

The Two Sides of Interventionist Causation

Peter W. Evans* and Sally Shrapnel†

School of Historical and Philosophical Inquiry
University of Queensland
Australia

Abstract

Pearl and Woodward are both well-known advocates of interventionist causation. What is less well-known is the interesting relationship between their respective accounts. In this paper we discuss the different perspectives of causation these two accounts present and show that they are two sides of the same coin. Pearl's focus is on leveraging *global* network constraints to correctly identify local causal relations. The rules by which global causal structures are composed from distinct causal relations are precisely defined by the global constraints. Woodward's focus, however, is on the use of *local* manipulation to identify single causal relations that then compose into global causal structures. The rules by which this composition takes place emerge as a result of local interventionist constraints (or so the claim goes). We contend that the complete picture of causality to be found between these two perspectives from the interventionist tradition must recognise both the global constraints of the sort identified by Pearl and the local constraints of the sort identified by Woodward, and the interplay between them: Pearl requires the possibility of local interventions and Woodward requires a global statistical framework within which to build composite causal structures.

1 Introduction

Motivated by the relative scarcity of formal tools available for characterising causal reasoning in statistics, medicine, economics, social science, and especially the fields of artificial intelligence and cognitive science, Pearl (2009) has developed an account of causality based upon probabilistic directed acyclic graphical models, or Bayesian networks. Such networks are given a causal interpretation through the application of two assumptions that apply at a statistical level: the causal Markov and faithfulness assumptions. These

*email: p.evans@uq.edu.au

†email: s.shrapnel@uq.edu.au

assumptions provide a *global* constraint—in the sense that they apply across the complete set of variables in the network—on sets of probability distributions to ensure that such distributions are appropriate for characterising causal structure. A further significant element of Pearl’s framework is the integral role played by interventions, in what Pearl calls the ‘do-calculus’, to provide a means for deciding between underdetermined graphical models.¹

Building upon Pearl’s framework, Woodward (2003) has embarked on the project of providing a more general philosophical foundation for the notion of an intervention as the basis for a manipulability theory of causation. According to manipulability theories, some event *C* is a cause of some other event *E* just in those cases where manipulating *C* would be an appropriate means for manipulating *E*. Woodward’s project can be characterised as an attempt to supply a precise and explicit characterisation of the conditions under which such manipulation is an appropriate probe of causal relations, where the resulting notion is very much commensurate with Pearl’s calculus of interventions—indeed, ‘intervention’ has emerged as the standard parlance for this formalised notion. These conditions provide a *local* constraint—in the sense that they apply only to a specific proximate subset of variables—on the statistical dependences and independences between interventions and the relevant variables in the model.

There is considerable overlap between the approaches of Pearl and Woodward, and their respective characterisations of interventions could quite plausibly be referred to indistinctly. It seems to us, however, a worthwhile project to point out the similarities and differences between the two approaches.² Notably for our present purposes, Pearl (2009, p.23) realises the importance of the possibility of local interventions (that is, interventions that do not disrupt other causal relationships in the model—more on this below) and that such interventions play an integral role in guaranteeing that the causal Markov condition and faithfulness are satisfied by the relevant graphical representation. Likewise, Woodward (2003, p.64) is aware of the significance of the causal Markov condition and faithfulness in specifying (albeit non-reductively) whatever connection there might be between causation and the conditional probability distributions defined over a Bayesian network. However, despite these conspicuous similarities there is also a strong sense in which the approaches of Pearl and Woodward differ in the emphasis they each put on the global and local constraints, respectively. Here is Ismael (forthcoming, p.23, fn.iv) observing this very point:

for Pearl, once you know what the causal mechanisms are, you can say which interactions constitute interventions. Woodward thinks that this limits the utility of interventions to discover causal mechanisms (among other things) and wants to characterize the

¹While Spirtes *et al.* (2000) develop a similar approach to causality in terms of Bayesian networks, and certainly consider interventions as an important part of their account, their work contains less of an emphasis on the calculus of interventions.

²For a recent similar project that reconstructs Woodward’s interventionist account of causation within the formalism of causal Bayesian networks, see Gebharter (2017, ch.5).

notion of an intervention independently so that it can be used as a probe for causal structure. To some extent this in house dispute reflects a difference in focus. From a metaphysical perspective, it is natural to take the underlying causal structure as basic. It is what explains the surface regularities and patterns of counterfactual dependence. But Woodward is interested in using interventions as a route in, so to speak. He wants to be able to identify interventions (perhaps provisionally) before we have a detailed understanding of the causal structure and use them to probe.

This “difference in focus”, as Ismael puts it, we take to be suggestive of a deeper point about the picture of causality obtained from the casual modelling and interventionist programmes. We contend that the complete picture of causality to be found here must recognise both the global constraints of the sort identified by Pearl and the local constraints of the sort identified by Woodward, and the interplay between them. On such a view, then, it would be wrong to consider what we have called here global constraints, such as the casual Markov condition and faithfulness, as capturing the essence of causality over and above the conditions for local interventions, and it would likewise be wrong to consider the conditions for local interventions as capturing the essence of causality over and above global constraints, such as the casual Markov condition and faithfulness. Rather, a globally constrained network is able to attach causal meaning to local edges *just when* one commits to the possibility of local interventions and, conversely, a collection of locally identified causal links can be composed to produce causal structures *just when* global constraints on the network provide the framework for composition. We will proceed as follows.

We begin in Section 2 by outlining the causal discovery schema of Pearl (2009) and the role of interventions within that schema. For Pearl, causal discovery begins at the level of statistical constraints on observable data. His framework rests on a hierarchy of assumptions. To begin with, the axioms of probability theory and two statistical assumptions, the causal Markov condition and faithfulness, underpin the causal interpretation of conditional dependences and independences. This interpretation then permits the possibility of interventions providing causal information. Pearl’s foundation is thus basic probability theory, from which conditional dependences and independences arise from statistical data and which then obtain a causal interpretation. In Section 3 we discuss Woodward’s (2003) interventionist account of causation. For Woodward, the identification of causal relations begins at the level of local interventions: Woodward’s foundation is the set of conditions that defines an intervention, from which we elicit causal relations given the right sort of assumptions for the statistical behaviour of variables. We consider in Section 4 the interrelation between interventions and the causal assumptions that underpin the identification of causal relations; we argue that this interrelation resists explication in terms of one or the other element as foundation.

2 Causal Bayesian networks as oracles for interventions

The primary formal tool that Pearl uses for causal representation is the Bayesian network. A Bayesian network is a DAG that coherently captures local and global probabilistic information about some system and whose nodes represent the distinct variables X_1, \dots, X_n that we take to model that system. For each node X_i of some Bayesian network G there is an associated conditional probability distribution that denotes the dependence of the value of X_i on its parents in the graph, $P(X_i|pa_G(X_i))$, where parent nodes and their child nodes are connected by directed edges. A necessary condition for a DAG G to be a Bayesian network is the statistical assumption that the joint probability distribution P factorises into a product of conditional probability distributions via the chain rule $P(X_1, \dots, X_n) = \prod_i P(X_i|pa_G(X_i))$.

Given a set of distributions that are all compatible with a DAG G , we can abstract away from the statistical detail of these distributions and characterise the set by the conditional independences that each distribution must satisfy. This is simply a property of the graph G : a graphical criterion known as *d-separation* establishes an equivalence between the structure of the graph and the list of correct conditional independences that hold in the distribution. That is, if a joint probability distribution P factorises over a graph G , and d-separation relative to a further set $Z_{dsep}(G(X, Y|Z))$ holds, then P satisfies $X \perp\!\!\!\perp Y \mid Z$.³ Thus the essential goal of d-separation is to determine when two sets of variables are independent given a third set and searching for such sets forms a critical step in constructing causal DAGs from probabilistic data.

If the set of all independences derived via the d-separation rule from the graph G is satisfied by P , then G accurately maps the independences of P and is thus an independence map, or *I-map*, of P . In general, I-maps are not unique. DAGs are underdetermined by observational data alone and so several graph structures can return the same probabilistic independences; for instance, the three graph structures in Fig. 1 all capture the independence $X \perp\!\!\!\perp Z \mid Y$. Such graphs are said to be independence equivalent, or *Markov equivalent*. Thus there is usually an equivalence class of graphs that contain valid independence relations that hold in the distribution. However, if the independences in the graph precisely capture *all and only* the independences in the probability distribution P , then we say G is a perfect map (*P-map*) of P . Ideally, P-maps are preferred to I-maps, as if we do not include all the independences possible in the distribution then our model is unnecessarily complicated. It is worth emphasising, however, that there is not a P-map for every probability distribution.

So far we have yet to describe how a Bayesian network can be given a causal interpretation; that is, how the directed edges can come to imply causal relations. For this Pearl

³ $(X \perp\!\!\!\perp Y \mid Z)$ denotes that X and Y are conditionally independent given Z :

$$P(X, Y|Z) = P(X|Z) \cdot P(Y|Z).$$




Figure 1: Casual underdetermination of probabilistic independence $X \perp\!\!\!\perp Z \mid Y$.

employs two assumptions that apply at the statistical level and serve to constrain the causal relations of any ensuing DAG. The first of these is the causal Markov condition: any variable is conditionally independent of its non-descendants given its parents. The causal Markov condition follows from the compatibility of G and P (Pearl, 2009, p. 19), which itself follows from the factorisation condition—a necessary condition for Bayesian networks. The second of these is the faithfulness assumption which requires that the independences encoded in the (Markovian) graph are the *only* independences that exist in the distribution. For unfaithful models there are conditional independences obtained from the statistics that are not directly due to the causal structure. By combining faithfulness with the causal Markov condition we can see that G^* must be a perfect map for P^* . That is, the causal Markov condition only provides a necessary condition for edges in the graph to convey causal meaning, but the faithfulness assumption is required to ensure sufficiency. Faithfulness excludes the possibility of causal graphs with causal paths that cancel exactly to create statistical independences between causally connected nodes.

Pearl suggests that two physically motivated premises underpin his claim that Bayesian networks that are Markovian and faithful can be used to represent causal structure. Firstly, he assumes that local conditional probability distributions (between parent nodes and their respective child node) encode information about local (in time and space) stochastic mechanisms, where the stochasticity is due to our ignorance about other variables that we do not include in the model. This means that the relationships between the values of parent and child nodes can be expressed as a function $P(X_i) = f_i(pa(X_i), U_i)$, where the U_i represent all the unmodelled influences on the variable values (noise). Crucially, Pearl assumes the U_i to be uncorrelated. For every real world system, Pearl assumes that there is a ‘true’ joint probability distribution P^* that can be derived from the ‘true’ causal graph G^* that represents the causal structure of the system.

Secondly, the act of setting a variable to a particular value—an intervention—can deterministically override the ‘natural’ causal mechanisms of the model, providing us with new information, by disrupting only the local mechanism associated with that node. This latter point is significant: a local intervention does not disrupt other causal relationships/conditional probability distributions in the model, they act to replace the original causal mechanism with one that determines the child variable value $X = x$ with prob-

ability 1. For Pearl, given some Bayesian network representation of a series of data, postulating that any associated U_i are uncorrelated and that it is possible to intervene locally on the system guarantees that the causal Markov condition is satisfied by the representation. This then permits a causal interpretation of the Bayesian network. Moreover, such a causal interpretation renders the Bayesian networks “oracles for interventions” (Pearl, 2009, p. 22); that is, they provide the ability to correctly infer the effects of various interventions on some system—this is a key strength of causal modelling via Bayesian networks.

Pearl introduces the do-calculus to formalise the act of intervention, enabling us to identify the kind of situations in which we can accurately predict the effects of interventions given the structure of the DAG and the accompanying joint distribution. Given a set of variables (X, Y, Z) we may wish to tell what the distribution of Y is given that X has been intervened upon from outside and had its value set deterministically to a single value x . This conditional distribution is represented as $P(Y|do(X = x))$ and it is crucial to recognise that this distribution is not always the same as the general conditional $P(Y|X)$. An integral part of Pearl’s schema is the process of identifying when we can correctly infer the effects of interventions given observations on other nodes for a given DAG and factorising joint distribution.

We get, then, the following picture of causation from Pearl’s schema. The basic axioms of probability theory let us derive conditional dependences and independences from statistical data, and they place certain constraints on the conditional probability relations between the variables that characterise the system data. As a result of this we develop DAG representations that are Bayesian networks satisfying the factorisation constraint. Postulating that (i) there are no unmeasured common causes (any U_i are uncorrelated) and (ii) local interventions are possible gives us grounds for assuming that the causal Markov condition and faithfulness hold, and thus we can provide a causal interpretation of the Bayesian networks such that they are “oracles for interventions”. (An underdetermined graph still requires interventionist information but the fact that these interventions are meaningful at all rests upon this framework.) Thus the causal content of interventions is grounded upon the causal interpretation of Bayesian networks, which is built upon the probability axioms and some physically motivated premises. This facilitates a key strength of Pearl’s formalism, to provide a method to uncover causal relations when interventions are in fact *not* possible, for example, for practical or ethical reasons.

3 Woodward’s local interventions

The point of departure for Woodward’s (2003) interventionist account of causation is, of course, the interventions themselves.⁴ Woodward defines a distinctive type of variable

⁴See also Hausman and Woodward (1999, 2004a,b); Woodward (2007, 2008, 2013, 2015).

with which to model causal relations in a system called an intervention variable, I . What defines an intervention variable is the satisfaction by I of a series of conditions that constrain the nature of the probabilistic relation between I , X and Y , and any other relevant variables which we take to be causally related to X and Y (Woodward, 2003, p. 98):

- (i) I causes X ;
- (ii) I breaks the relation between X and the rest of its causes;
- (iii) I is not (directly or indirectly) causally related to Y except (if at all) through X ;
- (iv) I is statistically independent of any variable that is both a cause of Y and is part of a causal chain that does not include X .

The essential idea behind these conditions is to place I in total control of the value of X and eliminate any correlations between X and Y that are not a function of the intervention on X . The interventionist account of causation is then a manipulability account: put simply, the relation between X and Y is causal if and only if there exists a possible intervention on X that changes the probability distribution of the possible values of Y , holding fixed all other variables relevant to the system.

There are some important consequences that follow as a result of these conditions that provide a more complete illustration of the notion of causation according to the interventionist account. Firstly, it is plausible that some possible interventions on X will bring about relevant changes in Y while other possible interventions will not. A causal relation is said to be *invariant* when the functional relation between X and Y , $Y = f(X)$, that establishes the causal relation holds for at least some range of possible interventions, $X = x_1 \dots x_n$. Secondly, there will likewise be a range of background conditions under which possible interventions on X will bring about relevant changes in Y while other possible interventions will not. A causal relation is said to be *stable* when there is at least some range of background conditions under which the relation between X and Y is causal. Both the notions of invariance and stability are relative notions: there might exist a causal relation between X and Y under a certain range of possible interventions and background conditions that breaks down under other possible interventions and background conditions. Only under the appropriate domains for both interventions and background conditions where the functional relation can be established can we claim that X causes Y ; this is because it is only within these domains that we can think of manipulating X as an appropriate means for manipulating Y .

This consideration provides a constrained framework within which the notion of an intervention is an appropriate precursor to the discovery of a causal relation. Indeed, if we ignore the influence of the background context on the functional relationship between the values of purportedly connected variables in a system, the claim that interventions are a good guide to causation is rather weak: any change in Y given a change in X

would lead us to claim that X causes Y . Given sufficient information concerning the background context, we can begin to explore “*which* interventions on X will change Y , and *how* they will change Y , and under what background circumstances” (emphasis in original) (Woodward, 2007, p. 76), providing a more robust model by which to establish causal relations.

Another consequence of the above conditions is that the variables upon which one might intervene must be chosen such that they are sufficiently distinct from each other: an intervention must be ‘surgical’ so as to ensure that the variable upon which the intervention is being made is the only variable influenced by the intervention. We say that an intervention is “fat-handed” or confounding when it either directly affects Y , or indirectly affects Y by affecting other variables that are not on the $I - X - Y$ path, in addition to affecting Y through X . In a similar vein, for a given set of functional relations between a set of variables to correctly represent the causal facts concerning some system, the interventionist account requires that the functional relations are *modular*; that is, an intervention I on some variable X does not alter the functional relation between the putative effect Y and any of its causes that are not on a directed path from X to Y (Hausman and Woodward, 1999, p. 543) (cf. Pearl’s local interventions). Modularity requires that some functional relation is invariant and stable over some range of interventions and background conditions (and thus is also a relative notion), and any other functional relations in the system remain unchanged when an intervention is carried out. This notion captures the idea of a system being constituted by distinct causal mechanisms.

Significantly, the invariance, stability and modularity of a set of functional relations relative to a range of interventions and background conditions is connected to the idea that the level of detail or generality of the variables that we take to characterise these functional relations—the *level of grain*—in a sense needs to be stipulated. So long as we stipulate a level of grain for the variables and relations of a system such that (i) we are able to intervene on the system as per the above criteria, (ii) the functional relations between the variables are sufficiently modular, and (iii) there are appropriate ranges of invariance and stability under which the functional relations hold, then we can take the model to represent causal relations. Variables and functional relations with these properties may manifest themselves at finer or coarser grains. The appropriate level of grain at which to model a system is dependent upon the sort of causal information one wishes to obtain by way of intervention: “The choice of grain associated with the causal analysis of a situation is intimately related to the contrastive character of causal claims. As we alter the grain, we alter the potential contrastive foci that are available” (Woodward, 2007, p. 85n). Likewise, whether a system can be characterised at all as being constituted by causal relations will itself depend upon the particular coarse-graining that is chosen.

Of interest given our focus in this work, Woodward claims the interventionist account does not require the causal Markov condition as an underlying assumption to elucidate the content of causal claims. Furthermore, Hausman and Woodward (2004a,b) argue that the

causal Markov condition follows as a result of the interventionist framework along with some more basic assumptions. Since this point is important in the present context, let us rehearse those assumptions with a view to understanding how they underpin the causal Markov condition.

Hausman and Woodward begin by making an ‘error-variable idealisation’, whereby all omitted causes from the model (the error variables) are direct causes of exactly one of the variables in the model and are causally connected with exclusively that variable and its effects (cf. Pearl’s uncorrelated U_i). They then assume the possibility of a surgical intervention on some system; that is, some intervention I on X that is a completely exogenous cause of X which sets the value or probability distribution of X without regard to the other causes of X , for some set of values of I . In addition, it must be the case that I satisfies the error-variable idealisation. The third assumption is that the causal relations of the target system are modular. We stated above that modularity captures the idea of a system being constituted by distinct causal mechanisms: an intervention I on some variable X does not alter the functional relation between the putative effect Y and any of its causes that are not on a directed path from X to Y . Put more simply, if X does not cause Y , then the probability distribution of Y is unchanged when there is an intervention with respect to X . This last statement is the explicit assumption that Hausman and Woodward make. The contrapositive of this statement of modularity is simply a sufficient condition for causation in terms of manipulations: if an intervention with respect to X changes the probability distribution of some other variable Y , then X causes Y .

The fourth assumption—the determination of probabilities assumption—is that the probability distribution of each of the variables of the model characterising the system is a deterministic function of its direct causes within the model and the omitted causes (cf. Pearl’s $P(X_i) = f_i(pa(X_i), U_i)$). Finally, Hausman and Woodward assume that correlations do not arise spontaneously—there is always a causal explanation of correlations. This assumption rules out the possibility of cases in which the variation in the values of two variables X and Y given the parents of X are coordinated under interventions on X in the absence of a causal relation from X to Y . (In fact, Hausman and Woodward (2004b, p. 852) claim that the assumption of no spontaneous correlation is already implicit in the assumption of modularity, since according to modularity any change in the probability distribution of some putative effect Y as a result of an intervention on X entails that X causes Y .) The claim is then that from these assumptions it follows that the causal Markov condition holds. That is, that if some variable X does not cause some other variable Y , then Y will be conditionally independent of X given the parents of X . Or, equivalently, the contrapositive: if Y is dependent on X conditional on X ’s parents, then X causes Y .

It should be clear that the assumptions that underlie this derivation of the causal Markov condition are essentially equivalent to the physically motivated premises utilised

by Pearl to achieve the same result. But the interconnecting physical assumptions and causal conditions are employed differently within each picture of causation. According to Woodward’s picture of causation, we begin with modular (local) interventions and, given a series of constraints governing the statistical relationships between the intervention variable, the system variables and the background variables, these provide the foundation for identifying local causal relations. Once we have such relations, we can begin to establish conditional dependences and independences across a wider range of variables and then compose the set of local causal relations into DAGs or causal Bayesian networks. On this view, the casual content of Bayesian networks is grounded upon the causal interpretation of interventions.

4 The two sides of interventionist causation

Each of Pearl’s and Woodward’s pictures provides a ‘statistical architecture’ for causation—that is, in each case causation is characterised in terms of a very specific set of statistical dependence and independence conditions that must be satisfied by the relevant cause, effect and environment variables. In Pearl’s causal discovery framework these dependence and independence conditions take the form of global constraints on the Bayesian network: the marginal probabilities must factorise, from which it follows that each variable is independent of its non-descendants conditional on its parents, no causal dependence betrays a statistical independence and there are no unmeasured common causes. As we have seen, this permits us to interpret the Bayesian network causally, and the globally constrained causal Bayesian network then underpins the causal content of local interventions. In Woodward’s interventionist account of causation the dependence and independence conditions take the form of local constraints on the intervention itself: in short, the intervention renders the putative cause independent of its usual causes, is independent of the putative effect except via the cause, and is independent of any other causes of the effect. Again, as we have seen, such local constraints are sufficient to underpin the causal interpretation of Bayesian networks.

As we noted in the introduction, we can understand the discrepancy between the pictures of Pearl and Woodward as a difference in focus, rather than any substantial difference in the content of causal claims. In both pictures, the key to causality emerges from changes in conditional probability distributions under the effect of local interventions, but the constraints that endow such conditional probability distributions with causal content come, so to speak, from the outside in for Pearl and from the inside out for Woodward. We take this to be suggestive of a deeper insight into the nature of causation as it can be understood in this tradition. Rather than taking one or the other of these foci as providing a naturally more fundamental platform for understanding causation than the other, it is the interplay between the two perspectives that provides the complete picture of the causal content of the causal modelling and interventionist framework; that is, there

is a single statistical architecture to causation with contrasting but necessarily consistent global and local aspects, which are then simply two sides of the same coin. Just as observations are never devoid of a conceptual background (the well-known theory-ladenness of observation) so too is our interventionist inquiry rarely isolated from hypotheses of causal structure, or our causal modelling rarely isolated from our practical interventions on a system. The discovery of causal relations is an iterative process that requires both global and local statistical considerations to be taken into account.

A more complete picture of the interplay between the global and local perspectives is given, we think, in the work of Pearl, who explicitly employs a calculus of interventions as a means of reducing any underdetermination of graphical structures. As part of this picture, the global perspective assumes a primitive causal substrate to be responsible for the statistical dependences and independences which, when revealed, provides the structure by which we can ascertain the results of interventions. (However, Pearl's particular adherence to the causal Markov condition and faithfulness as the appropriate global statistical constraints commits him to an exclusively classical conception of causation.) On the other hand, it seems that Woodward's focus on local interventions assumes no such primitive causal substrate; indeed, according to his view, the causal structure can be established only as a result of local interventions. We can imagine local interventions to provide a focus only on the immediate surroundings of the intervention—on the variables whose statistical dependences and independences are constrained by the conditions that qualify an intervention—much as if the target causal relation were lit by a narrow spotlight with the rest of the causal structure of the system remaining in darkness (the global statistical constraints allow one to turn on the room lights).

A question arises at this stage concerning the exact way in Woodward's picture that individual causal relations identified by local interventions compose to provide an account of the causal structure of some broader causal network. As we have seen, according to Hausman and Woodward (2004a,b), the causal Markov condition arises as a result of some more fundamental assumptions concerning the modularity of causal relations, the decoupling of omitted variables and other considerations already built into the structure of local interventions, and thus it would seem that, if such a view were correct, the rules for composition emerge as a result of local interventionist constraints (not to mention rendering the global aspects of Pearl's picture of causation redundant). There is reason, however, to be suspicious of this move.

The analysis of Hausman and Woodward (2004a,b) is an attempt to decompose the causal Markov condition into a series of more straightforward and intuitive assumptions that then follow naturally from the interventionist framework. It is not clear, however, whether the interventionist framework is not itself built upon the causal intuitions we find formalised in the causal Markov condition. Consider, for instance, the assumption of no spontaneous correlation: there is always a causal explanation of correlations. As Hausman and Woodward (2004a, p.155) themselves point out, this assumption "is at least close to

the factorization claim—that the joint distribution $P(X_1 \dots X_n)$ equals the product of $P(X_i|pa_i)$ for all i —which is provably equivalent to the causal Markov condition.” This similarity should be clear, as the no spontaneous correlation assumption claims that we should always expect there to be a causal explanation of any correlation, and factorisability claims that variables with disjoint causal pasts should not be correlated. Add to this consideration the error-variable idealisation—that any causes omitted from the model are uncorrelated—and modularity—that when there is no cause connecting variables then there should not be a correlation between them—and we see that the set of assumptions Hausman and Woodward take to underlie the causal Markov condition are not clearly distinguishable from the causal Markov condition itself.

Steel (2006) makes a similar point. His claim is that the causal Markov condition of Hausman and Woodward (2004a,b) is not exactly the causal Markov condition, which states that every variable X is independent of its nondescendants given its parents. Rather, Hausman and Woodward prove a pairwise version of the more general condition. To achieve the more general derivation, Steel argues that an extra assumption is required that ensures that each error-variable is probabilistically independent of any subset of variables from the model that is causally distinct from the error-variable. This assumption is usually motivated by the generalised version of the principle of the common cause, that any set of variables amongst which there are no causal connections is probabilistically independent, which itself is entailed by the causal Markov condition. This suggests to us that the global statistical constraints emphasised in Pearl’s picture do indeed play a significant role in constraining the composition of causal relations in Woodward’s picture.⁵

Our claim is then the following: a coherent account of causality requires that the statistical constraints on the local interventions of nodes in a Bayesian network be consistent with the statistical constraints on the relations between conditional probability distributions across the Bayesian network. We could then think of, say, the precise notion of an intervention or the factorisation of marginal conditional probability distributions as individual tools, part of a wider toolkit, that provide effective strategies for identifying causal relations. These individual tools may or may not in the end be the best possible tools (whatever that means) for such a task, but effectively identifying causal relations requires that the use of each tool is consistent with the use of the others.

5 Concluding remarks

It should be abundantly clear that between the approaches of Pearl and Woodward we find a consistent notion of causation. While it has been noted previously that there is a difference in focus between the approaches, as per the opening Ismael quote, we have

⁵The decomposed assumptions of Hausman and Woodward seem to be a good example of what Schaffer (201, p.67) calls in the context of causal representations *aptness constraints*. We can imagine the causal Markov condition and faithfulness to be sophisticated kinds of these sorts of constraints.

set out here a more explicit account of what that difference constitutes: Pearl’s focus is on leveraging *global* network constraints to correctly identify local causal relations, while Woodward’s focus is on the use of *local* manipulation to identify single causal relations that then compose into global causal structures. We claim, however, that neither of these foci can be taken individually as a more fundamental platform for understanding causation—Pearl requires the possibility of local interventions and Woodward requires a global statistical framework within which to build composite causal structures. What we find then between the approaches of Pearl and Woodward is a single statistical architecture of causation with contrasting but necessarily consistent global and local aspects; indeed, we have argued that Pearl’s account of causal Bayesian networks and Woodward’s account of interventionist causation are simply two sides of the same coin.

References

- Gebharder, A. (2017). *Causal Nets, Interventionism, and Mechanisms: Philosophical Foundations and Applications*. Springer International, Cham. doi:10.1007/978-3-319-49908-6.
- Hausman, D. M. and Woodward, J. (1999). Independence, invariance and the causal Markov condition. *British Journal for the Philosophy of Science* **50**(4): 521–583. doi:10.1093/bjps/50.4.521.
- (2004a). Modularity and the Causal Markov Condition: A Restatement. *British Journal for the Philosophy of Science* **55**(1): 147–161. doi:10.1093/bjps/55.1.147.
- (2004b). Manipulation and the Causal Markov Condition. *Philosophy of Science* **71**(5): 846–856. doi:10.1086/425235.
- Ismael, J. (forthcoming). Against Globalism About Laws. In B. van Fraassen and I. Peschard (eds.), *The Experimental Side of Modeling*, University of Chicago Press, Chicago.
- Pearl, J. (2009). *Causality: Models, Reasoning, and Inference*. Cambridge University Press, New York, 2nd edition.
- Schaffer, J. (201). Grounding in the image of causation. *Philosophical Studies* **173**(1): 49–100. doi:10.1007/s11098-014-0438-1.
- Spirtes, P., Glymour, C. N. and Scheines, R. (2000). *Causation, Prediction, and Search*. MIT Press, New York, 2nd edition.
- Steel, D. (2006). Comment On Hausman & Woodward On The Causal Markov Condition. *British Journal for the Philosophy of Science* **57**(1): 219–231. doi:10.1093/bjps/axi154.

- Woodward, J. (2003). *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press, New York.
- (2007). Causation with a Human Face. In H. Price and R. Corry (eds.), *Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited*, Oxford University Press, Oxford, chapter 4, pp. 66–105.
- (2008). Mental Causation and Neural Mechanisms. In J. Hohwy and J. Kallestrup (eds.), *Being Reduced: New Essays on Reductive Explanation and Causation*, Oxford University Press, Oxford, chapter 12, pp. 218–262.
- (2013). Causation and Manipulability. In E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. plato.stanford.edu/archives/win2013/entries/causation-mani/.
- (2015). Interventionism and Causal Exclusion. *Philosophy and Phenomenological Research* **91**(2): 303–347. doi:10.1111/phpr.12095.