these claims cannot all be true, since they are collectively inconsistent with a principle of causal exclusion. In this paper I argue that the best formulation of the interventionist theory of causation entails the falsity of the exclusion principle. However, I argue that it does so at the cost of revealing a weakness in the interventionist theory itself.

---

\* Earlier versions of this paper were presented to a Mellon Mental Causation Workshop at Syracuse University in 2007, to the CUNY Cognitive Science Symposium in 2008, and to a workshop at Harvard University in 2019. I am grateful to the audiences on those occasions, and especially to Karen Bennett, Tony Dardis, David Rosenthal, Samuel Lee and Ned Hall. I also thank Arif Ahmed, Ben Blumson, David Braddon-Mitchell, Michael Brent, John Campbell, Brandon Carey, Shamik Dasgupta, Sommer Hodson, Adam Kay, Doug Kutach, David Macarthur, Kevin McCain, Peter Menzies, Huw Price, and Nandi Theunissen for discussion or feedback, and Michael Baumgartner and Michael Rescorla for very helpful correspondence, on earlier versions. More recently, I thank Michael Tye for sending me his work on this topic, and Alastair Wilson, Katie Robertson, Henry Taylor, Mike Hicks, Ben Curry and two referees for this volume for their advice on preparing the final version.

## Contents

<b>1</b>	<b>Introduction</b>	<b>2</b>
<b>2</b>	<b>Interventionism Introduced</b>	<b>4</b>
<b>3</b>	<b>Exclusion Reformulated</b>	<b>12</b>
3.1	Non-Reductionism . . . . .	12
3.2	Completeness . . . . .	14
3.3	Mental Causation . . . . .	20
3.4	Exclusion . . . . .	22
3.5	The Interventionist Exclusion Problem . . . . .	24
<b>4</b>	<b>Compatibilism Examined</b>	<b>25</b>
4.1	Subvenience Sufficiency . . . . .	25
4.2	Internalism . . . . .	27
4.3	Externalism . . . . .	36
4.4	A Weakness in Interventionism . . . . .	39
<b>5</b>	<b>Conclusion</b>	<b>41</b>

## 1 Introduction

I take the exclusion problem to be the problem of providing a principled reason to reject at least one of the following inconsistent claims<sup>1</sup>:

---

<sup>1</sup>While this initial formulation of the problem involves the causal explanation of events in terms of properties, everything I say below could be reformulated depending on your preferred view of the causal relata. Sometimes **Completeness** is weakened, so that it does not presuppose that all events have complete causal explanations. I employ the stronger principle for simplicity, as it will not make any difference to my argument. I assume throughout that “explanation” is a factive term, and that in a causal explanation all explanans properties are causes of the explanandum. If you prefer not to formulate the problem as involving explanation at all, but rather as involving complete or sufficient causes, be patient: explanation will not appear in the final formulation of the problem I reach in Section 3.5.

- **Non-Reductionism.** Mental properties are distinct from, though metaphysically necessitated by, physical properties.
- **Completeness.** Every event has a complete causal explanation in terms of physical properties.
- **Mental Causation.** There exist causal explanations of events in terms of mental properties.
- **Exclusion.** If an event has a complete causal explanation in terms of one set of properties then it has no causal explanation in terms of any other properties.

In this paper, I examine the prospects for a principled rejection of **Exclusion**. Following Horgan (1997, p. 166) and Bennett (2003, p. 473, 2008, p. 283), I will refer to this position as *compatibilism*<sup>2</sup>. I will refer to the conjunction of **Non-Reductionism**, **Completeness** and **Mental Causation** as *non-reductive physicalism*.

Compatibilism is a popular position<sup>3</sup>. However, it has frequently been defended in the absence of an independently justified general framework for thinking about causation and causal explanation. That began to change after the development of a justly influential theory of causation and explanation by James Woodward (2003), which has come to be referred to as *interventionism*. The development of interventionism generated a robust debate concerning whether an interventionist is entitled to reject **Exclusion**, and it is this question I explore in what follows<sup>4</sup>. My central claim is that there is a significant blindspot in the existing discussion,

---

<sup>2</sup>Bennett restricts the term to those who say that mental causation is possible without causal overdetermination. I use the term in Horgan's more general sense.

<sup>3</sup>Bennett (2003) cites Goldman (1969), Blackburn (1991), Pereboom and Kornblith (1991), Yablo (1992), Burge (1993), Mellor (1995, pp. 103–104), Horgan (1997), Noordhof (1997) and Yablo (1997), to take just a few of the more prominent adherents.

<sup>4</sup>For arguments broadly sympathetic to interventionism on this score, see Shapiro and Sober (2007), Shapiro (2010), Raatikainen (2010), Polger, Shapiro, and Stern (2018), Woodward (2008a, 2015a, 2017) and Stern and Eva (2023). For arguments broadly critical, see Baumgartner (2009, 2010, 2013) (see also Baumgartner 2018), Hoffman-Kolss (2014) and Gebharder (2017a).

concerning the nature of the relationship between physical and mental properties. Attention to this blindspot reveals that while the best formulation of the interventionist theory of causation entails the falsity of the exclusion principle, it does so at the cost of revealing a weakness in the interventionist theory itself.

The structure of the paper is as follows. In Section 2 I introduce interventionism. In Section 3 I consider how to formulate the exclusion problem in interventionist terms, addressing each component of the problem in turn. In Section 4 I turn to arguments for **Exclusion**. In Section 4.1 I introduce a principle, *subvenience sufficiency*, concerning the relationship between physical and mental properties. The existing discussion has universally accepted the principle, thereby accepting a position I call *internalism*. I consider exclusion arguments from that standpoint in Section 4.2. In Section 4.3 I formulate an exclusion argument under the assumption that subvenience sufficiency is false, a position I call *externalism*. I argue that while interventionism has a response to the argument, it is one that reveals a limitation in the interventionist theory itself. I conclude in Section 5.

## 2 Interventionism Introduced

Central to the interventionist framework is the notion of a causal model<sup>5</sup>. A causal model is a representational device for encoding counterfactual relationships between variables. Counterfactual relationships are represented by equations which specify the way in which the value of a single variable on the left hand side would change as a function of changes to the values of the variables on the right hand side. More formally then, a causal model is an ordered pair  $\langle \mathcal{V}, \mathcal{E} \rangle$ , where  $\mathcal{V}$  is a set of variables and  $\mathcal{E}$  a set of equations, and every variable appears on the left hand side of exactly one equation.

---

<sup>5</sup>Here I present just enough to set up the discussion that follows. For more extended introductions to causal models, see Hitchcock (2009, 2023).

For example, a model  $\mathcal{M}_1$  might contain equations representing that variable  $Y$  depends on variables  $X_2$  and  $X_3$ , that variable  $X_2$  depends on variable  $X_1$ , and that variables  $X_1$  and  $X_3$  took values 1 and 0 respectively:

$$Y := X_2 \vee X_3$$

$$X_2 := X_1$$

$$X_1 := 1$$

$$X_3 := 0$$

Here “ $\vee$ ” should be interpreted as a function returning 1 if either side is 1 and 0 otherwise. Equations such as the last two, that simply assign a specific actual value to a variable, are *exogenous*. Equations such as the first two, that assign values as a function of other variables, are *endogenous*. I will assume that the equations are all deterministic, in which case the equations for a model entail the actual values of all variables in the model. In  $\mathcal{M}_1$  for example, the equations entail that  $X_2$  and  $Y$  both took value 1.

Variables in a causal model must represent entities capable of being changed by interventions, but the framework is otherwise consistent with a range of different metaphysical views concerning the nature of the causal relata. For simplicity, I will sometimes say that variable values represent properties and sometimes say that they represent events. All of this could be translated into whatever view of the causal relata is correct<sup>6</sup>. I will refer to a possible assignment of values to a set of variables as a *state* of that set, and will talk freely of actual and possible variable values, changes to variable values, states and changes of state of models. I will also talk about causal relations obtaining between variables and values of variables. This sort of talk should be interpreted throughout as reflecting corresponding actual or possible changes in, and causal

---

<sup>6</sup>For the complexities that this simplification evades, see Schaffer (2016, §1) and Gallow (2022, §1).

relations obtaining between, what is represented by the model. I will assume throughout that a causal model must be veridical, in the sense that every counterfactual relationship specified by the model is true.

The counterfactuals represented by causal models concern *interventions*. An intervention is an exogenous change to the value of a variable in a model, in the sense that the values of the other variables in the model are not themselves causes or effects of the change, unless they are effects of the variable intervened on. Moreover, it is required that interventions be *surgical*, in the sense that the usual causes of the variable in question are suspended, so that the value of the variable depends only on the intervention. I will consider the nature of interventions in more detail in Section 4.2.

In the literature on causation it has been common to distinguish between type-causal relations and token-causal relations. An analogous distinction can be made between causal relations between variables, and causal relations between variable values. While the terminology is slightly misleading, I will follow Woodward (2003) and refer to causal relations between variables as *type-level* causal relations, and between values of variables as *token-level* causal relations<sup>7</sup>.

In the remainder of this section I introduce the definitions in the interventionist framework that will be important for what follows<sup>8</sup>.

First we need the type-level notion of a *direct cause* (Woodward 2003, p. 55):

- **DC.** *X* is a *direct cause* of *Y* in model  $\mathcal{M}$  iff there is a possible intervention on *X* that would change *Y* when all other variables in

---

<sup>7</sup>For a discussion of the relationship between type-causal relations, token-causal relations, and causal relations between variables, see Hausman (2005).

<sup>8</sup>While I provide references to Woodward throughout, the precise formulations I give are sometimes simplified or expanded, and sometimes make use of definitions introduced in this paper. One important simplification is that I am setting aside the generalisation to the case of probabilistic causation, on which see Fenton-Glynn (2021).

$\mathcal{M}$  besides  $X$  and  $Y$  are held fixed at some combination of values by interventions<sup>9</sup>.

It is a necessary and sufficient condition for  $X$  to be a direct cause of  $Y$  in  $\mathcal{M}$  that  $X$  appear on the right hand side of the equation for  $Y$  in  $\mathcal{M}$ . So for example in model  $\mathcal{M}_1$ ,  $X_1$  is a direct cause of  $X_2$ , and  $X_2$  and  $X_3$  are direct causes of  $Y$ .

Second we need the type-level notion of a *directed path* (*ibid*, 42). This can be defined in terms of the properties of graphs associated with causal models. A *directed graph* for  $\mathcal{M}$  is an ordered pair  $\langle V, E \rangle$  where  $V$  is a set of vertices that correspond to the set of variables in  $\mathcal{M}$  and  $E$  a set of *directed edges* connecting these vertices, where there is a directed edge from vertex  $X$  to vertex  $Y$  iff  $X$  directly causes  $Y$  in  $\mathcal{M}$ . The definition is then:

- **P.** A sequence of variables  $\{V_1 \dots V_n\}$  is a *directed path* from  $V_1$  to  $V_n$  in  $\mathcal{M}$  iff for all  $i(1 \leq i < n)$  there is a directed edge from  $V_i$  to  $V_{i+1}$  in the directed graph for  $\mathcal{M}$ .

From here on, *path* should be read as equivalent to *directed path*. A path is simply a sequence of direct causes, but the graph-theoretic definition is useful because paths in a model can be easily discerned by constructing a diagram with the same structure as the associated directed graph. When presenting diagrams of this sort, I will follow the usual convention of using circles to represent vertices (variables) and arrows to represent directed edges (direct causes). So for example, by inspecting the diagram for  $\mathcal{M}_1$  in Figure 1, it is easy to see that  $X_1$  is a direct cause of  $X_2$ , that  $X_2$  and  $X_3$  are direct causes of  $Y$ , and that there is a path from  $X_1$  to  $Y$ , from  $X_2$  to  $Y$ , and from  $X_3$  to  $Y$ .

I will provide diagrams of this sort when they are helpful. However, it is important to keep in mind that not all of the information relevant to

---

<sup>9</sup>In interpreting the condition in this way I agree with Baumgartner (2009). Woodward (2015a) confirms that this was his intended interpretation.




Figure 1: Diagram for  $\mathcal{M}_1$ .

causation in the interventionist framework can be read off these diagrams. In particular, to know whether the next two definitions are satisfied, you need to know the particular equations that relate the variables.

Third we need the type-level notion of a *contributing cause* (*ibid*, 59):

- **CC.**  $X$  is a *contributing cause* of  $Y$  in model  $\mathcal{M}$  iff for some path  $P$  from  $X$  to  $Y$  in  $\mathcal{M}$ , there is an intervention on  $X$  that will change  $Y$  when all variables in  $\mathcal{M}$  not on  $P$  are held fixed at some combination of values by interventions.

In model  $\mathcal{M}_1$  for example,  $X_1$ ,  $X_2$  and  $X_3$  are all contributing causes of  $Y$ . When  $X_3 = 0$ , an intervention setting  $X_1$  from 0 to 1 would result in  $Y$  changing from 0 to 1. Likewise for  $X_2$ . And when  $X_1 = 0$  and  $X_2 = 0$ , an intervention setting  $X_3$  from 0 to 1 would result in  $Y$  changing from 0 to 1.

Finally, we need the token-level notion of an *actual cause*. The precise way to define actual causation in the interventionist framework remains



a matter of lively debate. However, as I show in Weslake ([unpublished](#)), many of the proposed definitions can be formulated as instances of the following schema:

- **AC.**  $X = x$  is an *actual cause* of  $Y = y$  relative to model  $\mathcal{M}$  iff:
  - **ACT.** The actual value of  $X = x$  and the actual value of  $Y = y$ .
  - **PATH.** There exists a path  $P$  from  $X$  to  $Y$  in  $\mathcal{M}$  for which an intervention on  $X$  would change the value of  $Y$ , when all variables  $V_1 \dots V_n$  in  $\mathcal{M}$  that are not on  $P$  are held fixed at some combination of values satisfying *<conditions specifying permissible values  $v_1 \dots v_n$  for  $V_1 \dots V_n$ >*.

The conditions specifying permissible values can be thought of as specifying the set of possible values of the off-path variables relative to which an intervention constitutes a test for actual causation along that path. All definitions of this form in the literature agree that *one* such permissible set is that in which all off-path variables have their actual values. So they all agree that a sufficient condition for  $X = x$  to be an actual cause of  $Y = y$  is for there to be a path from  $X$  to  $Y$  such that holding all off-path variables fixed at their actual values, there is an intervention setting  $X = x'$  where  $x \neq x'$  that would result in  $Y = y'$  where  $y \neq y'$ . In effect, that is, these theories agree that counterfactual dependence (of this sort) is sufficient for causation. Fortunately, for the purposes of the arguments I make below the differences between the various theories of actual causation on offer do not make any difference. So I will work with the following definition of actual causation, which also takes counterfactual dependence (of this sort) to be necessary for causation<sup>10</sup>:

- **AC<sub>A</sub>.**  $X = x$  is an *actual cause* of  $Y = y$  relative to model  $\mathcal{M}$  iff:
  - **ACT.** The actual value of  $X = x$  and the actual value of  $Y = y$ .

---

<sup>10</sup>The definition is equivalent (modulo some irrelevant differences) to Woodward's (AC) (2003, p. 77), and the definition of causation defined in terms of "Act" in Hitchcock (2001, pp. 286–287). In Weslake ([unpublished](#)) I argue against this and all other theories that fit the schema, but the arguments that follow also work for the theory I prefer.

- **PATH<sub>A</sub>**. There exists a path  $P$  from  $X$  to  $Y$  in  $\mathcal{M}$  for which an intervention on  $X$  would change the value of  $Y$ , when all variables  $V_1 \dots V_n$  in  $\mathcal{M}$  that are not on  $P$  are held fixed at their actual values.

In model  $\mathcal{M}_1$  for example,  $X_1 = 1$  and  $X_2 = 1$  are actual causes of  $Y = 1$ , but  $X_3 = 0$  is not. When we hold fixed  $X_3 = 0$ , an intervention setting  $X_1$  from 1 to 0 would result in  $Y$  changing from 1 to 0. Likewise for  $X_2$ . But when we hold fixed  $X_1 = 1$  and  $X_2 = 1$ , an intervention setting  $X_3$  from 0 to 1 would not result in  $Y$  changing value from 1.

There are several consequences of these definitions that will be important in what follows. First, notice that if  $X = x$  is an actual cause of  $Y = y$  in  $\mathcal{M}$  then  $X$  is a contributing cause of  $Y$  in  $\mathcal{M}$ . Second, notice that there may be more than one path that satisfies **PATH<sub>A</sub>**. When **AC<sub>A</sub>** is satisfied in virtue of **PATH<sub>A</sub>** being satisfied by path  $P$ , I will say that  $X = x$  is an actual cause of  $Y = y$  *along path  $P$* . Third, notice that each of these definitions is relativised to a causal model. The corresponding de-relativised definitions are as follows<sup>11</sup>:

- $X$  is a *contributing cause* of  $Y$  *simpliciter iff* there exists a model in which  $X$  is a contributing cause of  $Y$ ; and
- $X = x$  is an *actual cause* of  $Y = y$  *simpliciter iff* there exists a model in which  $X = x$  is an *actual cause* of  $Y = y$ .

I will return to the relationship between the relativised and de-relativised definitions in Section 4.3. However, because it will be important later, note that the de-relativised definitions do *not* require that for two variables to be causally related in either of these senses, *every* model containing those variables must represent them as causally related. One is enough<sup>12</sup>.

<sup>11</sup>Here I follow Hitchcock (2007, p. 503) and Woodward (2008b). For an argument that interventionist definitions should not be de-relativised in this way, see Statham (2018).

<sup>12</sup>In the terms employed by Stern and Eva (2023), interventionism so understood adopts the *Weak Causation Principle* but not the *Strict Causation Principle*.

Because it will simplify the discussion later, I will also introduce the following de-relativised definition here:

- **Causal Chain.** There is a *causal chain* containing  $X$  and  $Y$  iff there exists a model in which  $X$  and  $Y$  are members of the same path.

This outline is sufficient to exhibit some of the key features of interventionism. First, the theory does not provide an analysis or reduction of causation but rather an explication of causal claims in terms of interventions. The concept of an intervention is itself clearly causal in character, and in the interventionist framework it is explicitly defined in causal terms. What is important for present purposes is that the truth of causal claims can be established independently of any such analysis or reduction—it is whether or not it is true that mental properties sometimes causally explain physical events that is at issue in the exclusion problem, not whether these explanations can be grounded in a reductionist account of causation. Second, this is a kind of counterfactual account of causation—causal claims involve what *would* happen given some particular intervention, not what *actually* or *will* happen. Third, causal claims are model-relative in the sense that they are only well-defined with respect to the variables in a particular model. However, as should be clear, this is not a version of causal anti-realism. Causal claims are not made true or false by causal models, they are made true or false by the counterfactuals regarding experimental interventions that are represented by those models<sup>13</sup>. Moreover, because the counterfactuals are explicitly formulated in terms of interventions, it is typically transparent how they can be tested empirically. Nevertheless, as is clear from the definitions above, interventionism does entail that necessarily, if some causal claim is true, then there exists a model in which it is so represented.

---

<sup>13</sup>This may seem obvious, but the following mistake is routinely made (in this case, by a Nobel Prize winner): “A model is in the mind. As a consequence, causality is in the mind” (Heckman 2005, p. 2).

### 3 Exclusion Reformulated

In the interventionist setting, the exclusion problem can be initially formulated as follows:

- **Non-Reductionism<sub>i</sub>**: The values of mental variables are distinct from, though metaphysically necessitated by, the values of physical variables.
- **Completeness<sub>i</sub>**: For every event, there exists a causal model containing only physical variables which specifies a complete explanation of that event.
- **Mental Causation<sub>i</sub>**: There exists a causal model in which a mental variable explains an event.
- **Exclusion<sub>i</sub>**: If there exists a causal model specifying a complete explanation for an event, there exists no other causal model containing distinct variables specifying an explanation for that event.

In the remainder of this section I clarify and refine these notions in turn, and then more precisely reformulate the exclusion problem in the interventionist setting.

#### 3.1 Non-Reductionism

**Non-Reductionism<sub>i</sub>** requires clarification of the distinction between mental and physical variables. It also requires clarification of the notion of a variable value being distinct from another variable value.

It is a standard presupposition in the debate over exclusion that there is a distinction between physical properties and mental properties in the sense required to generate the problem. I take no stand on how this distinction should be drawn. But to keep the relationship between models and what they represent clear, I assume that corresponding to it is a distinction between sets of variables that represent physical properties, and sets of variables that represent mental properties. I will refer to these

sets of variables as involving *vocabularies*, where a vocabulary is a set of variables with variable values that all represent a single property type. So a *physical vocabulary*  $P$  contains only variables with values representing physical properties, and a *mental vocabulary*  $M$  contains only variables representing mental properties.

In order for **Non-Reductionism** <sub>$i$</sub>  to occupy the proper role in the exclusion problem it needs to express a claim about the world, not about our ways of representing the world. So I will assume that variable values are distinct if and only if they represent distinct properties:

- **Value Distinctness.** Variable values are distinct *iff* they represent distinct properties.

In addition, note that there is a necessary condition on two variables appearing in the same model that follows from the definition of direct causation provided in Section 2. Recall that whether  $X$  is a direct cause of  $Y$  in  $\mathcal{M}$  depends, according to **DC**, on whether there exists an intervention on  $X$  that will change  $Y$  when all other variables in  $\mathcal{M}$  are held fixed at some combination of values by interventions. This implies an independence condition on variables coexisting in a model: if  $X$  is a direct cause of  $Y$  in  $\mathcal{M}$  then there must be possible values  $x$  and  $x'$  for  $X$  such that an intervention on  $X$  from  $x$  to  $x'$  is possible when all other variables except  $Y$  in  $\mathcal{M}$  are set to some combination of possible values by independent interventions. There is a natural generalisation of this independence condition standardly assumed to hold in causal models, which can be motivated by the idea that for any set of variables appearing together in a model it must be possible to non-trivially *test*, for every pair, whether **DC** holds. According to Woodward (2003, §3.5, 2015a), the relevant sense of possibility here is *at a minimum* metaphysical possibility. The corresponding independence condition on variables coexisting in a model  $\mathcal{M}$  is this:

- **Independent Manipulability.** It is metaphysically possible that ev-

ery proper subset of the variables in  $\mathcal{M}$  be set to every combination of their possible values by independent interventions<sup>14</sup>.

**Independent Manipulability** reflects the natural idea that it is only variables not related by metaphysical necessity that are candidates for being related causally. It is well known that counterfactual theories of causation are inadequate if we allow dependencies between events that are related by metaphysical necessity (Kim 1973), and **Independent Manipulability** can be seen as the constraint that implements this restriction in the interventionist framework. When I refer to the interventionist theory of causation in what follows, I will take it to include all of the definitions provided in Section 2, as well as **Value Distinctness** and **Independent Manipulability**<sup>15</sup>.

My final formulation of **Non-Reductionism** is therefore:

- **Non-Reductionism**<sub>*i*</sub>; Mental variables are distinct from physical variables in the sense that they are drawn from distinct vocabularies  $M$  and  $P$ , and the values of the  $M$ -variables are metaphysically necessitated by the values of the  $P$ -variables.

### 3.2 Completeness

**Completeness**<sub>*i*</sub> requires clarification of the notion of a complete explanation for an event. The exclusion problem is often framed in terms of causal sufficiency rather than completeness, so any defensible notion of completeness must bear some close relationship to a notion of causal sufficiency. To begin, note that this should not be interpreted as being

---

<sup>14</sup>Baumgartner (2009, p. 167) calls a related condition *Fixability*, and Woodward (2015a, p. 316, 2017, p. 255) a related condition *Independent Fixability*. See Hoffmann-Kolss (this volume) for an additional line of argument for imposing the condition.

<sup>15</sup>For more detailed discussion of the reasons for imposing constraints of these sorts, and proposals for further necessary conditions on variables, see Hitchcock (2001, 2004, 2012), Halpern and Hitchcock (2010), Blanchard and Schaffer (2017), Halpern (2016), Woodward (2016), McDonald (2023, forthcoming) and Hoffmann-Kolss (this volume).

equivalent to the claim that every event can be completely explained in terms of some fundamental physical theory. This is so for at least two reasons. First, it is an open question whether reasonable candidates for fundamental physical theories should be interpreted causally<sup>16</sup>. Second, even if reasonable candidates for fundamental physical theories should be interpreted causally, it is not the case that the structure of fundamental physical theories is identical to the structure of causal models<sup>17</sup>. So for example, sufficiency is often defined along the following lines:

- **Physical Sufficiency.** An event *A* is physically sufficient for an event *B* iff the occurrence of *A* and the laws of physics together guarantee that *B* will occur (or fix a probability for *B* such that there are no further events conditioning on which would change the probability of *B*)<sup>18</sup>.

However, **Physical Sufficiency** makes no reference to causal models, and it is not clear how it should be translated into those terms. Since I am proceeding under the interventionist assumption that causation is to be defined with respect to models, this notion is therefore inadequate for formulating the exclusion problem<sup>19</sup>. In making this point I do not mean to weaken the support that causal completeness assumptions rightly draw from the promise of complete explanations of events in terms of funda-

---

<sup>16</sup>See Russell (1913), Field (2003) and the essays in Price and Corry (2007).

<sup>17</sup>One reason is that causal models do not allow the representation of continuous processes (Strevens 2007, pp. 242–244). Strevens puts the point by saying that interventionist causal models “represent less of causal reality than is actually out there” (243), but an interventionist may consistently claim both that every interventionist model omits some causal truth, and that all causal truths are represented by some interventionist model or other (Woodward 2008b, pp. 210–211).

<sup>18</sup>See for example Papineau (2001, p. 8, 2002, p. 17). Note that the relevant notion of event here must be liberal enough to allow events involving all physical properties instantiated across the entire cross-section of a lightcone in spacetime, if any events are going to turn out to be physically sufficient for any others (Field 2003; Ismael 2009, 2011).

<sup>19</sup>See Yates (this volume) for critical discussion of a set of principles closely related to **Physical Sufficiency**.

mental physical theories. My point is simply that there is an inference involved from the success of fundamental physics to the existence of a complete causal model in the sense required to formulate the exclusion problem in the interventionist setting.

Having clarified what **Completeness<sub>i</sub>** does not say, let us examine what it does say. There are a number of different notions of causal sufficiency that can be discriminated within the interventionist framework, only one of which, I will argue, is suitable for formulating the exclusion problem.

Consider first the following definition:

- **Sufficiency in the Circumstances in a Model.** A cause is sufficient in the circumstances for an effect in a model *iff* it is an actual cause of the effect in that model.

Since **Sufficiency in the Circumstances in a Model** collapses the notion of a sufficient cause and the notion of an actual cause, it is clearly too weak to play a role in formulating an appropriate causal closure principle. As a step in the right direction, consider next:

- **Weak Sufficiency in a Model.** Call an actual cause of an effect in a model a *weakly sufficient actual cause* *iff* it is an actual cause of the effect along path *P*, and there is no possible combination of interventions on variables not on *P* that would change the effect, if the actual cause were held fixed to its actual value by an intervention. A cause is weakly sufficient for an effect in a model *iff* it is a *weakly sufficient actual cause* of that effect in that model<sup>20</sup>.

Notice that this definition, like the preceding one, is a model internal

---

<sup>20</sup>Related but distinct notions are *sustenance* in Pearl (2000, §10.2), *switch* in Woodward (2003, pp. 96–97), *strongly causes* in Halpern and Pearl (2005, p. 855) and *sufficient condition* in McDermott (1995, p. 533, 2002, pp. 96–97). To keep the formulations as simple as possible, I give definitions on which only values of single variables can be sufficient for others. The generalisation to multiple variables is obvious and does not make any difference to the arguments that follow.



one, in the sense that sufficiency is defined with respect to a single causal model, and makes no reference to facts apart from those represented by that model<sup>21</sup>. This makes trouble. For suppose that we have a model  $\mathcal{M}_P$  framed in variables drawn from physical vocabulary  $P$  which specifies a cause that is weakly sufficient for some effect. That is consistent with supposing that there exists some model  $\mathcal{M}_{PM}$  constructed by adding variables from mental vocabulary  $M$  to  $\mathcal{M}_P$ , in which the  $M$ -variables specify an actual cause for the effect and the  $P$ -variables do not specify a cause that is weakly sufficient for the effect. What this possibility reveals is that the model internal definitions of sufficiency do not adequately capture the idea, central to any closure principle, that when one class of properties is causally closed with respect to another class the latter do not make any *additional* causal difference. I conclude that an adequate closure principle must be at least as strong as:

- **Weak Sufficiency in a Weakly Closed Model.** Call a model  $\mathcal{M}_F$  framed in variables drawn from vocabulary  $F$  *weakly closed* with respect to variables drawn from vocabulary  $G$  with respect to an effect *iff*  $\mathcal{M}_F$  contains a weakly sufficient actual cause of the effect, and there exists no model  $\mathcal{M}_{FG}$ , constructed by adding variables from  $G$  to  $\mathcal{M}_F$ , in which any weakly sufficient causes in  $\mathcal{M}_F$  are not also weakly sufficient causes in  $\mathcal{M}_{FG}$ . A cause is weakly sufficient for an effect in a weakly closed model  $\mathcal{M}_F$  with respect to  $G$  *iff* it is weakly sufficient for the effect in  $\mathcal{M}_F$ .

Notice however that **Weak Sufficiency in a Weakly Closed Model** is compatible with the actual values of  $\mathcal{M}_P$  specifying a weakly sufficient actual cause of an effect, and yet it being the case that no *alternative* values of the  $P$ -variables would have specified weakly sufficient causes of *alternative* values of the effect. That is, it is compatible with the actually instantiated physical properties sufficing for some event, while any

---

<sup>21</sup>To re-emphasise a point I made in Section 2, remember that a notion being defined in a model-internal way does not imply that the corresponding fact is in any way model-dependent.

alternatively instantiated physical properties would not have sufficed for any alternative event. So **Weak Sufficiency in a Weakly Closed Model** does not yet capture the sort of closure the successes of our scientific theorising typically license us to endorse, where for some class of properties proprietary to a theory, *whichever of those properties were instantiated* would have sufficed for all outcomes of a certain type. I conclude that the closure principle appropriate to formulating the exclusion problem in the interventionist setting is:

- **Strong Sufficiency in a Strongly Closed Model.** Call a cause *strongly sufficient* for an effect in a model *iff* it is weakly sufficient for the effect, and all alternative values of the cause would also be weakly sufficient for the value of the effect in any possible state of the model. Call a model  $\mathcal{M}_F$  framed in variables drawn from vocabulary  $F$  *strongly closed* with respect to variables drawn from vocabulary  $G$  with respect to an effect *iff*  $\mathcal{M}_F$  contains a strongly sufficient actual cause of the effect, and there exists no model  $\mathcal{M}_{FG}$ , constructed by adding variables from  $G$  to  $\mathcal{M}_F$ , in which any strongly sufficient causes in  $\mathcal{M}_F$  are not also strongly sufficient causes in  $\mathcal{M}_{FG}$ . A cause is strongly sufficient for an effect in a strongly closed model  $\mathcal{M}_F$  with respect to  $G$  *iff* it is strongly sufficient for the effect in  $\mathcal{M}_F$ .

It is important to note an immediate consequence of this definition. If a cause is strongly sufficient for an effect in a strongly closed model  $\mathcal{M}_F$  with respect to  $G$ , then in any model  $\mathcal{M}_{FG}$ , constructed by adding variables from  $G$  to  $\mathcal{M}_F$ , there are no paths from any variables in  $G$  to the effect.

**Strong Sufficiency in a Strongly Closed Model** is a model external definition but still a relative one, in the sense that a cause could be strongly sufficient in a strongly closed model with respect to one set of variables, but not with respect to a different set of variables<sup>22</sup>. While it would

---

<sup>22</sup>This is also true of the closely related probabilistic conception of completeness in

not make a difference to the argument below if we strengthened our understanding of completeness yet again, so that it involved the idea of a model strongly closed with respect to *all* other variables, I prefer the present formulation. This is because understanding the problem in this way captures the great variability in the way completeness assumptions are formulated. Sometimes the worrying complete or sufficient explanation is supposed to be provided by physics, sometimes by biology, sometimes by neuroscience, sometimes by (at least the “syntactical” explanations appearing within) cognitive science<sup>23</sup>. In my view the exclusion problem can be posed in terms of these different sciences precisely because it is reasonable to believe that there exist strongly sufficient causes, in models framed in variables drawn from the vocabularies of each of these sciences, which are strongly closed with respect to the *M*-variables<sup>24</sup>. If I had all of the physical information relevant to you my knowledge of what you will and would do would not be increased by knowing any further mental information about you—and likewise if I had all of the biological information, or all of the neuroscientific information, or all of the (“syntactic”) cognitive scientific information<sup>25</sup>. Moreover, once we understand completeness in the way I have suggested, it can be seen that the exclusion problem generalises—the physical causal model is strongly closed with respect to the variables of the biological causal model, the biological causal model is strongly closed with respect to the variables of the neuroscientific causal model, and so on up the hierarchy of the

---

Sober (1999a, p. 139).

<sup>23</sup>The emphasis on physical causes is familiar from Kim (1998, 2005). As is made clear in Kim (1989), Kim was generalising from an argument initially formulated by Malcolm (1968) in terms of a “neurophysiological theory”. The emphasis on syntactic causes is familiar from Field (1978) and Stich (1983).

<sup>24</sup>For a historical survey of how completeness in physics and biology became compelling, see Papineau (2001). For a comparison with the assumptions that generated earlier problems with mental causation, see Patterson (2005).

<sup>25</sup>See Loewer (2008, 2009). Note that this is not to say that my *explanations* would not be improved by the possession of this information. Indeed, I think they would be (Weslake 2010).

sciences and never *vice versa*<sup>26</sup>. And so if the exclusion problem arises for mental variables it also arises for any variables not appearing in some maximally strongly closed causal model<sup>27</sup>.

Because my central interest is in **Exclusion**, in what follows I will not defend these claims, and will simply proceed under the assumption that a closure principle concerning physical and mental variables formulated in terms of **Strong Sufficiency in a Strongly Closed Model** is true<sup>28</sup>. My final formulation of **Completeness** is therefore:

- **Completeness<sub>j</sub>** For every event, there exists a causal model with variables drawn from  $P$ , which is strongly closed with respect to  $M$ , in which there is a strongly sufficient actual cause  $P_i$  for that event.

### 3.3 Mental Causation

My initial formulation of **Mental Causation** is straightforward:

- **Mental Causation<sub>j</sub>** There exists a causal model with variables drawn from  $M$  in which a mental variable  $M_i$  is an actual cause of an event.

---

<sup>26</sup>It is a *hierarchy* in part *because* this relation is asymmetric in this way.

<sup>27</sup>Here I side with Bontly (2002) against Kim (1997, 1998).

<sup>28</sup>There are two options available to someone who accepts **Non-Reductionism<sub>j</sub>** but wishes to deny **Completeness<sub>j</sub>**. One is to deny **Physical Sufficiency**, and with it **Completeness<sub>j</sub>**. In my view the most interesting arguments of this sort are those made by Cartwright (1983, 1994) as developed in the case of chemistry by Hendry (2006, 2010a,b, 2017). See Sklar (2003) for the general line of response that I think blocks these arguments. The other is to accept **Physical Sufficiency** but deny that it entails **Completeness<sub>j</sub>**. One strategy here turns on the idea that causes must be “proportional” to their effects (List and Menzies 2009; Menzies and List 2010; Raatikainen 2010; Yablo 1992; and for critical discussion Weslake 2017). Another strategy, the “dual explanandum” or “intralevelist” solution to the exclusion problem, turns on the idea that the effects of mental causes are individuated mentally rather than physically (the position dates at least to Putnam 1975; see also Gibbons 2006; Marras 1998; Schlosser 2009; Thomasson 1998; and for critical discussion Buckareff 2011, 2012; Sober 1999b).

Three comments on this formulation, before I introduce a revision in the following section.

First, interventionism is attractive not only as a theory of causation generally, but as a theory of mental causation specifically. In particular, it has been defended as providing a good framework for understanding causal explanation in psychology (Campbell 2007; see also Kaiserman 2020; Rescorla 2018), in psychiatry (Campbell 2008; Kendler and Campbell 2009), and in folk psychology (Menzies 2010). If interventionism is true, there is no special problem in understanding how a mental variable can be a cause.

Second, **Mental Causation**<sub>j</sub> might be granted, and yet it might be argued that only a model containing physical variables *really* represents causes, and that models containing mental variables merely specify *explanations*, or some other weak cousin of causation. This might be because only the physical model promises to be maximally predictively accurate and therefore maximally strongly closed (Davidson 1963, 1967, 1970, 1995), or because the physical model is a truthmaker for the mental model (Crane 2008; Robb, Heil, and Gibb 2023, §5.3), or for more recherché metaphysical reasons (Jackson and Pettit 1988, 1990a,b). In my view the arguments for these claims are unsound, but for present purposes I simply note that if they succeed they are arguments against interventionism in general and therefore should be addressed as independent claims about the nature of causation and causal explanation. Proceeding under the presumption of the truth of interventionism, I here set them to one side<sup>29</sup>.

Third, while this will also make no difference to the argument below, note that **Mental Causation**<sub>j</sub> does not require that in order for mental variables to causally explain, they must be either weakly or strongly sufficient for their effects. It is unclear to me why the exclusion problem is often framed so that mental causes must be sufficient for their effects in a stronger

---

<sup>29</sup>See Burge (2007b) and Woodward (2008a, pp. 244–249) for arguments against some of these lines of objection.

sense than sufficiency in the circumstances. Bennett (2003, §5) thinks that anything less than sufficiency would endanger the “full-fledged causal efficacy of the mental” (481), granting it merely “a derivative efficacy” (482). I cannot see the motivation for claims of this form if sufficiency is supposed to be stronger than circumstantial sufficiency—especially given the metaphors that are often used to characterise the exclusion problem<sup>30</sup>. If an event has a complete physical cause, mental causes are often said to have “no work left to do” (Kim 1998, 35, 37, 54, 110, 126 n. 6), “no gaps left to fill” (Menzies 2003), no opportunity to “inject themselves” into the causal order (Kim 1998, p. 41, 2005, p. 16); if there is no lowest level of causation, we are supposed to worry that causal powers will “drain away” (Block 2003; Kim 2003). But if there *were* work left to do, a gap to be filled, an injection to be provided or a drain to be plugged, presumably the context would be almost sufficient, and the additional impetus plus context would be wholly sufficient. The work, filler, injection or plug would not itself be wholly sufficient, but rather would be sufficient in the circumstances. Now perhaps these are all just poor metaphors for what is supposed to be at issue here; but metaphors aside, the claim in question would be that any actual causes that are not at least weakly sufficient must have merely derivative efficacy. Given that sufficiency in the circumstances is the sort of efficacy most causes have in most scientific theories, I say that derivative causes in this idiosyncratic sense would be causes enough for mental causation<sup>31</sup>.

### 3.4 Exclusion

It may seem that the formulation of **Exclusion** is now straightforward: it should simply be the strongest principle that is inconsistent with **Non-Reductionism<sub>j</sub>**, **Completeness<sub>j</sub>** and **Mental Causation<sub>j</sub>**. However, for the principle to have any *prima facie* plausibility, it needs to be weaker. To translate a point first made by Goldman (1969, pp. 470–473, 1970,

<sup>30</sup>I do not suggest Bennett endorses the position I here criticise.

<sup>31</sup>For a more detailed argument for this claim, see Woodward (2008a, pp. 245–249).

pp. 159–161) into this context, **Completeness<sub>j</sub>** is perfectly compatible with **Mental Causation<sub>j</sub>** in a case where there is a path from the mental variable that is an actual cause of the event to the physical variable that is strongly sufficient for the event, or in a case where the mental variable is on a path from the physical variable to the event. A principle that says that if an event has a sufficient cause it has no other causes is clearly far too strong to be plausible, for it is inconsistent with the existence of causal chains.

My final formulation of **Exclusion** is therefore:

- **Exclusion<sub>j</sub>** If there exists a causal model with variables drawn from a vocabulary  $F$ , which is strongly closed with respect to variables drawn from vocabulary  $G$ , and in which a variable  $F_1$  is a strongly sufficient actual cause for an event, then there exists no causal model in which a variable  $G_1$  drawn from vocabulary  $G$  is an actual cause of that event, unless there is a causal chain containing  $F_1$  and  $G_1$ .

An attractive feature of this formulation is that it reveals a way in which someone who rejects **Non-Reductionism<sub>j</sub>** can evade the exclusion problem. Here I have in mind Lowe (2000, 2003), who has argued that all causal closure principles with strong empirical support are logically consistent with non-physicalist theories on which mental causes occupy a place in causal chains *between* sufficient physical causes and their effects. While I think we have overwhelmingly strong reasons to reject theories of this sort, my formulations of **Completeness<sub>j</sub>**, **Exclusion<sub>j</sub>**, and **Mental Causation<sub>j</sub>** support Lowe's claim.

The non-reductive physicalist, on the other hand, is in no position to make a similar move. They would thereby be committed to a position on which mental causes occupy a place in causal chains between sufficient physical causes and their effects, but are metaphysically necessitated by *different* physical variables, which *are not themselves sufficient for those*

*effects*. It is rare to find a position in logical space no philosopher is willing to occupy, but this must be one of them<sup>32</sup>. My final formulation of **Mental Causation<sub>j</sub>** is therefore the following, which closes this loophole and renders the propositions that form the exclusion problem logically inconsistent:

- **Mental Causation<sub>j</sub>** There exists a causal model with variables drawn from  $M$  in which a mental variable  $M_1$  is an actual cause of an event, and there is no causal chain containing  $P_1$  and  $M_1$ .

### 3.5 The Interventionist Exclusion Problem

Putting this all together, I conclude that the exclusion problem in the interventionist framework should be formulated as follows:

- **Non-Reductionism<sub>j</sub>** Mental variables are distinct from physical variables in the sense that they are drawn from distinct vocabularies  $M$  and  $P$ , and the values of the  $M$ -variables are metaphysically necessitated by the values of the  $P$ -variables.
- **Completeness<sub>j</sub>** For every event, there exists a causal model with variables drawn from  $P$ , which is strongly closed with respect to  $M$ , in which there is a strongly sufficient actual cause  $P_1$  for that event.
- **Mental Causation<sub>j</sub>** There exists a causal model with variables drawn from  $M$  in which a mental variable  $M_1$  is an actual cause of an event, and there is no causal chain containing  $P_1$  and  $M_1$ .
- **Exclusion<sub>j</sub>** If there exists a causal model with variables drawn from a vocabulary  $F$ , which is strongly closed with respect to variables drawn from vocabulary  $G$ , and in which a variable  $F_1$  is a strongly sufficient actual cause for an event, then there exists no causal model in which a variable  $G_1$  drawn from vocabulary  $G$  is an actual

---

<sup>32</sup>See Kim (1998, pp. 37, 40, 44). As Kim notes, the non-reductive physicalist will invariably be committed to versions of physicalism and closure on which this option is ruled out.



cause of that event, unless there is a causal chain containing  $F_1$  and  $G_1$ .

One of these claims must go. We are finally in a position to consider arguments for **Exclusion<sub>j</sub>**.

## 4 Compatibilism Examined

In this section I evaluate two arguments for **Exclusion<sub>j</sub>**. The first is familiar from discussion of the exclusion problem in the interventionist setting, but the second is not. This is because the discussion has almost universally assumed a particular conception of the relationship between physical and mental variables, according to which the mental cause  $M_1$  that figures in **Mental Causation<sub>j</sub>** is metaphysically necessitated by the strongly sufficient actual cause  $P_1$  that figures in **Completeness<sub>j</sub>**. I will call this assumption *subvenience sufficiency*, the non-reductive physicalist position that accepts it *internalism* and the non-reductive physicalist position that denies it *externalism*. As I will show, the distinction is important. I begin with a discussion of subvenience sufficiency itself, and then consider internalism and externalism in turn. I side with those who take the interventionist to have a good response to the argument for **Exclusion<sub>j</sub>** under the assumption of internalism. But I go on to argue that the response the interventionist has to the argument for **Exclusion<sub>j</sub>** under the assumption of externalism serves to expose a weakness in interventionism itself.

### 4.1 Subvenience Sufficiency

As formulated, the exclusion problem invites us to consider two causal models. Each model contains a variable  $E_1$  that is a candidate effect for a mental cause. The first model,  $\mathcal{M}_P$ , the existence of which is entailed by **Completeness<sub>j</sub>**, contains (in addition to  $E_1$ ) only variables drawn from  $P$ , is strongly closed with respect to  $M$ , and contains a strongly sufficient

actual cause  $P_1$  for  $E_1$ . The second model,  $\mathcal{M}_M$ , the existence of which is entailed by **Mental Causation<sub>j</sub>**, contains (in addition to  $E_1$ ) variables drawn from  $M$ , and contains an actual cause  $M_1$  of  $E_1$ . As is also entailed by **Mental Causation<sub>j</sub>**, there is no causal chain containing  $P_1$  and  $M_1$ . I will make the simplifying assumptions that  $\mathcal{M}_M$  contains (in addition to  $E_1$ ) only variables drawn from  $M$ , and that both  $\mathcal{M}_P$  and  $\mathcal{M}_M$  only contain variables on paths terminating in variable  $E_1$ .

With respect to models of this sort, subvenience sufficiency can be defined as follows:

- **Subvenience Sufficiency.** Given two models  $\mathcal{M}_F$  and  $\mathcal{M}_G$ , where each model only contains variables on paths terminating in variable  $E_1$ ,  $\mathcal{M}_F$  is *subvenience sufficient* with respect to  $\mathcal{M}_G$  and  $E_1$  iff the values of all other variables in  $\mathcal{M}_G$  are metaphysically necessitated by the values of strongly sufficient causes of  $E_1$  in  $\mathcal{M}_F$ .

It is important to see that  $\mathcal{M}_P$  being subvenience sufficient with respect to  $\mathcal{M}_M$  is a substantive assumption that is not itself entailed by **Non-Reductionism<sub>j</sub>** either alone or in conjunction with **Completeness<sub>j</sub>** and **Mental Causation<sub>j</sub>**. In particular, while **Non-Reductionism<sub>j</sub>** merely requires that the values of the  $M$ -variables are metaphysically necessitated by the values of the  $P$ -variables,  $\mathcal{M}_P$  being subvenience sufficient with respect to  $\mathcal{M}_M$  imposes the much stronger constraint that the  $M$ -variables are metaphysically necessitated by the very  $P$ -variables that are strongly sufficient for their effects. The assumption is vividly illustrated by what Loewer (2015, p. 60) calls “Kim’s Favourite Diagram” (2003, p. 159), a way to represent the exclusion argument that is ubiquitous in Kim’s work, in which one and the same physical event is represented as both the cause of a given effect, and as the subvenience basis for the mental event which Kim takes it to exclude (see Figure 2).

I will refer to non-reductive physicalism in conjunction with **Subvenience Sufficiency** as *internalism* and non-reductive physicalism without

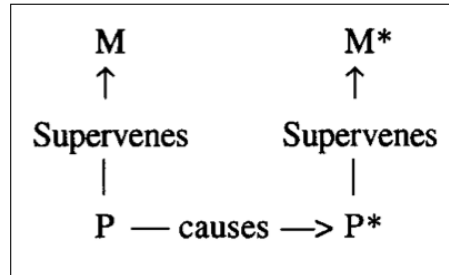


Figure 2: Kim's Favourite Diagram.

**Subvenience Sufficiency as externalism.** The fact that internalism has been so frequently assumed in the discussion of the exclusion problem would be unremarkable if it were not the case that most non-reductive physicalists are committed to rejecting it, and if it did not make a difference to the arguments available to the non-reductive physicalist for rejecting **Exclusion<sub>j</sub>**. As Bennett (2003) notes, internalism is inconsistent with content externalism, with functionalism in general, and with conceptual role semantics in particular<sup>33</sup>. These are the most prominent of the theories of mental properties that motivate non-reductionism in the first place, so it is hardly open to the non-reductive physicalist to ignore their consequences. I begin, however, with internalism.

## 4.2 Internalism

Baumgartner (2009, 2010) has argued that if interventionism is true, then whenever variables stand in relationships of metaphysical necessitation, the necessitated variable cannot have any of the same effects as the necessitating variable. If that is right, then internalism is incoherent. For the internalist is committed, by virtue of the claim that  $\mathcal{M}_P$  is subvenience sufficient with respect to  $\mathcal{M}_M$ , to the existence of variables that

<sup>33</sup>In addition, Worley (1993, §5) argues that internalism is inconsistent with anomalous monism and the folk-psychological platitude that a single mental state may be a cause of different effects on different occasions

stand in exactly this relationship. In this section I argue that the correct formulation of interventionism blocks this argument<sup>34</sup>.

As was clear from the definitions introduced in Section 2, all of the fundamental interventionist causal concepts are defined in terms of interventions. Baumgartner's argument depends on the way in which interventions are defined by Woodward (2003, p. 98). Woodward first introduces the type-level notion of an *intervention variable*:

- **IV.** *I* is an *intervention variable* for *X* with respect to *Y* iff:
  - **I<sub>1</sub>.** *I* is a contributing cause of *X*.
  - **I<sub>2</sub>.** There is a model in which *I* has at least one value that is weakly sufficient for the value of *X*.
  - **I<sub>3</sub>.** Every causal chain from *I* to *Y* contains *X*.
  - **I<sub>4</sub>.** *I* is statistically independent of every contributing cause of *Y* on causal chains that do not contain *X*.

This is then used to define the token-level notion of an *intervention*:

- **IN.** *I = i* is an *intervention* on *X* with respect to *Y* iff *I* is an intervention variable for *X* with respect to *Y* and there is a model in which *I = i* is a weakly sufficient cause of the value of *X*.

Note that I have presented definitions that are weaker than Woodward's in one respect, since **I<sub>2</sub>** is weaker than Woodward's condition. The distinction amounts to whether interventions must be *hard*, so that they override all other causal connections to the variable intervened on, or whether they may be *soft*, and merely make an additional causal impact to the variable intervened on. I opt for the weaker definitions because Baumgartner's argument works either way, and because Woodward himself accepts both formulations (2015b, p. 3584, 2015a, p. 321 fn. 15, 2017, p. 254 fn. 3)<sup>35</sup>.

---

<sup>34</sup>I am indebted here to correspondence with Michael Baumgartner.

<sup>35</sup>For discussion of hard and soft interventions, see Korb et al. (2004); Markowetz, Grossmann, and Spang (2005); Eberhardt and Scheines (2007); Eberhardt (2014) and

Baumgartner’s argument is simple, with each premise following from the definitions of the relevant notions, or from claims the internalist is committed to accepting. For  $M_1$  to be an actual cause of  $E_1$  in  $\mathcal{M}_M$ , it must be a contributing cause of  $E_1$  in  $\mathcal{M}_M$  (**AC<sub>A</sub>**). For it to be a contributing cause of  $E_1$  in  $\mathcal{M}_M$ , there must be an intervention on  $M_1$  with respect to  $E_1$  (**CC**). For there to be an intervention on  $M_1$  with respect to  $E_1$ , there must be an intervention variable  $I$  for  $M_1$  with respect to  $E_1$  (**IN**). For there to be an intervention variable  $I$  for  $M_1$  with respect to  $E_1$ , every causal chain from  $I$  to  $E_1$  must contain  $M_1$  (**IV, I<sub>3</sub>**), and  $I$  must be statistically independent of every contributing cause of  $E_1$  on causal chains that do not contain  $M_1$  (**IV, I<sub>4</sub>**). But internalism is committed to the claim that  $\mathcal{M}_P$  is subvenience sufficient with respect to  $\mathcal{M}_M$ . This entails that there is no way to make a change to  $M_1$  without also changing  $P_1$ , which in turn entails both that there is a causal chain from  $I$  to  $P_1$  to  $E_1$  that does not contain  $M_1$ , and that  $P_1$  is not statistically independent of  $M_1$ . So there is no such intervention variable  $I$ ,  $M_1$  is not an actual or contributing cause of  $E_1$ , and there is no such model  $\mathcal{M}_M$  as required by the internalist<sup>36</sup>.

In response to this argument, Woodward (2015a, p. 323) helpfully distinguishes three questions. First, are the definitions that lead to this result adequate interpretations of Woodward (2003)? Second, must any interventionist theory that deserves the name adopt definitions that lead to this result? Third, in order for variables to be causes, must they make a difference to their effects beyond the differences made by variables that metaphysically necessitate them? I set Woodward’s first question aside<sup>37</sup>. On Woodward’s second question, some authors have considered ways in

---

in the psychological context Campbell (2007); Kaiserman (2020).

<sup>36</sup>As Baumgartner (2009, p. 171) notes, the argument does not depend on a premise concerning causal closure: it can be used to show that, on these definitions, no variables related by metaphysical necessity can share any effects. Gebharder (2017b) argues that this is also the case with respect to the argument in Gebharder (2017a), and Stern and Eva (2023) agree.

<sup>37</sup>Woodward (2015a, pp. 324–325, 2017, p. 257) has argued that the answer is “no”. As he says, it is the least interesting of the three.

which the basic interventionist framework can be expanded, so that variables related both causally and by metaphysical necessitation can appear in the same model. The debate then becomes whether the principles that should govern models of this sort should permit or prohibit causation by necessitated variables (Baumgartner 2010; Gebharder 2017a; Stern and Eva 2023; Woodward 2015a). If we wish to restrict our focus to causal relationships, however, there exists a more conservative amendment of the interventionist framework that is sufficient to block Baumgartner’s argument<sup>38</sup>.

The amendment is as follows:

- **IV\***. *I* is an *intervention variable* for *X* with respect to *Y* iff:
  - **I<sub>1</sub>**. *I* is a contributing cause of *X*.
  - **I<sub>2</sub>**. There is a model in which *I* has at least one value that is weakly sufficient for the value of *X*.
  - **I<sub>3</sub>\***. Every path from *I* to *Y* goes through *X* in every model containing *I*, *X* and *Y*.
  - **I<sub>4</sub>\***. *I* is statistically independent of every contributing cause of *Y* on paths that do not contain *X* in every model containing *I*, *X* and *Y*.

The difference between **IV** and **IV\*** concerns the third and fourth conditions. Condition **I<sub>3</sub>** requires that there are no paths, in any model, from *I* to *Y* without *X*. Condition **I<sub>3</sub>\*** relaxes this requirement, requiring that there are no paths from *I* to *Y* without *X* in any model that contains those variables. Likewise, condition **I<sub>4</sub>** requires that *I* is statistically dependent of contributing causes of *Y*, in all models, that are on paths without *X*. **I<sub>4</sub>\*** relaxes this requirement, requiring that *I* is statistically dependent of contributing causes of *Y*, on paths without *X*, in any model that contains those variables.

---

<sup>38</sup>The basic strategy I develop in the remainder of this section is also proposed by Eronen and Brooks (2014), who cite an earlier version of this paper. I do not endorse their arguments for it.

The difference between these definitions shows up as a consequence of **Independent Manipulability** (Section 3.1), according to which it is a condition on variables coexisting in a model that it be metaphysically possible for every proper subset to be set to every combination of their possible values by independent interventions. This entails:

- **Non-Necessitation.** A causal model cannot contain a variable with a possible value that is metaphysically necessitated by a possible combination of values of any proper subset of the other variables in the model.

An immediate consequence of this, if internalism is true, is that there are no causal models that contain both  $P_1$  and  $M_1$ . This in turn blocks Baumgartner's argument. According to  $\mathbf{IV}^*$ , the presence of a causal chain from  $I$  to  $P_1$  to  $E_1$  that does not contain  $M_1$ , and the fact that  $P_1$  is not statistically independent of  $M_1$ , do not threaten the satisfaction of  $\mathbf{I}_3^*$  and  $\mathbf{I}_4^*$ . More generally, the fact that any change to  $M_1$  entails some corresponding change to  $P_1$  is no obstacle to there being well defined interventions on  $M_1$  (see Figure 3).

The answer to Woodward's second question, therefore, is "no". There is a coherent formulation of an interventionist theory of causation that does not generate the consequence that a necessitated variable cannot have any of the same effects as the necessitating variable. Moreover, it is a formulation that is perfectly suited to the non-reductive physicalist, in the following sense. As Bennett (2008) argues, the non-reductive physicalist would ideally like a solution to the exclusion problem that can play two roles. On the one hand, it should show that causal considerations do not force her into reductive physicalism. On other hand, it should show that causal considerations are still a problem for the dualist<sup>39</sup>. Intervention-

---

<sup>39</sup>In the current framework, *reductive physicalism* can be defined as the position on which mental variables are not distinct from physical variables, and *dualism* can be defined as the position on which mental variables are distinct from physical variables, and the values of the mental variables are not metaphysically necessitated by the values




Figure 3: Diagram for  $\mathcal{M}_M$ , if internalism is true. Solid lines represent variables and direct causes that are in the model. Dashed lines represent variables and direct causes that are not in the model. Double arrows represent metaphysical necessitation. According to **IV**,  $I$  cannot be an intervention variable for  $M_1$  with respect to  $E_1$ . According to **IV\***, it can.



ism formulated in terms of  $IV^*$  has both consequences, at least for the internalist. On the one hand, as I have just argued, the internalist can argue that their position is coherent, and compatible with mental causation. On the other hand, the internalist can point out that the dualist, in virtue of rejecting the metaphysical necessitation of the values of mental variables by the values of physical variables, cannot avail themselves of the same sorts of interventions on mental variables. Instead, they must commit to the existence of interventions on mental variables that do not entail any changes to physical variables, and are statistically independent of physical variables. But in that case, accepting **Completeness<sub>j</sub>** would force the dualist to admit that while there may be such interventions, they could not result in any downstream effects. As a result, the dualist who accepts **Completeness<sub>j</sub>** is committed to rejecting **Mental Causation<sub>j</sub>**. Here is a different way to see the point. **Completeness<sub>j</sub>** says that  $\mathcal{M}_p$  is strongly closed with respect to  $M$ . This means that there is no model containing all of the variables from  $\mathcal{M}_p$ , and any variables from  $M$ , in which any strongly sufficient causes in  $\mathcal{M}_p$  lose that status. But there are two ways this can be true. The first is for there to be *no such model* containing all of the variables from  $\mathcal{M}_p$  and any variables from  $M$ . This is what internalism is committed to, and it is consistent with the existence of a model in which  $M_1$  is a cause. The second is for there to be such a model. This is what dualism is committed to, and it is not consistent with  $M_1$  being a cause, in that model or any other<sup>40</sup>.

I turn now to Woodward's third question. What can be said to recommend interventionism formulated in terms of  $IV^*$  over interventionism formulated in terms of  $IV$ , besides the fact that it facilitates a coherent non-reductionism?

---

of the physical variables.

<sup>40</sup>Here I disagree with Shapiro and Sober (2007), who suggest that a well conceived argument for epiphenomenalism, under the assumption of interventionism, "should aim to show that one class of properties does not affect a second class, not that the first has no effects at all" (241). This underestimates how strong the constraints are that completeness puts on the sorts of properties that can be causes.

One argument derives from the very motivation for the theory. If there is a single idea at the heart of interventionism, it is that the best way to understand the nature of causation is to theorise through the lens of an ideal experiment for detecting it (Woodward 2003, p. 14). So for example, when I introduced **Independent Manipulability** (Section 3.1), I said that it is motivated by the idea that for variables to coexist in a model, it must be possible to non-trivially test, for every pair, whether they are related by direct causation. The same motivation can be given for  $IV^*$ . The idea behind conditions  $I_3$  and  $I_4$  is that interventions should be independent of potential confounding causes. But just as the interventionist should say that variables are only candidates for being causally related if they can be independently manipulated, they should say that variables are only candidates for being potential confounding causes if they can be independently manipulated.  $I_3$  and  $I_4$  do not entail this constraint, but  $I_3^*$  and  $I_4^*$  do.

This is not a new form of argument. In his discussion of causal completeness in the context of probabilistic theories of causation, Sober (1999a, §2) considers theories according to which a positive causal factor must raise the probability of an effect in at least one background context, and lower it in none:

- **Positive Causal Factor.**  $C$  is a *positive causal factor* for  $E$  iff  $P(E|C\&X_i) \geq P(E|\neg C\&X_i)$  for all background contexts  $X_i$ , with strict inequality for at least one  $X_i$ .

What counts as a background context? According to Sober, a necessary condition on a set of properties constituting a background context relative to a given cause and effect is that these probabilities are well defined. As he notes, this entails that when evaluating whether a necessitated property is a causal factor, necessitating properties cannot be part of any background context, since then  $P(\neg C\&X_i)$  would be 0 and  $P(E|\neg C\&X_i)$  not well defined. Sober's argument is identical to the argument I have just given

for **IV\***, transposed to the probabilistic case<sup>41</sup>.

The convergence of these arguments underscores that the issue concerning exclusion for difference-making theories of causation, under the assumption of internalism, concerns the contexts relative to which causes must make a difference. Must they make a difference controlling for *all* other causes, or must they make a difference controlling for all other *independent* causes? (Shapiro and Sober 2007, p. 241). I will briefly describe three other arguments that can be given for the second conception, before turning to externalism.

First, it is implicit in scientific practice that you do not need to control for necessitating variables in order to be justified in believing that necessitated variables are causes (Shapiro 2010; Shapiro and Sober 2007). As Sober puts it: “This fact about scientific practice stands on its own” (1999a, p. 147). In this connection, it is also important to note that allowing necessitated properties to be causes does not mean that they trivially satisfy the requirements to be causes, simply in virtue of their being necessitated by causes (Segal and Sober 1991; Shapiro and Sober 2007, pp. 256–259; Woodward 2015a, 2017). It is a substantive and difficult matter to determine whether a necessitated property meets the conditions for causation by the lights of interventionism under the assumption of **IV\***.

Second, it is clear that in examples involving logical or conceptual necessitation between variables, we do not need to hold fixed one variable in order to determine whether the other makes a difference (Woodward 2008a, 2015a). Indeed, since for any variable we can introduce others related to it in these ways, imposing this requirement would mean that no variables could possibly satisfy the requirements for being causally

---

<sup>41</sup>The same line of thought is arguably implicit in Eells (1991, p. 31). Similarly, Humphreys (1989, p. 74) requires that it be physically possible for the cause and its absence to occur relative to all background factors (for an application to the exclusion problem, see Henderson 1994). An analogous argument, in the context of a theory of causation along the lines of Mackie (1974), is given by Melnyk (2003, pp. 137–138).

related. It can then be argued either that metaphysical necessitation is relevantly similar to those forms of dependence, or that **IV\*** provides the correct theory in light of that fact.

Third, it can be argued that **IV** is, but **IV\*** is not, subject to the argument that if there were no fundamental causal level, causation would drain away (Block 2003; Kim 2003).

I do not claim that these arguments are collectively decisive. But I do claim that in interventionism formulated with **IV\***, the non-reductionist has a coherent and well motivated theory of causation that entails the falsity of **Exclusion<sub>j</sub>**. If the non-reductionist is an internalist, there is no obstacle to their endorsing **Mental Causation<sub>j</sub>**.

### 4.3 Externalism

As it happens, very few non-reductionists can rest content at this point. For most of the conceptions of mental properties that motivate non-reductionism in the first place entail that internalism is false. So we need to consider arguments that target the externalist conception of mental properties. I will begin by discussing the externalist position generally, and then discuss some of the more concrete forms it may take when they become relevant.

The first point to note is that the externalist cannot make use of the same line of reasoning available to the internalist, who can appeal to **Non-Necessitation** in order to argue that there is no model containing both  $M_1$  and  $P_1$ . Since the externalist by definition rejects the necessitation of  $M_1$  by  $P_1$ , there is no obstacle to the existence of a model that contains both variables. Since the externalist remains a physicalist, they must thereby be committed to the existence of other physical variables that, together with  $P_1$ , necessitate  $M_1$ . For simplicity, I will use a single variable  $P_2$  to represent these. So the externalist is committed to  $P_1$  and  $P_2$  together necessitating  $M_1$ , and neither  $P_1$  nor  $P_2$  alone necessitating  $M_1$ .

It now appears that epiphenomenalism looms. Consider model  $\mathcal{M}_{PM_1}$ , containing variables  $P_1$ ,  $M_1$  and  $E_1$ . As I noted in Section 3.2, **Completeness<sub>j</sub>** entails that there is no path from  $M_1$  to  $E_1$  in  $\mathcal{M}_{PM_1}$ . Moreover, this is not because there cannot be an intervention variable for  $M_1$  with respect to  $E_1$ . There can be, but it must involve changing  $M_1$  by changing  $P_2$  (which is permissible according to **IV\***). However, a difference of that sort cannot make any additional difference to  $E_1$ . At least in model  $\mathcal{M}_{PM_1}$ ,  $M_1$  is epiphenomenal (see Figure 4)<sup>42</sup>.

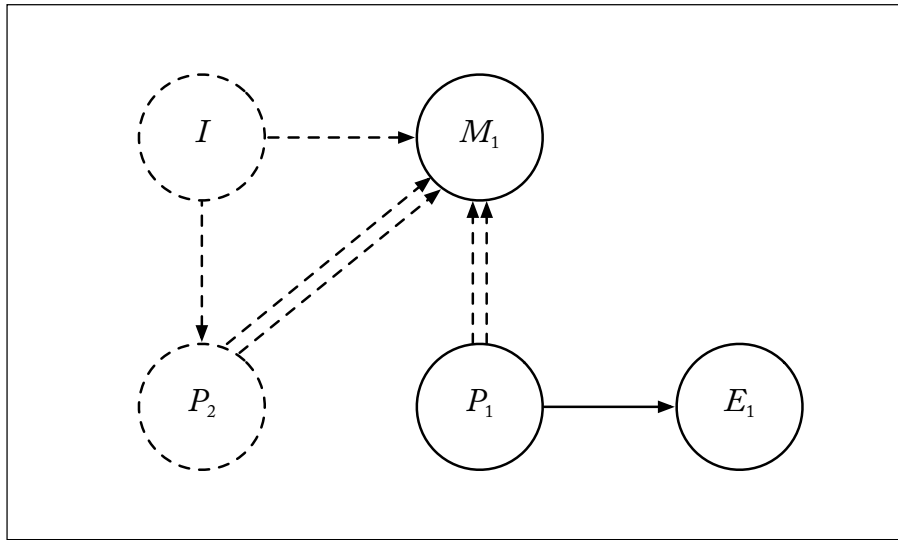


Figure 4: Diagram for  $\mathcal{M}_{PM_1}$ .

Can we conclude that  $M_1$  is epiphenomenal *simpliciter*? Not without additional argument. For according to the definitions provided in Section 2, any immediate inference from a variable not causing another in a model to its not causing it *simpliciter* is invalid. Recall that  $X = x$  is an actual cause of  $Y = y$  *simpliciter* iff there exists a model in which  $X = x$  is an actual cause of  $Y = y$ . It follows that there is an asymmetry in what it takes to show that a variable value is or is not an actual cause of

<sup>42</sup>Arguments of this form are discussed by Block (1990), Worley (1993) and Rescorla (2012, 2014, §7).

another. To show that a variable value *is* an actual cause, we simply need to identify a model in which it is. But to show that a variable value *is not* an actual cause, we need to show that there *does not exist* a model in which it is. What is needed to establish **Exclusion<sub>j</sub>** under the assumption of externalism is an argument that could establish the non-existence of a model in which  $M_1$  is an actual cause of  $E_1$ , on the basis that it is not an actual cause in  $\mathcal{M}_{PM_1}$ .

In fact, the externalist can do better than simply rejecting this inference. For they can exhibit a model in which  $M_1$  is a cause of  $E_1$ . Consider model  $\mathcal{M}_{PM_2}$ , containing variables  $P_2$ ,  $M_1$  and  $E_1$ . In this model, if we hold  $P_2$  fixed at some particular value, then any intervention on  $M_1$  must change  $P_1$ . So long as there exists a change of this sort that is associated with a change to  $E_1$ , then  $M_1$  will be a direct cause of  $E_1$ , and for at least one state of the model will be an actual cause of  $E_1$  (see Figure 5). Interventionism is therefore consistent not only with mental causation under the assumption of internalism, but with mental causation under the assumption of externalism.




Figure 5: Diagram for  $\mathcal{M}_{PM_2}$ .

For illustration, consider the case of content externalism. Here  $P_1$  can be interpreted as representing neurophysiological properties,  $P_2$  can be interpreted as representing content-fixing environmental properties, and  $M_1$  can be interpreted as representing externally-individuated mental properties, where the values of  $M_1$  are metaphysically necessitated by the values of  $P_1$  and  $P_2$ .  $\mathcal{M}_{PM_1}$  reveals the fact that if we hold fixed the neurophysiological properties, altering mental properties by altering the content-fixing environmental properties on which they partly depend would make no difference to behaviour.  $\mathcal{M}_{PM_2}$ , on the other hand, reveals the fact that if we hold fixed the content-fixing environmental properties, altering mental properties by altering the neurophysiological properties on which they partly depend may make a difference to behaviour. According to interventionism, the existence of the former model does not entail that mental properties are not causes of behaviour, while the existence of the latter model entails that mental properties are causes of behaviour. If the non-reductionist is an externalist, there is no obstacle to their endorsing **Mental Causation<sub>j</sub>**.

#### 4.4 A Weakness in Interventionism

I have argued that interventionism allows both internalists and externalists to consistently accept **Non-Reductionism<sub>j</sub>**, **Completeness<sub>j</sub>** and **Mental Causation<sub>j</sub>**. In other words, an interventionist is entitled to reject **Exclusion<sub>j</sub>**. However, in this section I suggest that attention to the externalist case reveals a weakness in interventionism.

The basic form of the problem is identified by Rescorla (2014, §11). As Rescorla notes, there are situations in which structurally identical models to  $\mathcal{M}_{PM_1}$  and  $\mathcal{M}_{PM_2}$  apply, and yet in which it is not the case that  $M_1$  is a cause of  $E_1$ . Take, for example, a simple pocket calculator (Haugeland 1985, pp. 121–123; Rescorla 2014, pp. 180–181). The semantic properties instantiated by the calculator during the course of a calculation (let these be represented by  $M_1$ ) are jointly determined by two factors: the physical

properties it instantiates (let these be represented by  $P_1$ ) and the interpretation to which they are subject (let this be represented by  $P_2$ ). In this context, claims parallel to those concluding the previous section can now be introduced. Consider some particular output of the calculator (let this be represented by  $E_1$ ).  $\mathcal{M}_{PM_1}$  reveals the fact that if we hold fixed the physical properties of the calculator, altering its semantic properties by altering the interpretation to which its physical properties are subject would make no difference to the output.  $\mathcal{M}_{PM_2}$ , on the other hand, reveals the fact that if we hold fixed the interpretation to which its physical properties are subject, altering its semantic properties by altering its physical properties may make a difference to the output. But semantic properties don't cause the outputs of pocket calculators (Rescorla 2012). Something has gone wrong.

Moreover, the problem cannot be evaded by simply rejecting externalism. For there are many other situations in which structurally identical models to  $\mathcal{M}_{PM_1}$  and  $\mathcal{M}_{PM_2}$  apply, and in which  $M_1$  is a cause of  $E_1$ . For example, consider a match struck in the presence of air, causing it to light. Let the presence of air be represented by  $M_1$ , the presence of oxygen be represented by  $P_1$ , the presence of all other constituents of air be represented by  $P_2$ , and the match lighting be represented by  $E_1$ .  $\mathcal{M}_{PM_1}$  reveals the fact that if we hold fixed the presence of oxygen, altering the presence of air by altering the presence of the other constituents would make no difference to the match lighting.  $\mathcal{M}_{PM_2}$ , on the other hand, reveals the fact that if we hold fixed the presence of the other constituents, altering the presence of air by altering the presence of oxygen would make a difference to the match lighting<sup>43</sup>.

---

<sup>43</sup>For discussion of this example in the context of mental causation, see Segal and Sober (1991), Tye (1991), Peacocke (1993a,b), and Segal (2004, 2009). Note that examples of this sort place pressure on a condition Woodward (2008a, 2021a,b, 2022) calls *realisation independence*, which requires that interventions must have the same effect no matter how they are realised. This condition seems to entail that in any case in which structurally identical models to  $\mathcal{M}_{PM_1}$  and  $\mathcal{M}_{PM_2}$  apply,  $M_1$  is not a cause of  $E_1$ . See Hoffman-Kolss (2014, §5) for a different argument against realisation independence.



In sum,  $M_1$  is a cause of  $E_1$  in only some of the cases in which it appears that  $\mathcal{M}_{PM_2}$  applies, and the interventionist therefore owes us an account both of the difference between the cases, and why we should believe that mental properties fall on the right side of the line<sup>44</sup>.

## 5 Conclusion

It has been almost 25 years since the publication of Jaegwon Kim's *Mind in a Physical World: An Essay on the Mind-Body Problem and Mental Causation* (1998), the canonical investigation of causal exclusion principles. In summarising his discussion of counterfactual theories of causation, Kim wrote (pp. 71–72):

[...] what the counterfactual theorists need to do is to give an *account* of just what makes those mind-body counterfactuals we want for mental causation true, and show that on that account those counterfactuals we don't want, for example, epiphenomenalist counterfactuals, turn out to be false. Merely to point to the apparent truth, and acceptability, of certain mind-body counterfactuals as a vindication of mind-body causation is to misconstrue the philosophical task at hand. [...] Such gestures only show that mind-body causation is part of what we normally take to be the real world; they go no further than a mere reaffirmation of our belief in the reality of mental causation. What we want—at least, what some of us are looking for—is a philosophical account of *how* it can be real in light of other principles and truths

---

<sup>44</sup>A referee for this volume suggests that the notion of *conditional independence* recently discussed by Woodward (2020, 2021a,b,c, 2022) may help here. I do not think so, for two reasons. First, and in my view correctly, Woodward does not propose that conditional independence is a necessary condition on causation. Second, nothing I have said entails whether or not conditional independence is satisfied in either the case of mental properties or the case of the calculator, and I do not see any principled reason for saying it must always hold in the former and never in the latter.

that seem to be forced upon us.

The “principles and truths” Kim refers to here are those that he took as the premises in his arguments for causal exclusion principles. I have argued that an interventionist is entitled to reject those principles. But I have also argued that interventionists have not yet discharged the obligations that Kim here describes. In particular, they need to explain what is defective about the application of  $\mathcal{M}_{PM_2}$  to the case of the pocket calculator. Work of this sort must proceed along two paths: the development of principled constraints on when a causal model is appropriate for a given situation, as in for example Hitchcock (2001, 2004, 2012), Halpern and Hitchcock (2010), Halpern (2016), Woodward (2016), Blanchard and Schaffer (2017), McDonald (2023, forthcoming) and Hoffmann-Kolss (this volume); and the application of these constraints to specific conceptions of the the relationship between physical and mental properties, as in for example Rescorla (2014). Only when this cumulative case has been made, for the difference between pocket calculators and minds, can an interventionist claim to have a fully principled basis for rejecting **Exclusion**.

28 July 2023

## References

- Michael Baumgartner. 2009. “Interventionist Causal Exclusion and Non-reductive Physicalism”, in *International Studies in the Philosophy of Science*, Vol. 23, No. 2, July 2009, pp. 161–178. URL: <http://doi.org/10.1080/02698590903006909>.
- . 2010. “Interventionism and Epiphenomenalism”, in *Canadian Journal of Philosophy*, Vol. 40, No. 3, Sept. 2010, pp. 359–383. URL: <https://doi.org/10.1080/00455091.2010.10716727>.
- . 2013. “Rendering Interventionism and Non-Reductive Physicalism Compatible”, in *Dialectica*, Vol. 67, No. 1, Mar. 2013, pp. 1–27. URL: <http://doi.org/10.1111/1746-8361.12008>.

- Michael Baumgartner. 2018. “The Inherent Empirical Underdetermination of Mental Causation”, in *Australasian Journal of Philosophy*, Vol. 96, No. 2, pp. 335–350. URL: <https://doi.org/10.1080/00048402.2017.1328451>.
- Brian Beakley and Peter Ludlow. 2006. *The Philosophy of Mind: Classical Problems/Contemporary Issues*, ed. by Brian Beakley and Peter Ludlow. 2nd ed. Cambridge MA: MIT Press.
- Helen Beebe, Christopher Hitchcock, and Huw Price. 2017. *Making a Difference: Essays on the Philosophy of Causation*, ed. by Helen Beebe, Christopher Hitchcock, and Huw Price. Oxford: Oxford University Press. URL: <http://doi.org/10.1093/oso/9780198746911.001.0001>.
- Karen Bennett. 2003. “Why the Exclusion Problem Seems Intractable, and How, Just Maybe, To Tract It”, in *Noûs*, Vol. 37, No. 3, Sept. 2003, pp. 471–497. URL: <http://doi.org/10.1111/1468-0068.00447>.
- . 2008. “Exclusion Again”, in *Being Reduced: New Essays on Reduction, Explanation and Causation*, ed. by Jesper Kallestrup and Jakob Hohwy, Oxford University Press, Oxford, pp. 280–305. URL: <http://doi.org/10.1093/acprof:oso/9780199211531.003.0015>.
- Simon Blackburn. 1991. “Losing Your Mind: Physics, Identity, and Folk Burglar Prevention”, in *The Future of Folk Psychology: Intentionality and Cognitive Science*, ed. by John D. Greenwood, Cambridge University Press, Cambridge, pp. 196–225. Reprinted in Blackburn (1993, pp. 229–254).
- . 1993. *Essays in Quasi-Realism*, Oxford: Oxford University Press.
- Thomas Blanchard and Jonathan Schaffer. 2017. “Cause without Default”, in *Making a Difference: Essays on the Philosophy of Causation*, ed. by Helen Beebe, Christopher Hitchcock, and Huw Price, Oxford University Press, Oxford, pp. 175–214. URL: <http://doi.org/10.1093/oso/9780198746911.001.0001>.

- Ned Block. 1980. *Readings in Philosophy of Psychology*, ed. by Ned Block. Vol. 1. Cambridge MA: Harvard University Press.
- \_\_\_\_\_. 1981. *Readings in Philosophy of Psychology*, ed. by Ned Block. Vol. 2. Cambridge MA: Harvard University Press.
- \_\_\_\_\_. 1990. “Can the Mind Change the World?”, in *Meaning and Method: Essays in Honor of Hilary Putnam*, ed. by George Boolos, Cambridge University Press, Cambridge, pp. 137–170.
- \_\_\_\_\_. 2003. “Do Causal Powers Drain Away?”, in *Philosophy and Phenomenological Research*, Vol. 67, No. 1, July 2003, pp. 133–150. URL: <http://doi.org/10.1111/j.1933-1592.2003.tb00029.x>.
- Thomas D. Bontly. 2002. “The Supervenience Argument Generalizes”, in *Philosophical Studies*, Vol. 109, No. 1, May 2002, pp. 75–96. URL: <http://doi.org/10.1023/A:1015786809364>.
- Andrei A. Buckareff. 2011. “Intralevel Mental Causation”, in *Frontiers of Philosophy in China*, Vol. 6, No. 3, Sept. 2011, pp. 402–425. URL: <http://doi.org/10.2307/44259314>.
- \_\_\_\_\_. 2012. “An Action-Theoretic Problem for Intralevel Mental Causation”, in *Philosophical Issues*, Vol. 22, pp. 89–105. URL: <http://doi.org/10.2307/41683062>.
- Tyler Burge. 1993. “Mind-Body Causation and Explanatory Practice”, in *Mental Causation*, ed. by John Heil and Alfred Mele, Oxford University Press, Oxford, pp. 97–120. Reprinted with postscript in Burge (2007a, pp. 344–362).
- \_\_\_\_\_. 2007a. *Foundations of Mind*, Vol. 2. Philosophical Essays. Oxford: Oxford University Press.
- \_\_\_\_\_. 2007b. “Postscript: Mind-Body Causation and Explanatory Practice”, in *Foundations of Mind*. Vol. 2, Philosophical Essays, Oxford University Press, Oxford, pp. 363–382.
- John Campbell. 2007. “An Interventionist Approach to Causation in Psychology”, in *Causal Learning: Psychology, Philosophy, Computation*, ed. by Alison Gopnik and Laura Schulz, Oxford University Press, New York, pp. 58–66.

- John Campbell. 2008. "Causation in Psychiatry", in *Philosophical Issues in Psychiatry: Explanation, Phenomenology and Nosology*, ed. by Kenneth S. Kendler and Josef Parnas, Johns Hopkins University Press, Baltimore MD, pp. 196–215.
- Nancy Cartwright. 1983. *How the Laws of Physics Lie*, Oxford: Oxford University Press. URL: <http://doi.org/10.1093/0198247044.001.0001>.
- \_\_\_\_\_. 1994. "Fundamentalism vs. the Patchwork of Laws", in *Proceedings of the Aristotelian Society*, Vol. 94, No. 1, June 1994, pp. 279–292. Reprinted in Papineau (1996, pp. 314–326) and Cartwright (1999, pp. 23–34). URL: <https://doi.org/10.1093/aristotelian/94.1.279>.
- \_\_\_\_\_. 1999. *The Dappled World: A Study of the Boundaries of Science*, Cambridge: Cambridge University Press. URL: <https://doi.org/10.1017/CBO9781139167093>.
- Andy Clark and Peter Millican. 1999. *Connectionism, Concepts, and Folk Psychology: The Legacy of Alan Turing*, ed. by Andy Clark and Peter Millican. Vol. 2. Oxford: Oxford University Press.
- Tim Crane. 2008. "Causation and Determinable Properties: On the Efficacy of Colour, Shape, and Size", in *Being Reduced: New Essays on Reduction, Explanation and Causation*, ed. by Jesper Kallestrup and Jakob Hohwy, Oxford University Press, Oxford, pp. 176–195. URL: <http://doi.org/10.1093/acprof:oso/9780199211531.003.0011>.
- Donald Davidson. 1963. "Actions, Reasons, and Causes", in *The Journal of Philosophy*, Vol. 60, No. 23, Nov. 1963, pp. 685–700. Reprinted in Davidson (2001, pp. 3–19). URL: <http://doi.org/10.2307/2023177>.
- \_\_\_\_\_. 1967. "Causal Relations", in *The Journal of Philosophy*, Vol. 64, No. 21, Nov. 1967, pp. 691–703. Reprinted in Davidson (2001, pp. 149–162). URL: <https://doi.org/10.2307/2023853>.
- \_\_\_\_\_. 1970. "Mental Events", in *Experience and Theory*, ed. by Lawrence Foster and Joe William Swanson, Univer-

- sity of Massachusetts Press, Amherst MA, pp. 79–101. Reprinted in Block (1980, pp. 107–119), Moser and Trout (1995, pp. 11–126) and Davidson (2001, pp. 207–227). URL: <http://doi.org/10.1093/0199246270.003.0011>.
- \_\_\_\_\_. 1995. “Laws and Cause”, in *Dialectica*, Vol. 49, No. 2-4, June 1995, pp. 263–279. URL: <https://doi.org/10.1111/j.1746-8361.1995.tb00165.x>.
- \_\_\_\_\_. 2001. *Essays on Actions and Events*, 2nd ed. Oxford: Oxford University Press. URL: <http://doi.org/10.1093/0199246270.001.0001>.
- Frederick Eberhardt. 2014. “Direct Causes and the Trouble with Soft Interventions”, in *Erkenntnis*, Vol. 79, No. 4, Apr. 2014, pp. 755–777. URL: <https://doi.org/10.1007/s10670-013-9552-2>.
- Frederick Eberhardt and Richard Scheines. 2007. “Interventions and Causal Inference”, in *Philosophy of Science*, Vol. 74, No. 5, Dec. 2007, pp. 981–995. URL: <https://doi.org/10.1086/525638>.
- Ellery Eells. 1991. *Probabilistic Causality*, Cambridge: Cambridge University Press. URL: <https://doi.org/10.1017/CBO9780511570667>.
- Markus I. Eronen and Daniel S. Brooks. 2014. “Interventionism and Supervenience: A New Problem and Provisional Solution”, in *International Studies in the Philosophy of Science*, Vol. 28, No. 2, pp. 185–202. URL: <https://doi.org/10.1080/02698595.2014.932529>.
- Herbert Feigl and May Brodbeck. 1953. *Readings in the Philosophy of Science*, ed. by Herbert Feigl and May Brodbeck. New York: Appleton-Century-Crofts.
- Luke Fenton-Glynn. 2021. *Causation*, Cambridge: Cambridge University Press. URL: <https://doi.org/10.1017/9781108588300>.
- Hartry Field. 1978. “Mental Representation”, in *Erkenntnis*, Vol. 13, No. 1, Jan. 1978, pp. 9–61. Reprinted with postscript in Block (1981, pp. 78–116), in Stich and Warfield (1994, pp. 34–78), and with

- an updated postscript in Field (2001, pp. 30–67). URL: <http://doi.org/10.1007/BF00160888>.
- \_\_\_\_\_. 2001. *Truth and the Absence of Fact*, Oxford: Oxford University Press. URL: <http://doi.org/10.1093/0199242895.001.0001>.
- \_\_\_\_\_. 2003. “Causation in a Physical World”, in *The Oxford Handbook of Metaphysics*, ed. by Michael J. Loux and Dean W. Zimmerman, Oxford University Press, Oxford, pp. 435–460. URL: <http://doi.org/10.1093/oxfordhb/9780199284221.003.0015>.
- J. Dmitri Gallow. 2022. “The Metaphysics of Causation”, in *Stanford Encyclopedia of Philosophy*, ed. by Edward N. Zalta, Stanford University, Stanford. URL: <https://plato.stanford.edu/archives/fall2022/entries/causation-metaphysics/>.
- Alexander Gebharder. 2017a. “Causal Exclusion and Causal Bayes Nets”, in *Philosophy and Phenomenological Research*, Vol. 95, No. 2, Sept. 2017, pp. 353–375. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/phpr.12247>.
- \_\_\_\_\_. 2017b. “Causal Exclusion Without Physical Completeness and No Overdetermination”, in *Abstracta*, Vol. 10, pp. 3–14.
- John Gibbons. 2006. “Mental Causation without Downward Causation”, in *The Philosophical Review*, Vol. 115, No. 1, Jan. 2006, pp. 79–103. URL: <http://doi.org/10.1215/00318108-115-1-79>.
- Alvin I. Goldman. 1969. “The Compatibility of Mechanism and Purpose”, in *The Philosophical Review*, Vol. 78, No. 4, Oct. 1969, pp. 468–482. URL: <https://doi.org/10.2307/2184199>.
- \_\_\_\_\_. 1970. *A Theory of Human Action*, Englewood Cliffs NJ: Prentice Hall.
- Joseph Y. Halpern. 2016. “Appropriate Causal Models and the Stability of Causation”, in *The Review of Symbolic Logic*, Vol. 9, No. 1, Mar. 2016, pp. 76–102. URL: <http://doi.org/10.1017/S1755020315000246>.
- Joseph Y. Halpern and Christopher Hitchcock. 2010. “Actual Causation and the Art of Modeling”, in *Heuristics, Probability and*

- Causality: A Tribute to Judea Pearl*, ed. by Rina Dechter, Hector Geffner, and Joseph Y. Halpern, College Publications, London, pp. 383–406.
- Joseph Y. Halpern and Judea Pearl. 2001. “Causes and Explanations: A Structural-Model Approach. Part I: Causes”, in *Uncertainty in Artificial Intelligence: Proceedings of the Seventeenth Conference*, ed. by Jack Breese and Daphne Koller, Morgan Kaufmann, San Francisco, pp. 194–202. Revised version published as Halpern and Pearl (2005). URL: <http://arxiv.org/abs/cs/0011012v2>.
- . 2005. “Causes and Explanations: A Structural-Model Approach. Part I: Causes”, in *British Journal for the Philosophy of Science*, Vol. 56, No. 4, Dec. 2005, pp. 843–887. Revised version of Halpern and Pearl (2001). URL: <http://doi.org/10.1093/bjps/axi147>.
- John Haugeland. 1985. *Artificial Intelligence: The Very Idea*, Cambridge MA: MIT Press.
- Daniel M. Hausman. 2005. “Causal Relata: Tokens, Types, or Variables?”, in *Erkenntnis*, Vol. 63, No. 1, July 2005, pp. 33–54. URL: <http://doi.org/10.1007/s10670-005-0562-6>.
- James J. Heckman. 2005. “The Scientific Model of Causality”, in *Sociological Methodology*, Vol. 35, No. 1, Aug. 2005, pp. 1–98. URL: <http://doi.org/10.1111/j.0081-1750.2006.00163.x>.
- David K. Henderson. 1994. “Accounting for Macro-Level Causation”, in *Synthese*, Vol. 101, No. 2, Nov. 1994, pp. 129–156. URL: <https://doi.org/10.1007/BF01064014>.
- Robin Hendry. 2006. “Is There Downward Causation in Chemistry?”, in *Philosophy of Chemistry: Synthesis of a New Discipline*, ed. by Davis Baird, Eric Scerri, and Lee McIntyre. Vol. 242, Boston Studies in the Philosophy of Science, Springer, Dordrecht, pp. 173–189. URL: [https://doi.org/10.1007/1-4020-3261-7\\_9](https://doi.org/10.1007/1-4020-3261-7_9).
- . 2010a. “Emergence vs. Reduction in Chemistry”, in *Emergence in Mind*, ed. by Cynthia Macdonald and Graham Macdonald,



- Oxford University Press, Oxford, pp. 205–221. URL: <http://doi.org/10.1093/acprof:oso/9780199583621.003.0014>.
- Robin Hendry. 2010b. “Ontological Reduction and Molecular Structure”, in *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, Vol. 41, No. 2, May 2010, pp. 183–191. URL: <https://doi.org/10.1016/j.shpsb.2010.03.005>.
- . 2017. “Prospects for Strong Emergence in Chemistry”, in *Philosophical and Scientific Perspectives on Downward Causation*, Routledge, New York, pp. 146–163. URL: <https://doi.org/10.4324/9781315638577-9>.
- Christopher Hitchcock. 2001. “The Intransitivity of Causation Revealed in Equations and Graphs”, in *The Journal of Philosophy*, Vol. 98, No. 6, June 2001, pp. 273–299. URL: <https://doi.org/10.2307/2678432>.
- . 2004. “Routes, Processes, and Chance-Lowering Causes”, in *Cause and Chance: Causation in an Indeterministic World*, ed. by Phil Dowe and Paul Noordhof, Routledge, London, pp. 138–152. URL: <https://doi.org/10.4324/9780203494660>.
- . 2007. “Prevention, Preemption, and the Principle of Sufficient Reason”, in *Philosophical Review*, Vol. 116, No. 4, Oct. 2007, pp. 495–532. URL: <http://doi.org/10.1215/00318108-2007-012>.
- . 2009. “Causal Modelling”, in *The Oxford Handbook of Causation*, ed. by Helen Beebe, Christopher Hitchcock, and Peter Menzies, Oxford University Press, Oxford, pp. 299–314. URL: <http://doi.org/10.1093/oxfordhb/9780199279739.003.0015>.
- . 2012. “Events and Times: A Case Study in Means-Ends Metaphysics”, in *Philosophical Studies*, Vol. 160, No. 1, Aug. 2012, pp. 79–96. URL: <http://doi.org/10.1007/s11098-012-9909-4>.
- . 2023. “Causal Models”, in *Stanford Encyclopedia of Philosophy*, ed. by Edward N. Zalta, Stanford University, Stanford. URL:

<https://plato.stanford.edu/archives/spr2023/entries/causal-models/>.

- Vera Hoffman-Kolss. 2014. “Interventionism and Higher-Level Causation”, in *International Studies in the Philosophy of Science*, Vol. 28, No. 1, pp. 49–64. URL: <http://doi.org/10.1080/02698595.2014.915653>.
- Terence Horgan. 1997. “Kim on Mental Causation and Causal Exclusion”, in *Noûs*, Vol. 31, Supplement: Philosophical Perspectives, 11, Mind, Causation, and World 1997, pp. 165–184. URL: <https://doi.org/10.1111/0029-4624.31.s11.8>.
- Paul W. Humphreys. 1989. *The Chances of Explanation*, Princeton: Princeton University Press. URL: <https://doi.org/10.1515/9781400860760>.
- Jenann Ismael. 2009. “Probability in Deterministic Physics”, in *Journal of Philosophy*, Vol. 106, No. 2, Feb. 2009, pp. 89–108. URL: <https://doi.org/10.5840/jphil2009106214>.
- . 2011. “A Modest Proposal about Chance”, in *Journal of Philosophy*, Vol. 108, No. 8, Aug. 2011, pp. 416–442. URL: <https://doi.org/10.5840/jphil2011108822>.
- Frank Jackson and Philip Pettit. 1988. “Functionalism and Broad Content”, in *Mind*, Vol. 97, No. 387, July 1988, pp. 381–400. Reprinted in Pessin and Goldberg (1996, pp. 219–237) and Jackson, Pettit, and Smith (2004, pp. 95–118). URL: <http://doi.org/10.1093/mind/XCVII.387.381>.
- . 1990a. “Causation in the Philosophy of Mind”, in *Philosophy and Phenomenological Research*, Vol. 50, Supplement, Autumn 1990, pp. 195–214. Reprinted with postscript in Clark and Millikan (1999, pp. 75–100) and Jackson, Pettit, and Smith (2004, pp. 45–68). URL: <http://doi.org/10.2307/2108039>.
- . 1990b. “Program Explanation: A General Perspective”, in *Analysis*, Vol. 50, No. 2, Mar. 1990, pp. 107–117. Reprinted in

- Jackson, Pettit, and Smith (2004, pp. 119–130). URL: <http://doi.org/10.2307/3328853>.
- Frank Jackson, Philip Pettit, and Michael Smith. 2004. *Mind, Morality, and Explanation: Selected Collaborations*, Oxford: Oxford University Press.
- Alex Kaiserman. 2020. “Interventionism and Mental Surgery”, in *Erkenntnis*, Vol. 85, No. 4, Aug. 2020, pp. 919–935. URL: <https://doi.org/10.1007/s10670-018-0059-8>.
- Jesper Kallestrup and Jakob Hohwy. 2008. *Being Reduced: New Essays on Reduction, Explanation and Causation*, ed. by Jesper Kallestrup and Jakob Hohwy. Oxford: Oxford University Press. URL: <http://doi.org/10.1093/acprof:oso/9780199211531.001.0001>.
- Kenneth S. Kendler and John Campbell. 2009. “Interventionist Causal Models in Psychiatry: Repositioning the Mind-Body Problem”, in *Psychological Medicine*, Vol. 39, No. 6, June 2009, pp. 881–887. URL: <http://doi.org/10.1017/S0033291708004467>.
- Jaegwon Kim. 1973. “Causes and Counterfactuals”, in *The Journal of Philosophy*, Vol. 70, No. 17, Oct. 1973, pp. 570–572. Reprinted in Tooley (1999, pp. 190–193). URL: <https://doi.org/10.2307/2025312>.
- . 1989. “Mechanism, Purpose, and Explanatory Exclusion”, in *Philosophical Perspectives*, Vol. 3, Philosophy of Mind and Action Theory 1989, pp. 77–108. Reprinted in Kim (1993, pp. 237–264). URL: <http://doi.org/10.2307/2214264>.
- . 1993. *Supervenience and Mind: Selected Philosophical Essays*, Cambridge: Cambridge University Press. URL: <https://doi.org/10.1017/CBO9780511625220>.
- . 1997. “Does the Problem of Mental Causation Generalize?”, in *Proceedings of the Aristotelian Society*, Vol. 97, No. 3, pp. 281–297. URL: <http://doi.org/10.1111/1467-9264.00017>.

- Jaegwon Kim. 1998. *Mind in a Physical World: An Essay on the Mind-Body Problem and Mental Causation*, Cambridge MA: MIT Press. URL: <https://doi.org/10.7551/mitpress/4629.001.0001>.
- . 2003. “Blocking Causal Drainage and Other Maintenance Chores with Mental Causation”, in *Philosophy and Phenomenological Research*, Vol. 67, No. 1, July 2003, pp. 151–176. URL: <http://doi.org/10.1111/j.1933-1592.2003.tb00030.x>.
- . 2005. *Physicalism, Or Something Near Enough*, Princeton: Princeton University Press. URL: <https://doi.org/10.1515/9781400840847>.
- Kevin B. Korb, Lucas R. Hope, Ann E. Nicholson, and Karl Axnick. 2004. “Varieties of Causal Intervention”. In *PRICAI 2004: Trends in Artificial Intelligence*. Ed. by Chengqi Zhang, Hans W. Guesgen, and Wai-Kiang Yeap. Berlin: Springer, pp. 322–331. URL: [https://doi.org/10.1007/978-3-540-28633-2\\_35](https://doi.org/10.1007/978-3-540-28633-2_35).
- Christian List and Peter Menzies. 2009. “Nonreductive Physicalism and the Limits of the Exclusion Principle”, in *Journal of Philosophy*, Vol. 106, No. 9, Sept. 2009, pp. 475–502. URL: <http://doi.org/10.2307/20620197>.
- Barry Loewer. 2008. “Why There Is Anything Except Physics”, in *Being Reduced: New Essays on Reduction, Explanation and Causation*, ed. by Jesper Kallestrup and Jakob Hohwy, Oxford University Press, Oxford, pp. 149–163. URL: <http://doi.org/10.1093/acprof:oso/9780199211531.001.0001>.
- . 2009. “Why is There Anything Except Physics?”, in *Synthese*, Vol. 170, No. 2, Sept. 2009, pp. 217–233. URL: <http://doi.org/10.1007/s11229-009-9580-2>.
- . 2015. “Mental Causation: The Free Lunch”, in *Qualia and Mental Causation in a Physical World: Themes from the Philosophy of Jaegwon Kim*, ed. by David Sosa, Terence Horgan, and Marcelo Sabatés, Cambridge University Press, Cambridge, pp. 40–63. URL: <https://doi.org/10.1017/CBO9781139939539.004>.

- E. J. Lowe. 2000. "Causal Closure Principles and Emergentism", in *Philosophy*, Vol. 75, No. 294, Oct. 2000, pp. 571–585. Reprinted in Lowe (2008, pp. 41–57). URL: <https://doi.org/10.1017/S003181910000067X>.
- . 2003. "Physical Causal Closure and the Invisibility of Mental Causation", in *Physicalism and Mental Causation: The Metaphysics of Mind and Action*, ed. by Sven Walter and Heinz-Dieter Heckmann, Imprint Academic, Exeter. Reprinted in Lowe (2008, pp. 58–78).
- . 2008. *Personal Agency: The Metaphysics of Mind and Action*, Oxford: Oxford University Press. URL: <https://doi.org/10.1093/acprof:oso/9780199217144.001.0001>.
- William G. Lycan and Jesse J. Prinz. 2008. *Mind and Cognition: An Anthology*, ed. by William G. Lycan and Jesse J. Prinz. 3rd ed. Malden MA: Blackwell.
- Cynthia Macdonald and Graham Macdonald. 2010. *Emergence in Mind*, ed. by Cynthia Macdonald and Graham Macdonald. Oxford: Oxford University Press. URL: <http://doi.org/10.1093/acprof:oso/9780199583621.001.0001>.
- John L. Mackie. 1974. *The Cement of the Universe: A Study of Causation*, Oxford: Oxford University Press. URL: <http://doi.org/10.1093/0198246420.001.0001>.
- Norman Malcolm. 1968. "The Conceivability of Mechanism", in *The Philosophical Review*, Vol. 77, No. 1, Jan. 1968, pp. 45–72. URL: <https://doi.org/10.2307/2183182>.
- Florian Markowetz, Steffen Grossmann, and Rainer Spang. 2005. "Probabilistic Soft Interventions in Conditional Gaussian Networks". In *IN 10TH AI/STATS*. Ed. by Robert Cowell and Zoubin Ghahramani. Savannah Hotel, Barbados: The Society for Artificial Intelligence and Statistics, Jan. 2005. URL: <http://www.gatsby.ucl.ac.uk/aistats/>.

- Ausonio Marras. 1998. “Kim’s Principle of Explanatory Exclusion”, in *Australasian Journal of Philosophy*, Vol. 76, No. 3, pp. 439–451. URL: <https://doi.org/10.1080/00048409812348551>.
- Michael McDermott. 1995. “Redundant Causation”, in *British Journal for the Philosophy of Science*, Vol. 46, No. 4, pp. 523–544. URL: <http://doi.org/10.1093/bjps/46.4.52>.
- . 2002. “Causation: Influence versus Sufficiency”, in *The Journal of Philosophy*, Vol. 99, No. 2, Feb. 2002, pp. 84–101. URL: <https://doi.org/10.5840/jphil200299219>.
- Jenn McDonald. 2023. “Causal Models and Contrastivism”. Unpublished manuscript.
- . Forthcoming. “Essential Structure for Apt Causal Models”, in *Australasian Journal of Philosophy*.
- D. H. Mellor. 1995. *The Facts of Causation*, London: Routledge.
- Andrew Melnyk. 2003. *A Physicalist Manifesto: Thoroughly Modern Materialism*, Cambridge: Cambridge University Press. URL: <https://doi.org/10.1017/CBO9780511498817>.
- Peter Menzies. 2003. “The Causal Efficacy of Mental States”, in *Physicalism and Mental Causation: The Metaphysics of Mind and Action*, ed. by Sven Walter and Heinz-Dieter Heckmann, Imprint Academic, Exeter, pp. 195–224. Reprinted in Monnoyer (2004).
- . 2010. “Reasons and Causes Revisited”, in *Naturalism and Normativity*, ed. by Mario De Caro and David Macarthur, Columbia University Press, New York, pp. 142–170.
- Peter Menzies and Christian List. 2010. “The Causal Autonomy of the Special Sciences”, in *Emergence in Mind*, ed. by Cynthia Macdonald and Graham Macdonald, Oxford University Press, Oxford, pp. 108–128. URL: <http://doi.org/10.1093/acprof:oso/9780199583621.003.0008>.
- Jean-Maurice Monnoyer. 2004. *La Structure du Monde: Objets, Propriétés, États de Choses*, ed. by Jean-Maurice Monnoyer. Paris: Vrin.

- Paul K. Moser and J. D. Trout. 1995. *Contemporary Materialism: A Reader*, ed. by Paul K. Moser and J. D. Trout. London: Routledge. URL: <https://doi.org/10.4324/9780203427262>.
- Paul Noordhof. 1997. "Making the Change: the Functionalist's Way", in *British Journal for the Philosophy of Science*, Vol. 48, No. 2, pp. 233–250. URL: <http://doi.org/10.1093/bjps/48.2.233>.
- David Papineau. 1996. *The Philosophy of Science*, ed. by David Papineau. Oxford: Oxford University Press.
- . 2001. "The Rise of Physicalism", in *Physicalism and its Discontents*, ed. by Carl Gillett and Barry Loewer, Cambridge University Press, Cambridge, pp. 3–36. Reprinted in Stone and Wolff (2000, pp. 174–208). URL: <https://doi.org/10.1017/CBO9780511570797.002>.
- . 2002. *Thinking about Consciousness*, Oxford: Oxford University Press. URL: <http://doi.org/10.1093/0199243824.001.0001>.
- Sarah Patterson. 2005. "Epiphenomenalism and Occasionalism: Problems of Mental Causation, Old and New", in *History of Philosophy Quarterly*, Vol. 22, No. 3, July 2005, pp. 239–257.
- Christopher Peacocke. 1993a. "Externalist Explanation", in *Proceedings of the Aristotelian Society*, Vol. 93, No. 3, June 1993, pp. 203–230. URL: <https://doi.org/10.1093/aristotelian/93.1.203>.
- . 1993b. "Review of *The Imagery Debate*", in *Philosophy of Science*, Vol. 60, No. 4, Dec. 1993, pp. 675–677. URL: <https://doi.org/10.1086/289774>.
- Judea Pearl. 2000. *Causality*, Cambridge: Cambridge University Press.
- Derk Pereboom and Hilary Kornblith. 1991. "The Metaphysics of Irreducibility", in *Philosophical Studies*, Vol. 63, No. 2, Aug. 1991, pp. 125–145. URL: <http://doi.org/10.1007/BF00381684>.
- Andrew Pessin and Sanford Goldberg. 1996. *The Twin Earth Chronicles: Twenty Years of Reflection on Hilary Putnam's "The Meaning of 'Meaning'"*, ed. by Andrew Pessin and Sanford Goldberg. Armonk NY: M. E. Sharpe.

- Thomas W. Polger, Lawrence A. Shapiro, and Reuben Stern. 2018. “In Defense of Interventionist Solutions to Exclusion”, in *Studies in History and Philosophy of Science Part A*, Vol. 68, Apr. 2018, pp. 51–57. URL: <https://doi.org/10.1016/j.shpsa.2018.01.012>.
- Huw Price and Richard Corry. 2007. *Causation, Physics and the Constitution of Reality: Russell’s Republic Revisited*, ed. by Huw Price and Richard Corry. Oxford: Oxford University Press.
- Hilary Putnam. 1975. “Philosophy and our Mental Life”, in *Mind, Language and Reality*. Vol. II, Philosophical Papers, Cambridge University Press, Cambridge, pp. 291–303. Reprinted in Block (1980, pp. 134–143) and Moser and Trout (1995, pp. 127–137). URL: <https://doi.org/10.1017/CBO9780511625251.016>.
- Panu Raatikainen. 2010. “Causation, Exclusion, and the Special Sciences”, in *Erkenntnis*, Vol. 73, No. 3, Nov. 2010, pp. 349–363. URL: <https://doi.org/10.1007/s10670-010-9236-0>.
- Michael Rescorla. 2012. “Are Computational Transitions Sensitive to Semantics?”, in *Australasian Journal of Philosophy*, Vol. 90, No. 4, Dec. 2012, pp. 703–721. URL: <https://doi.org/10.1080/00048402.2011.615333>.
- . 2014. “The Causal Relevance of Content to Computation”, in *Philosophy and Phenomenological Research*, Vol. 88, No. 1, Jan. 2014, pp. 173–208. URL: <http://doi.org/10.1111/j.1933-1592.2012.00619.x>.
- . 2018. “An Interventionist Approach to Psychological Explanation”, in *Synthese*, Vol. 195, No. 5, May 2018, pp. 1909–1940. URL: <https://doi.org/10.1007/s11229-017-1553-2>.
- David Robb, John Heil, and Sophie Gibb. 2023. “Mental Causation”, in *Stanford Encyclopedia of Philosophy*, ed. by Edward N. Zalta, Stanford University, Stanford. URL: <https://plato.stanford.edu/archives/spr2023/entries/mental-causation/>.



- Alexander Rosenberg and Robert Arp. 2010. *Philosophy of Biology: An Anthology*, ed. by Alexander Rosenberg and Robert Arp. Malden MA: Blackwell.
- Bertrand Russell. 1913. “On the Notion of Cause”, in *Proceedings of the Aristotelian Society*, Vol. 13, No. 1, June 1913, pp. 1–26. Reprinted in Russell (1918, pp. 142–164) and Russell (2003, pp. 163–182). Simplified version published as Russell (1914). URL: <https://doi.org/10.1093/aristotelian/13.1.1>.
- . 1914. “On the Notion of Cause, with Applications to the Free Will Problem”, in *Our Knowledge of the External World: As a Field for Scientific Method in Philosophy*, Open Court, Chicago, pp. 214–246. Reprinted in Feigl and Brodbeck (1953, pp. 387–407).
- . 1918. *Mysticism and Logic*, London: Longmans Green.
- . 2003. *Russell on Metaphysics: Selections from the Writings of Bertrand Russell*, ed. by Stephen Mumford. London: Routledge.
- Jonathan Schaffer. 2016. “The Metaphysics of Causation”, in *Stanford Encyclopedia of Philosophy*, ed. by Edward N. Zalta, Stanford University, Stanford. URL: <https://plato.stanford.edu/archives/fall2016/entries/causation-metaphysics/>.
- Markus E. Schlosser. 2009. “Non-Reductive Physicalism, Mental Causation, and the Nature of Action”, in *Reduction: Between the Mind and the Brain*, ed. by Alexander Hieke and Hannes Leitgeb, Ontos Verlag, Frankfurt, pp. 73–90. URL: <https://doi.org/10.1515/9783110328851.73>.
- Gabriel Segal. 2004. “Reference, Causal Powers, Externalist Intuitions, and Unicorns”, in *The Externalist Challenge*, ed. by Richard Schantz, Walter de Gruyter, Berlin, pp. 329–344. Reprinted in Lycan and Prinz (2008). URL: <https://doi.org/10.1515/9783110915273.329>.

- Gabriel Segal. 2009. “The Causal Inefficacy of Content”, in *Mind and Language*, Vol. 24, No. 1, Feb. 2009, pp. 80–102. URL: <http://doi.org/10.1111/j.1468-0017.2008.01354.x>.
- Gabriel Segal and Elliott Sober. 1991. “The Causal Efficacy of Content”, in *Philosophical Studies*, Vol. 63, No. 1, July 1991, pp. 1–30. URL: <http://doi.org/10.1007/BF00375995>.
- Lawrence A. Shapiro. 2010. “Lessons from Causal Exclusion”, in *Philosophy and Phenomenological Research*, Vol. 81, No. 3, Nov. 2010, pp. 594–604. URL: <http://doi.org/10.1111/j.1933-1592.2010.00382.x>.
- Lawrence A. Shapiro and Elliott Sober. 2007. “Epiphenomenalism: The Dos and the Don’ts”, in *Thinking about Causes: From Greek Philosophy to Modern Physics*, ed. by Peter Machamer and Gereon Wolters, Pittsburgh-Konstanz Series in the Philosophy and History of Science, University of Pittsburgh Press, Pittsburgh, pp. 235–264.
- Lawrence Sklar. 2003. “Dappled Theories in a Uniform World”, in *Philosophy of Science*, Vol. 70, No. 2, Apr. 2003, pp. 424–441. URL: <http://doi.org/10.1086/375476>.
- Elliott Sober. 1999a. “Physicalism from a Probabilistic Point of View”, in *Philosophical Studies*, Vol. 95, No. 1-2, Aug. 1999, pp. 135–174. Excerpts reprinted in Beakley and Ludlow (2006, pp. 461–467). URL: <http://doi.org/10.1023/A:1004519608950>.
- . 1999b. “The Multiple Realizability Argument Against Reductionism”, in *Philosophy of Science*, Vol. 66, No. 4, Dec. 1999, pp. 542–564. Reprinted in Sober (2006, pp. 301–324) and Rosenberg and Arp (2010, pp. 235–250). URL: <https://doi.org/10.1086/392754>.
- . 2006. *Conceptual Issues in Evolutionary Biology: An Anthology*, ed. by Elliott Sober. 3rd ed. Cambridge MA: MIT Press.

- Georgie Statham. 2018. “Woodward and Variable Relativity”, in *Philosophical Studies*, Vol. 175, No. 4, pp. 885–902. URL: <https://doi.org/10.1007/s11098-017-0897-2>.
- Reuben Stern and Benjamin Eva. 2023. “Antireductionist Interventionism”, in *The British Journal for the Philosophy of Science*, Vol. 74, No. 1, Mar. 2023, pp. 241–267. URL: <https://doi.org/10.1086/714792>.
- Stephen Stich. 1983. *From Folk Psychology to Cognitive Science: The Case Against Belief*, Cambridge MA: MIT Press.
- Stephen Stich and Ted A. Warfield. 1994. *Mental Representation: A Reader*, ed. by Stephen Stich and Ted A. Warfield. Oxford: Basil Blackwell.
- M. W. F. Stone and Jonathan Wolff. 2000. *The Proper Ambition of Science*, ed. by M. W. F. Stone and Jonathan Wolff. Vol. 2. London Studies in the History of Philosophy. London: Routledge.
- Michael Strevens. 2007. “Review of Woodward, *Making Things Happen*”, in *Philosophy and Phenomenological Research*, Vol. 74, No. 1, Jan. 2007, pp. 233–249. URL: <http://doi.org/10.1111/j.1933-1592.2007.00012.x>.
- Amie Thomasson. 1998. “A Nonreductivist Solution to Mental Causation”, in *Philosophical Studies*, Vol. 89, No. 2-3, Mar. 1998, pp. 181–195. URL: <http://doi.org/10.1023/A:1004280812099>.
- Michael Tooley. 1999. *Laws of Nature, Causation, and Supervenience*, ed. by Michael Tooley. Vol. 1. Metaphysics. New York: Garland.
- Michael Tye. 1991. *The Imagery Debate*, Cambridge MA: MIT Press.
- Sven Walter and Heinz-Dieter Heckmann. 2003. *Physicalism and Mental Causation: The Metaphysics of Mind and Action*, ed. by Sven Walter and Heinz-Dieter Heckmann. Exeter: Imprint Academic.
- Brad Weslake. 2010. “Explanatory Depth”, in *Philosophy of Science*, Vol. 77, No. 2, Apr. 2010, pp. 273–294. URL: <http://doi.org/10.1086/651316>.

- Brad Weslake. 2017. "Difference-Making, Closure and Exclusion", in *Making a Difference: Essays on the Philosophy of Causation*, ed. by Helen Beebe, Christopher Hitchcock, and Huw Price, Oxford University Press, Oxford, pp. 215–231. URL: <http://doi.org/10.1093/oso/9780198746911.003.0011>.
- \_\_\_\_\_. unpublished. "A Partial Theory of Actual Causation".
- James Woodward. 2003. *Making Things Happen: A Theory of Causal Explanation*, New York: Oxford University Press. URL: <http://doi.org/10.1093/0195155270.001.0001>.
- \_\_\_\_\_. 2008a. "Mental Causation and Neural Mechanisms", in *Being Reduced: New Essays on Reduction, Explanation and Causation*, ed. by Jesper Kallestrup and Jakob Hohwy, Oxford University Press, Oxford, pp. 218–262. URL: <http://doi.org/10.1093/acprof:oso/9780199211531.003.0013>.
- \_\_\_\_\_. 2008b. "Response to Strevens", in *Philosophy and Phenomenological Research*, Vol. 77, No. 1, Aug. 2008, pp. 193–212. URL: <http://doi.org/10.1111/j.1933-1592.2008.00181.x>.
- \_\_\_\_\_. 2015a. "Interventionism and Causal Exclusion", in *Philosophy and Phenomenological Research*, Vol. 91, No. 2, Sept. 2015, pp. 303–347. URL: <http://doi.org/10.1111/phpr.12095>.
- \_\_\_\_\_. 2015b. "Methodology, Ontology, and Interventionism", in *Synthese*, Vol. 192, No. 11, Nov. 2015, pp. 3577–3599. URL: <https://doi.org/10.1007/s11229-014-0479-1>.
- \_\_\_\_\_. 2016. "The Problem of Variable Choice", in *Synthese*, Vol. 193, No. 4, Apr. 2016, pp. 1047–1072. URL: <https://doi.org/10.1007/s11229-015-0810-5>.
- \_\_\_\_\_. 2017. "Intervening in the Exclusion Argument", in *Making a Difference: Essays on the Philosophy of Causation*, ed. by Helen Beebe, Christopher Hitchcock, and Huw Price, Oxford University Press, Oxford, pp. 251–268. URL: <http://doi.org/10.1093/oso/9780198746911.003.0013>.

- James Woodward. 2020. “Causal Complexity, Conditional Independence, and Downward Causation”, in *Philosophy of Science*, Vol. 87, No. 5, Dec. 2020, pp. 857–867. URL: <https://doi.org/10.1086/710631>.
- . 2021a. “Downward Causation and Levels”, in *Levels of Organization in the Biological Sciences*, ed. by Daniel S. Brooks, James DiFrisco, and William C. Wimsatt, MIT Press, Cambridge MA, pp. 175–194. URL: <https://doi.org/10.7551/mitpress/12389.003.0013>.
- . 2021b. “Downward Causation Defended”, in *Top-Down Causation and Emergence*, ed. by Jan Voosholz and Markus Gabriel, Springer, Cham, pp. 217–252. URL: [https://doi.org/10.1007/978-3-030-71899-2\\_9](https://doi.org/10.1007/978-3-030-71899-2_9).
- . 2021c. “Explanatory Autonomy: The Role of Proportionality, Stability, and Conditional Irrelevance”, in *Synthese*, Vol. 198, No. 1, Jan. 2021, pp. 237–265. URL: <https://doi.org/10.1007/s11229-018-01998-6>.
- . 2022. “Levels, Kinds and Multiple Realizability: The Importance of What Does Not Matter”, in *Levels of Reality in Science and Philosophy*, ed. by Stavros Ioannidis, Gal Vishne, Meir Hemmo, and Orly Shenker, Springer, Cham, pp. 261–292. URL: [https://doi.org/10.1007/978-3-030-99425-9\\_14](https://doi.org/10.1007/978-3-030-99425-9_14).
- Sara Worley. 1993. “Mental Causation and Explanatory Exclusion”, in *Erkenntnis*, Vol. 39, No. 3, Nov. 1993, pp. 333–358. URL: <https://doi.org/10.1007/BF01128507>.
- Stephen Yablo. 1992. “Mental Causation”, in *The Philosophical Review*, Vol. 101, No. 2, Apr. 1992, pp. 245–280. Reprinted in Yablo (2009, pp. 222–248). URL: <https://doi.org/10.2307/2185535>.
- . 1997. “Wide Causation”, in *Noûs*, Vol. 31, Supplement: Philosophical Perspectives, 11, Mind, Causation, and World 1997, pp. 251–281. Reprinted in Yablo (2009, pp. 275–306). URL: <https://doi.org/10.1111/0029-4624.31.s11.12>.

- Stephen Yablo. 2009. *Thoughts: Papers on Mind, Meaning, and Modality*, Oxford: Oxford University Press. URL: <http://doi.org/10.1093/acprof:oso/9780199266463.001.0001>.
- Edward N. Zalta. 2023. *Stanford Encyclopedia of Philosophy*, ed. by Edward N. Zalta. Stanford: Stanford University. URL: <http://plato.stanford.edu/>.