This is the final version of a paper that has been published in Austin, C.J., Marmodoro, A., and Roselli, A. (eds.), *Powers, Parts and Wholes: Essays on the Mereology of Powers* (2023), Routledge: 130-141.

Vera Hoffmann-Kolss

What Can Causal Powers Do for Interventionism? The Problem of Logically Complex Causes

Abstract: Analyzing causation in terms of Woodward's interventionist theory and describing the structure of the world in terms of causal powers are usually regarded as different projects in contemporary philosophy. Interventionists aim to give an account of how causal relations can be empirically discovered and described, without committing themselves to views about what causation really is. Causal powers theorists engage in precisely the latter project, aiming to describe the metaphysical structure of the world. This chapter argues that interventionism can benefit from incorporating considerations about causal powers. The chapter describes a previously undiscussed problem for Woodward's definition of an intervention variable that arises when interventionist causal models include disjunctive variables. This problem is solved not by excluding disjunctive (or logically compound) variables per se but by excluding variables whose values represent disjunctions of overly disparate causal powers. This suggests that interventionism and causal powers theories may be less distant from each other than is often assumed.

Introduction

Causal talk and causal reasoning are ubiquitous. In everyday reasoning, we take individual causal claims involving concrete individuals to be true, such as that Joanne's being late caused her to miss an important meeting or that Jimmy's leaving the cake in the oven too long caused Amy to be in a bad mood (because she does not like the taste of burnt cake). In scientific contexts, we often try to answer general questions about causal relations, such as whether COVID-19 infections cause heart problems or whether increased global temperatures cause forest fires. From a philosophical point of view, this raises the question of under what conditions we can say that certain causal claims are true.

This question can be answered in several different ways. One option is to analyze causal relations in terms of recent interventionist theories of causation à la Woodward (Woodward 2003), which are based on the formalism of Bayesian causal models (Hitchcock 2007; Pearl 2000; Spirtes, Glymour and Scheines 2000). Another option is theories based on the assumptions that there are causal powers in the world and that these causal powers are the driving force that makes things happen. According to this picture, causal relations occur because things instantiate powers, and under certain triggering conditions, these powers produce changes in the world; proponents of powers theories include Bird (2007), Cartwright (1989), Molnar (2003), Shoemaker (1980), and Vetter (2015).

These two approaches come from quite different philosophical camps. Interventionist theories of causation typically aim to be as ontologically undemanding as possible. They were originally designed to capture the way in which causal relations are discovered and described by practicing scientists, especially in fields that work with large data sets, such as econometrics, climate science, or neuroscience research using brain scan data. The general idea is that if the data under investigation show certain patterns (and if additional conditions are met), these patterns can be given a causal interpretation. This approach is typically neutral about the ontological question of whether there are fundamental powers in the world that give rise to causal necessitation relations. Powers theorists, in comparison, are by definition committed to the existence of causal powers, with the remaining question being what exactly is the connection between powers and causes or causal relations (see, e.g., Baltimore 2019; Mumford and Anjum 2011).

Interestingly, however, at least certain parts of the debate on causal models have begun to increasingly incorporate metaphysical considerations and concepts. One example is the debate about the relationship between interventionism, Bayesian models, and the metaphysical causal exclusion problem in the philosophy of mind (see, e.g., Baumgartner 2009; Gebharter 2017; Kroedel and Schulz 2016; Woodward 2015). Another example is the thesis that the metaphysical grounding relation can be analyzed using the formal framework of causal models (Schaffer 2016; Wilson 2018).

A topic that has not been explored in comparable detail, however, is the question of the metaphysical status of the relata of the causal relation. According to causal modeling approaches, causal relations can be described as dependence relations between variables.

These variables assume at least two and possibly infinitely many mutually exclusive values, and these values represent the relata of the causal relation under consideration. This leaves open how logically complex the relata of the causal relation can be, that is, whether only fundamental or logically simple events or properties can be causally related to each other, or whether the values of the variables can also stand for logically complex entities.

In this chapter, I argue that if arbitrary disjunctive variables can occur in causal structures, the interventionist criterion of causation becomes inadequate. I further argue that a possible solution to this problem is to exclude variables whose values represent disjunctions of overly disparate causal powers. I also suggest that this is an instance of a more general observation about interventionism, that is, the observation that interventionism needs at least some 'metaphysical input' in order to adequately represent causal structures in the world.

I begin with an introduction to interventionism. I then argue that the interventionist criterion of causal relevance is problematic if the variables in the causal models under consideration can be disjunctive. I suggest that this problem can be solved on the basis of considerations about causal powers and conclude by briefly pointing out that interventionism in general must rest, at least to some extent, on substantial metaphysical assumptions.

The Interventionist Theory of Causation

According to the interventionist account of causation, causal structures can be represented by directed causal graphs consisting of two elements: (a) a set V of vertices consisting of variables standing in causal relations to each other, and (b) a set of directed edges connecting these vertices. If a sequence of variables $\{X_1, \ldots, X_n\}$ is such that for any i with $1 \leq i < n$, there is a directed edge from X_i to X_{i+1}, then the sequence is a directed path leading from X_1 to X_n (Spirtes, Glymour and Scheines 2000; Woodward 2003). Directed paths between variables represent causal relevance relations.

The standard locution in this context is that *variables* stand in causal relevance relations to each other. However, this should be understood as shorthand for the claim that the entities represented by the values of variables stand in causal relevance relations to each other. These entities may be properties or events. Accordingly, the interventionist criterion of causation as it is presented here can be applied either to so-called *token* *causation*, that is, causal relations between concrete events, or to so-called *type causation*, that is, causal relevance relations between properties. The relationship between these two kinds of causation is complicated (see, e.g., Hausman 2005). However, the details are not important for the purposes of the present argument. Therefore, for the sake of simplicity, I will simply follow the convention of talking about causal relevance relations between *variables* and leave open whether the values of these variables stand for events or properties.

Against the background of a formal causal model, the interventionist criterion of causation classifies a variable X as causally relevant to a variable Y iff there is an intervention on the value of X that changes the value or the probability distribution of Y, provided that the values of all other variables that are not on the causal path between X and Y are held fixed by interventions (Woodward 2003; Hitchcock 2001, 2007).

The notion of an intervention is an essential component of Woodward's approach. Interventions are characterized by intervention variables, which are defined as follows: 'I is an intervention variable for X, with respect to Y, if it meets the following conditions:

(1) I is causally relevant to X.

(2) I is not causally relevant to Y through a route that excludes X

(3) I is not correlated with any variable Z that is causally relevant to Y through a route that excludes X, be the correlation due to I's being causally relevant to Z, Z's being causally relevant to I, I and Z sharing a common cause, or some other reason.

(4) I acts as a switch for other variables that are causally relevant to X. That is, certain values of I are such that when I attains those values, X ceases to depend upon the values of other variables that are causally relevant to X.'

(Woodward and Hitchcock 2003: 12–13)

The purpose of intervention variables is to do justice to the role of interventions in controlled experiments. To find out whether some variable X is causally relevant to some variable Y in an experimental situation, one must manipulate the value of X and hold all possible confounding factors fixed.

Confounding can occur when the intervention is directly causally relevant to the effect under consideration. For example, if the effect of a painkiller on headaches is being studied, an intervention could be to offer a painkiller to a person with headaches. However, if the intervention consists of offering a painkiller together with a glass of water, the intervention itself might be directly causally relevant to the effect variable, since headaches can sometimes occur when a person is (mildly) dehydrated, and therefore, drinking water may also have a positive effect on headaches. Condition (2) is designed to eliminate this type of confounding.

Sometimes confounding occurs because the intervention is correlated with some other cause of the effect variable. For example, if a person's headache is due to too much exposure to the sun, and the intervention is to ask her to come into an air-conditioned room to take a painkiller, the intervention is interfering with a possible cause of the headache. Condition (3) is designed to rule out such confounding.

In the next section, however, I argue that condition (3), and thus Woodward's definition of an intervention, becomes problematic if the variable Z, that is, the variable describing possible alternative causes of Y, can be disjunctive.

The Problem of Disjunctive Causes

Consider a causal structure with two variables, one describing whether a person has a headache and one describing whether the person has an allergic rash. Suppose, somewhat simplistically, that the variables are binary and defined as follows:

R:= 1 if Person P has allergic rash; 0 otherwise
H:= 1 if Person P has a headache; 0 otherwise

If an experimenter wants to find out whether H is causally relevant to R, the interventionist criterion requires that she carry out an intervention on the value of H and observe whether this intervention has an effect on the value of R. Suppose, in contrast to the experimental setup described in the previous section, where the goal was to investigate the effect of a painkiller on headaches, the experimenter now has a painkiller that she knows to be very effective against headaches. Accordingly, if the initial value of H is H = 1, that is, if Person P has a headache, an obvious intervention is to administer this painkiller.

Again, this intervention must satisfy the conditions of an intervention variable described in the second section. In particular, it must satisfy condition (2) and not be causally relevant to R via a causal path that excludes H. This condition would be violated, for example, if the painkiller also had an effect on allergic rashes. The intervention must also satisfy condition (3) and not be correlated with any other cause of R that is on a causal path that excludes H. For example, if Person P is allergic to apples and develops an allergic rash after eating apples or apple products, administering the painkiller with a glass of apple juice would not satisfy this condition.

However, if administering the painkiller to Person P does satisfy the conditions of an intervention, and if it does indeed have an effect on the value of R, then the experimenter is justified in concluding that H is causally relevant to R. Note that the converse is not true. The interventionist causal criterion has the structure of an existential condition: if there is an intervention on X that changes the value of Y, then X is causally relevant to Y. If a particular intervention on X does not produce a change in Y, then it cannot be excluded that there is some other intervention on X that does change the value of Y. If, however, it can be plausibly concluded that no intervention on X changes the value of Y, then X is not causally relevant to Y.

One of the strengths of the interventionist approach to causation is that it analyzes not only causal relations between two relata but also entire causal networks. Accordingly, interventionism is able to represent even complex causal structures with multiple interconnections. Taking into account the more complex causal structure is also crucial for the notion of intervention, since interventions must be independent of possible alternative causes of the effect in question. For example, as noted earlier, if Person P is allergic to apples and develops an allergic rash when consuming apples or apple products, then the intervention on Person P's headache should not interact with the following variable:

A:= 1 if Person P consumes apples or apple products; 0 otherwise.

Moreover, as was also pointed out earlier, an intervention on H with respect to R should be independent not only of A but also of any other variable that is causally relevant to R on a causal path that does not include H. Interventionism *per se* does not rule out that the variables to be taken into account are logically disjunctive or otherwise logically complex. This is consistent with the pragmatic approach of interventionism, according to which the variables included in a causal structure should be chosen according to the purpose of the scientific inquiry (Woodward 2016). Variables assuming logically compound values should be taken into consideration if this is in accordance with the purpose of scientific inquiry. Considerations concerning the metaphysical structure of the properties represented by the values of a variable are at best of secondary importance.

Variable A illustrates this consideration. The value A = 1 represents a disjunction of different properties, *eating (pure) apples* and *eating (processed) apple products*. In this context, it makes perfect sense to consider a disjunctive variable that refers to apples and all foods containing apples, since the relevant property is not how exactly the apples are prepared but whether the person somehow ingests substances contained in apples.

Not all disjunctive variables are like this, however. Consider another disjunctive variable that is defined as follows:

> (A v not-I) := 1 if Person P consumes apples or apple products or does not take a pain killer; 0 otherwise

The 'or'-operator in (A v not-I) is to be understood as nonexclusive; that is, the variable assumes the value (A v not-I) = 1 if P consumes both apples and an apple product and does not take a painkiller.

Obviously, (A v not-I) is a rather artificial variable. But it does cause problems for the interventionist criterion of causation. To see this, first note that (A v not-I) is causally relevant to both R and H. To see that (A v not-I) is causally relevant to R, suppose that Person P does not have an allergic rash, that is, R = 0. Further suppose that she has not yet taken any painkillers and has not consumed any apples or apple products. Given the disjunctive structure of (A v not-I), this implies that (A v not-I) = 1. If P is now offered a glass of apple juice, which she is willing to drink, this is an intervention that changes the value of (A v not-I) to (A v not-I) = 1. But this intervention also changes the value of R: if Person P drinks the apple juice, she is likely to develop an allergic rash.

Note that not every intervention that changes the value of (A v not-I) will have an effect on the value of R. If the value of (A v not-I) is manipulated by asking Person P to take a painkiller instead of drinking the apple juice, the value of R will remain unaffected. Here the fact that the interventionist causal criterion has the structure of an existential requirement becomes relevant again. There is an intervention on (A v not-I) that changes the value of R, and this is sufficient to conclude that (A v not-I) is causally relevant to R. Since it can be shown exactly analogously that (A v not-I) is also causally relevant to H, it follows that (A v not-I) is causally relevant to both R and H. This structure is illustrated in Figure 7.1.

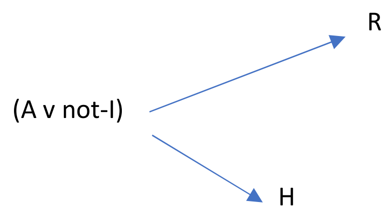

Figure 7.1 Causal structure in which the disjunctive variable (A v not-I) describing the use of painkillers (I) and the consumption of apples or apple products (A) is causally relevant to both allergic rash (R) and headaches (H).

If (A v not-I) is causally relevant to R, then the interventionist criterion of causation becomes problematic. Recall that the original question was whether H is causally relevant to R. An intervention variable on H with respect to R must satisfy condition (3) of the definition of an intervention: it must not be correlated with any variable Z that is causally relevant to R through a route that excludes H, be the correlation due to the intervention's being causally relevant to Z, Z's being causally relevant to the intervention, the intervention and Z sharing a common cause or some other reason (Woodward and Hitchcock 2003).

In the causal structure depicted in Figure 7.1, (A v not-I) is such a variable Z, since it is causally relevant to R through a route that excludes H. However, an intervention on H that consists of administering a painkiller to Person P is correlated with (A v not-I) *for some other reason*, the reason being that the intervention variable describing it is a logical component of (A v not-I). Therefore, administering a painkiller does not satisfy the conditions of an intervention on H with respect to R.

And even worse, no other possible intervention satisfies this condition! The not-I component of (A v not-I) can be replaced by any possible intervention variable I. (A v not-I) will then still be causally relevant to R, and I will not satisfy condition (3). But then there is no variable I that can be an intervention variable on H with respect to R, and given the logical structure of the interventionist criterion of causation, this implies that H cannot be causally relevant to R on purely formal grounds – *regardless of what the true causal structure is*.

Moreover, the problem does not depend on this particular example. I arises whenever one tries to investigate the causal influence of some variable X on some variable Y which has some other cause Z. Then any intervention variable I can be combined with Z to form a disjunctive variable (Z v not-I), and the situation has exactly the same structure as shown in Figure 7.1 (with X = H, Y = R, and (A v not-I) = (Z v not-I)). Accordingly, there is no variable I that satisfies the conditions of an intervention variable on X.

One possible response is to argue that condition (3) should not be interpreted too strictly. Although the intervention of administering a painkiller is correlated with the disjunctive variable (A v not-I), the correlation is not due to any causal connection but arises only from the logical structure of (A v not-I), and such correlations should be exempted from the conditions that an intervention variable must satisfy.

However, this would lead to even more problematic results. Suppose that H is in fact causally independent of R, that is, there is no causal relationship between Person P's headache and her allergic rash. Suppose the initial value of H is H = 1, that is, Person P has a headache, and the initial value of R is R = 0, that is, Person P does not have an allergic rash. Further assume, as earlier, that she has not taken any painkillers and has not consumed any apples or apple products, that is, (A v not-I) = 1. Now suppose that P is given a painkiller that effectively cures her headache, so that the value of H changes to H = 0. The interventionist criterion of causation requires that an intervention on the cause variable under consideration must hold the values of all other variables not on a causal path between the cause and effect variables fixed. Accordingly, if Person P is given a painkiller, the value of (A v not-I) must be held fixed. Given that (A v not-I) = 1, this means that P must consume an apple or an apple product. But then she is likely to develop an allergic rash. It follows that if the value of H is changed by an intervention while the value of (A v not-I) is held fixed, the value of R will also change – even though H and R are causally unrelated. But this seems absurd.

The real problem seems to arise from the assumption that artificial variables with the structure of (A v not-I) must be taken into account when carrying out interventions. How can this be avoided? In the next section, I argue that causal powers might prove helpful in this regard.

Logically Complex Variables: A New Field of Work for Powers?

As argued in the previous section, the interventionist criterion of causation is problematic when causal structures include logically complex variables that contain intervention variables as components. An immediate response might be to simply exclude disjunctive variables. It is well known in formally oriented metaphysics and philosophy of science that disjunctive entities cause problems here and there, and it may be a perfectly legitimate move, at least in some contexts, to simply ban them from theory.

It is also well known that not everything that can be formally written as a disjunction is a disjunctive entity. The property of *eating an apple* is logically equivalent to the property of *either eating an apple and eating a piece of chocolate or eating an apple and not eating a piece of chocolate* (i.e., (*eating an apple* & *eating chocolate*) v (*eating an apple* & *not eating chocolate*)). However, this does not mean that it is a disjunctive property.

Often, disjunctiveness is defined in terms of naturalness. The classical metaphysical view in this respect is that of Lewis, according to which there is a class of perfectly natural properties that are included in the fundamental laws of nature (Lewis 1983, 1986). Other properties can be constructed from perfectly natural properties by logical operations. The more logical operations are required to construct a property from perfectly natural properties, the less natural it is. This means that properties can be ordered according to their degree of naturalness. A property is disjunctive if it is less natural than any of the disjuncts into which it can be decomposed (Langton and Lewis 1998).

Would it help to exclude variables some of whose values represent disjunctive properties? Obviously, the problematic variable (A v not-I) would have to be excluded. One of its values represents the state that Person P consumes apples or apple products or does not take a painkiller, and this state is clearly constituted by a property that is disjunctive in Lewis's sense.

However, as noted previously, a strict prohibition on disjunctive variables in causal structures would be too strong a requirement. In investigating whether Person P's headaches are causally related to her developing a rash, one must keep in mind other possible causes of her rash. And, given the context, the disjunctive variable A, representing whether she consumes apples or apple products, is one of those variables.

One might object that instead of the disjunctive variable A, one could consider a variable representing whether Person P consumes apples and another variable representing whether she consumes apple products. However, it is plausible to assume that the property of *consuming apple products* is also disjunctive and consists of an infinite set of disjuncts representing the various apple products one might consume: apple juice, apple pie, apple ice cream, red cabbage with apples, and other possible apple dishes that have not even been invented. Given that causal structures are supposed to consist of a finite set of variables, one cannot include all these disjuncts into a causal structure. However, all these disjuncts are captured under the umbrella property of *consuming apple products*.

The reason why it makes sense to use the disjunctive umbrella variable A is that all the disjuncts of *consuming apples or apple products* share a causal power that is relevant in the given context – the power to produce a rash in people with a certain condition. The crucial difference between variable A and variable (A v not-I) is not that the latter is even more disjunctive than the former but that the values of (A v not-I) represent even more diverse causal powers than the values of A.

Here, then, is my hypothesis about how the interventionist framework can benefit from considerations about causal powers. The values of variables in a causal model represent properties or events that are constituted by properties. According to Shoemaker's (1980) causal conception of properties, properties confer conditional causal powers on their bearers. For example, the property of being a painkiller, instantiated by a pill, gives the pill the causal power to relieve headaches on the condition that it is swallowed by a person who has headaches.

Now consider condition (3) of Woodward's definition of an intervention. According to this condition, an intervention variable I for X, with respect to Y, must not be correlated with any variable Z that is causally relevant to Y through a route that excludes X. Suppose that Z is disjunctive in the sense that at least one of its values has multiple disjuncts. According to Shoemaker, each of these disjuncts confers a set of conditional powers on its bearers. Consider the intersection S of all these sets. The intervention variable I should be required to be independent of Z iff S is not empty and at least some of the powers contained in S are relevant to Y. Otherwise, if S is empty or if all the powers contained in S are not relevant to Y, the intervention variable I need not be independent of Z.

In our example scenario involving Person P, her headaches, and her allergic rash, an intervention on H should be independent of variable A. If A = 1, then the disjuncts of this value share a conditional causal power, the power to cause rashes in people who are allergic to one of the components of apples and who ingest apples. Moreover, this power is directly relevant to the effect variable R. In contrast, an intervention on H need not be independent of (A v not-I). It is plausible to assume that the intersection of the conditional causal powers of eating apples, eating apple products, and not taking a painkiller is empty. Therefore, (A v not-I) does not fall into the category of variables that need to be considered as possible confounders.

If Z is non-disjunctive, then each of its values is non-disjunctive. Suppose that Z_i is one of the values of Z. Then Z_i confers a characteristic set of conditional causal powers on its bearers. Since Z_i is non-disjunctive, it has only one disjunct, and the set S is identical to the conditional causal powers of Z_i. It is plausible to assume that Z is causally relevant to Y if some of the conditional powers of any of its values are relevant to Y. Therefore, the limiting case in which Z is non-disjunctive collapses into the original definition of an intervention: an intervention variable I for X, with respect to Y, must not be correlated with any variable Z that is causally relevant to Y via a route that excludes X.

This consideration is in accordance with a restriction on variables imposed by Woodward (2016), who argues that the variables appearing in a causal model should have unambiguous effects on the other variables in the model. One way to understand the notion that variables have ambiguous effects is to assume that the powers of their values are too diverse: if the value of (A v not-I) is manipulated by administering a painkiller to Person P, the possible effects of this manipulation will be very different from the possible effects of a manipulation that consists of feeding apples to P. The powers-based account proposed in this section provides a metaphysical explanation of what ambiguous manipulations might be and why they should be avoided.

Conclusion

The aim of this argument was to show that supplementing the interventionist criterion of causation with a condition based on causal powers can solve a problem arising from logically complex variables. Of course, the argument does not show that powers ontology is the only way to solve the disjunctive variable problem described in the second section. Moreover, it leaves open the extent to which the interventionist criterion might benefit from considerations about causal powers. However, I have already argued elsewhere that the interventionist criterion of causation must rely on stronger metaphysical assumptions than is often assumed, and in particular must presuppose the distinction between metaphysical and nomological necessity (Hoffmann-Kolss 2022). The argument of this chapter can therefore be seen as an example of a more general claim: interventionist causation can only count as a fully-fledged philosophical theory of causation if it relies, at least to some extent, on substantive metaphysical claims about causation and causal relations – be they claims about necessity, causal powers, or something else.

References

Baltimore, J. (2019), 'Expanding the Vector Model for Dispositionalist Approaches to Causation', *Synthese* 196(12): 5083-5098.

Baumgartner, M. (2009), 'Interventionist Causal Exclusion and Non-reductive Physicalism', *International Studies in the Philosophy of Science* 23(2): 161-178.

Bird, A. (2007), *Nature's Metaphysics: Laws and Properties*: Oxford University Press.

Cartwright, N. (1989), *Nature's Capacities and Their Measurement*, Oxford University Press.

Gebharter, A. (2017), 'Causal Exclusion and Causal Bayes Nets', *Philosophy and Phenomenological Research* 95(2): 353-375.

Hausman, D.M. (2005), 'Causal Relata: Tokens, Types, or Variables?', *Erkenntnis* 63(1): 33-54.

Hitchcock, C. (2001), 'The Intransitivity of Causation Revealed in Equations and Graphs', *Journal of Philosophy* 98(6): 273-299.

——— (2007), 'Prevention, Preemption, and the Principle of Sufficient Reason', *Philosophical Review* 116(4): 495-532.

Hoffmann-Kolss, V. (2022), 'Interventionism and Non-Causal Dependence Relations: New Work for a Theory of Supervenience', *Australasian Journal of Philosophy* 100(4): 679-694.

Kroedel, T. and Schulz, M. (2016), 'Grounding Mental Causation', *Synthese* 193(6): 1909-1923.

Langton, R. and Lewis, D. (1998), 'Defining "Intrinsic"', *Philosophy and Phenomenological Research* 58(2): 333-345.

Lewis, D. (1983), 'New Work for a Theory of Universals', *Australasian Journal of Philosophy* 61(4): 343-377.

——— (1986), *On the Plurality of Worlds*, Blackwell Publishers.

Molnar, G. (2003), *Powers: A Study in Metaphysics*, edited by S. Mumford, Oxford University Press.

Mumford, S. and Anjum, R.L. (2011), 'Spoils to the Vector - How to Model Causes If You Are a Realist about Powers', *The Monist* 94(1): 54-80.

Pearl, J. (2000), *Causality: Models, Reasoning, and Inference*, Cambridge University Press.

Schaffer, J. (2016), 'Grounding in the Image of Causation', *Philosophical Studies* 173(1): 49-100.

Shoemaker, S. (1980), 'Causality and Properties', in P. Van Inwagen (ed.), *Time and Cause. Essays Presented to Richard Taylor*: Reidel: 109-136; quoted from: S. Shoemaker (1984), *Identity, Cause, and Mind*: Cambridge University Press: 206-233.

Spirtes, P., Glymour, C. and Scheines, R. (2000), *Causation, Prediction and Search*, 2nd ed.: The MIT Press.

Vetter, B. (2015), *Potentiality: From Dispositions to Modality*: Oxford University Press.

Wilson, A. (2018), 'Metaphysical Causation', *Nous* 52(4): 723-751.

Woodward, J. (2003), *Making Things Happen: A Theory of Causal Explanation*: Oxford University Press.

——— (2015), 'Interventionism and Causal Exclusion', *Philosophy and Phenomenological Research* 91(2): 303-347.

——— (2016), 'The Problem of Variable Choice', *Synthese* 193(4): 1047-1072.

Woodward, J. and Hitchcock, C. (2003), 'Explanatory Generalizations, Part I: A Counterfactual Account', *Noûs* 37(1): 1-24.