Manipulation and the Causal Markov Condition

Daniel Hausman and James Woodward†

This paper explores the relationship between a manipulability conception of causation and the causal Markov condition (CM). We argue that violations of CM also violate widely shared expectations—implicit in the manipulability conception—having to do with the absence of spontaneous correlations. They also violate expectations concerning the connection between independence or dependence relationships in the presence and absence of interventions.

1. Introduction. The causal Markov condition (CM) relates probability distributions to the causal structures that generate them. Given the direct causal relationships among the variables in some set **V** and an associated probability distribution P over **V**, CM says that conditional on its parents (its direct causes in **V**) every variable is independent of every other variable, except its effects. Writing pa_i for the parents of X_i and "$X_i \perp X_j/pa_i$" for "X_i and X_j are probabilistically independent conditional on pa_i":

> **CM.** For all X_i, X_j, $i \neq j$ in **V**, if X_i does not cause X_j, then $X_i \perp X_j/pa_i$.

CM is a generalization of familiar screening off conditions. CM implies that the full set of common causes screens off joint effects from one another, that direct causes screen off effects from distal causes, and that variables that are not related as cause and effect or as effects of a common cause are probabilistically independent.

CM may fail if the variable set **V** is not acyclic or omits common causes, if variables are not distinct in the right way, if variables are too coarse grained, if the probability distribution P is a mixture of distinct distri-

†To contact the authors, please write to: Daniel Hausman, University of Wisconsin, Department of Philosophy, 600 N. Park Street, Madison, WI 53706; e-mail: dhausman@wisc.edu. James Woodward, California Institute of Technology, Humanities & Social Sciences 101-40, Pasadena, CA 91125-0001; e-mail: jfw@hss.caltech.edu.

Philosophy of Science, 71 (December 2004) pp. 846–856. 0031-8248/2004/7105-0016$10.00
Copyright 2004 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/425235 Published online by Cambridge University Press

butions resulting from different causal structures, or if the population in which P holds is selected by conditioning on a common effect of variables in **V**, or is biased or unrepresentative in some other way. In what follows, we assume that none of these "defeating" conditions obtain.

Given the absence of these defeating conditions, what reasons are there to suppose that CM holds? When causal relationships are pseudo-indeterministic and the "error terms" are probabilistically independent of one another and of the other causal variables in the equations in which they appear, CM must hold. (Pseudo-indeterminism holds when the underlying causal relationships are deterministic, but for each variable X_i in **V**, some direct causes of X_i, the combined effect of which is represented by an additive error term U_i', are omitted from **V**.) However, this observation says nothing about whether CM holds in indeterministic contexts, and at least one prominent critic, Nancy Cartwright, maintains that CM is frequently violated under indeterminism. Moreover, the rationale for the independence assumptions described above is that variables are probabilistically independent if they are not related as cause and effect or as effects of a common cause, which is itself an implication of CM. Those who question CM are likely to question these assumptions. Finally, one would like some insight into whether any widely recognized features of causation make CM especially natural.

In this essay we argue for an intimate connection between a manipulability conception of causation and CM. We focus on this connection because, first, manipulation is crucial to our conception of causation and to the contrast between causation and mere correlation. When X and Y are correlated and X does not cause Y, one expects that when one manipulates X, the correlation will break down. By contrast, if X causes Y, one expects that for some range of values of X, if one is able to manipulate those values, one can thereby control the value of Y. Second, both graphical and equational representations of causal relationships assumed as background to CM are often given a interventionist interpretation. For example, theorists take the presence of an arrow directed from X to Y (indicating that X is a direct cause of Y) to mean that if X were to be appropriately manipulated while all other variables in the graph were held fixed, the value of Y would change (Glymour 1997 and Woodward 2003). Are there any reasons to expect CM to hold when "cause" is given such a manipulability interpretation?

2. Interventions, Modularity, and Causation. Some philosophers have attempted to use the connection between causation and manipulation to provide a reductive analysis of causation. We think this is a hopeless undertaking, because causal notions are required to explain what an in-

https://doi.org/10.1086/425235 Published online by Cambridge University Press

tervention is. Nonetheless, the connection between causation and manipulation can be genuinely illuminating.

The connection is not simple. It is not true that X causes Y if and only if every intervention on X results in some change in Y. Consider threshold effects: even though X causes Y, an intervention that changes the value of X from x to x^* may be "too small" to bring about a change in the value of Y. In addition, X may have causal effects on Y along different paths, which exactly cancel. In such failures of "faithfulness" (Spirtes, Glymour, and Scheines 2000), no intervention on X will change the value of Y, even though X is, in the terminology of Hitchcock 2001 and Woodward 2003, a contributing cause of Y. Moreover, some interventions on X may be outside the "domain of invariance" of the relationship between X and Y (see below). Finally, if X is an indeterministic cause of Y, then interventions on X may fail to change the value of Y, although, unless one faces one of the other difficulties just described, interventions on X should change Y's probability distribution.

There is nevertheless a plausible sufficient condition which for reasons that emerge shortly, we call MOD*: If some intervention with respect to X_i changes the probability distribution of some other variable X_j, then X_i causes X_j. MOD* is obviously *not* an analysis of causation. MOD* states only a sufficient condition, and it relies on the notion of an intervention which will be given an explicitly causal characterization.

Before turning to this characterization, we need to make our framework more explicit and take account of some complications. When causation is deterministic, a system of causal relationships may be represented by a set of equations of the form: $X_i = f_i(pa_i, U_i')$ for each variable X_i in the variable set $\mathbf{V}$; pa_i are the parents of X_i and U_i' represents the action of causes of X_i omitted from $\mathbf{V}$. Within this framework, if the value of X_i is "set" by an intervention and the equation system correctly predicts whether (and how) the values of all other variables change, then the equation system is "causally correct": it correctly describes the quantitative causal relationship between X_i and all the other variables. When an equation system is causally correct, one may derive the outcome of an intervention with respect to X_i by replacing the equation in which X_i is a dependent variable with a new equation specifying the value to which X_i was set and leaving all the other equations alone. Obviously, this procedure will work reliably only if interventions that disrupt the relations between X_i and its parents do not alter other causal relationships in the system and the equations that govern them. In other words, equations not containing the variable X_i as a dependent variable should be *invariant*—left unchanged—under interventions that set the value of X_i.

In earlier work (1999), we said that a set of equations meeting this invariance condition was "modular." The idea is if a set of equations is

https://doi.org/10.1086/425235 Published online by Cambridge University Press

causally correct, separate equations should represent distinct causal mechanisms or "modules," and if these are distinct it should be possible to interfere with each separately without disrupting the others. Modularity fits with a manipulability conception of causation: If an equation in which X_j is the dependent variable changes whenever there is an intervention with respect to X_i, then according to a manipulability view of causation there is a causal relationship between X_i and X_j, and the equation system fails to specify it correctly.

However, there is a complication: as noted above, causal relations are not invariant to *every* intervention. Consider a simple spring whose restoring force F for displacement Y is

$$F = -kY \quad (\text{for } 0 \leq Y \leq y^*). \tag{1}$$

(1) describes the spring's behavior if it is not stretched too far. If $Y > y^*$, the restoring force no longer increases proportionally to the displacement. (1) is invariant to only some interventions with respect to Y. Suppose that one replaces (1) with

$$\begin{aligned} F &= -kY \quad \text{for } 0 \leq Y \leq y^*; \\ F &= -g(Y) \quad \text{for } y^* < Y \leq y^{**}; \\ F &= 0 \quad \text{for } Y > y^{**}. \end{aligned} \tag{1'}$$

y^{**} is the point when the spring breaks. Even if (1′) correctly predicts the restoring force for springs that have not previously been stretched beyond y^*, stretching the spring beyond y^* will distort it in such a way that (1′) will no longer correctly describe the restoring force for extensions less than y^*. Moreover, once $Y > y^{**}$, $F = 0$ for all subsequent extensions.

One might try to accommodate this fact by writing down a yet more complicated set of equations (1″) incorporating information about the previous history of the spring. However, even with such temporal qualifying clauses, (1″) will still not provide correct predictions regarding all interventions—for example, (1″) will break down under manipulations of the extension under extreme heat. In systems such as the spring, we know of no causal equations that are invariant under all interventions.

Accordingly, we cannot accept Nancy Cartwright's proposal (in her criticism of Hausman and Woodward 1999) that the modularity requirement be built into the characterization of an intervention (Cartwright, 2002). She suggests that a manipulation of X_i should not count as a bona fide intervention on X_i unless it leaves undisturbed *all* causal relationships in the system of interest, except of course the relationship between X_i and its parents. Cartwright is worried by the possibility that a putative intervention on X_i (like an extension that breaks the spring) may disrupt the connection between X_i and its effects. If there is no change in an effect

https://doi.org/10.1086/425235 Published online by Cambridge University Press

X_j of X_i under such a disrupting intervention, one might mistakenly conclude that X_i does not cause X_j.

But the proper response is not to build modularity into the definition of an intervention. Doing so trivializes the connection between invariance under interventions and causal correctness (since the notion of an intervention builds in reference to the correct set of causal relationships) and consequently prevents one from using the notion of invariance under interventions (or the behavior of X_j under interventions on X_i) to provide any independent purchase on what it is for an equation linking X_i to X_j to be causally correct. Moreover, as examples like (1)–(1″) above illustrate, what most people would regard as interventions with respect to actual causal systems are invariant only to some range of interventions. Finally, the proposal requires that the notion of an intervention be relativized to a specific set of functional relationships: the same manipulation of Y may be an intervention with respect to (1) but not with respect to (1′). The proper response to Cartwright's worry that a manipulation that destroys the connection between X_i and its effects will lead to mistaken causal conclusions is to recognize that causal relationships have "domains of invariance"—ranges of interventions over which they continue to hold, even though they do not hold universally. Modularity should be understood as claiming that for each X_i there will be some range of interventions (interventions that are of the "right kind" and not "too big") that disrupt only the relationship between X_i and its direct causes and no others in the system of interest. We take this to be a substantive claim concerning the nature of causation, which is grounded in a manipulability view.

So far we have been talking only about deterministic causal relationships. When causation is indeterministic, we shall assume that for each X_i in **V**, $P(X_i) = g_i(pa_i, U_i')$, again where g_i is some function and U_i' a set of omitted causes that bears no causal relations to any other variables in **V** apart from those that result from its causal relation to X_i. That is, when causation is indeterministic, the *probability distribution* of X_i is determined by the values of its parents and U_i'. We will return to the status of this "determinism of probabilities" assumption below. Within this indeterministic framework there is a natural analogue to the modularity condition: a system of equations of the form $P(X_i) = g_i(pa_i, U_i')$ will be modular if for each variable X_i, there exists a range of interventions that disrupt only the equation linking X_i to its parents and U_i'—i.e., only the function g_i—while leaving all of the other equations in the system—all of the other functions g_j for $i \neq j$—undisturbed.

We are now ready to provide a more precise characterization of an intervention with respect to both deterministic and indeterministic relations. Let Z_i be a particular direct cause of X_i that is omitted from **V**—that is, Z_i is one of the omitted variables whose combined effect is sum-

https://doi.org/10.1086/425235 Published online by Cambridge University Press

marized by U_i'. Let U_i represent the combined effect of all remaining variables in U_i' except for Z_i. Then Z_i will constitute an (arrow-breaking) intervention with respect to the variable set **V** if and only if it meets following conditions:

1. Z_i causes X_i.
2. Z_i is not caused by any of the Xs, or other Zs.
3. Z_i does not cause any of the Us (nor, as (2) implies, any of the other Zs) and has no causes in common with any of the Us or other Zs.
4. For all X_j for $j \neq i$, if Z_i or any cause of Z_i causes X_j, then it does so only via a path passing through X_i.
5. If X_i is deterministically caused, then for some range of values of Z_i, $\mathbf{z}_i^*$, if $Z_i = z_i^* \in \mathbf{z}_i^*$, then $X_i = x_i^*$, regardless of the values of any of the Xs, Us, or Zs. If X_i is indeterministically caused, then for some range of values of Z_i, $\mathbf{z}_i^*$, if $Z_i = z_i^* \in \mathbf{z}_i^*$ $\mathrm{P}(X_i) = \mathrm{P}_i^*$ regardless of the values of any of the Xs, Us, or Zs. We shall say that $\mathbf{z}_i^*$ consists of the "on" values of Z_i. For other values of Z_i, X_i or $\mathrm{P}(X_i)$ is a function of pa_i and other omitted causes, the effect of which is summarized by U_i.

This characterization is not as complicated as it may look. 2 and 3 require that an intervention be exogenous—it is not caused by any of the other variables in the system and does not have causes in common with any of those other variables. 4 says that if an intervention on X_i causes some second variable X_j, it does so only through its influence on X_i and not in some other way. 5 says that an intervention acts as a "switch." When the intervention on X_i has an "off" value, the value of X_i, or in the indeterministic case, $\mathrm{P}(X_i)$ is determined entirely by its parents and the omitted causes whose action is represented by U_i. When the intervention on X_i is "on," the value of X_i (or, in the indeterministic case, $\mathrm{P}(X_i)$) is determined entirely by the value of the intervention, and the other causes of X_i have no influence on it. In other words, an intervention on X_i "breaks" all arrows directed into X_i except for the arrow coming from the intervention variable.

Let $\mathrm{P}_{[Z_i=\mathrm{on}]}(X_j)$ denote the probability distribution of X_j if the value of X_i were "set" by an intervention Z_i and similarly $\mathrm{P}_{[Z_i=\mathrm{off}]}(X_j)$ denotes $\mathrm{P}(X_j)$ when there is no intervention with respect to X_i. The generalization of modularity, which we call MOD* can then be expressed as follows:

> **MOD***. For all X_i and X_j in **V** ($i \neq j$) if $\mathrm{P}_{[Z_i=\mathrm{on}]}(X_j) \neq \mathrm{P}_{[Z_i=\mathrm{off}]}(X_j)$ then X_i causes X_j.

MOD* formulates a sufficient condition for causation in terms of the results of manipulations. It is also closely linked to CM. Suppose that X_i and X_j are correlated under an intervention on X_i. Given the definition

https://doi.org/10.1086/425235 Published online by Cambridge University Press

of an intervention, the correlation cannot be due to (i) some common cause of X_i and X_j or to (ii) X_j causing X_i. MOD* assumes the only other possibility is (iii) X_i causes X_j. Thus the principle that correlations don't arise spontaneously but always have a causal explanation (at least when interventions are involved) is already implicit in MOD* and in a manipulationist account of causation.

Of course, experiments, especially with a small number of subjects, often show correlations between the values of X_i produced by interventions and the values of X_j owing to accident or chance rather than because X_i causes X_j. But in such cases, one ordinarily expects that the association will disappear if the experiment is performed repeatedly with a large number of subjects.

What about the possibility that X_i and X_j remain dependent, although causally unrelated, no matter how one repeats the experiment? One response is to hold that if the relationship between X_i and X_j *looks* as though it is causal in the sense that correlations under interventions between X_i and X_j hold in (what seem to be) representative large populations, then this is all that is required for it to *be* causal. Alternatively, one may regard doubt at this point about whether X_i causes X_j as Cartesian or hyperbolic—as uninteresting as skepticism about the existence of the external world. Another response is more methodological: whether or not what is envisioned is logically possible, our experience is that inferences based on MOD* when the dependencies in question hold systematically in representative populations are highly reliable. Furthermore, if one rejects such inferences, one loses the possibility of learning anything about causal relationships from experiments—it is not as though there is some other competing principle one might employ instead. So one should accept MOD*.

3. From the Manipulability Conception to CM. One way of thinking about the connection between the manipulability conception and CM is to think of MOD* as describing what would happen in situations in which there are causal connections and intervention variables are "on" and CM as addressing what would happen in situations in which the same causal connections hold but intervention variables are "off." It is part of ordinary thinking about causation that there will be a systematic connection between these two sorts of situations. When people do experiments, they do them with the expectation that they will learn not just about how the causal agents they manipulate behave in experimental contexts, but how they behave more generally, in other contexts as well.

This connecting assumption can be expressed as follows: conditioning on the parents of a variable simulates the effect of an intervention on that variable in the sense that if X_i and X_j are dependent conditional on

https://doi.org/10.1086/425235 Published online by Cambridge University Press

pa_i, then X_i and X_j will be dependent under an intervention on X_i. This is the second premise in the argument in Section 4 below. We call it "Indication" since it claims that conditional dependence "indicates" the result of an intervention. In effect, Indication takes the idea that correlations do not arise spontaneously when interventions are "on," an idea which we found to be implicit in MOD*, and extends it to situations in which interventions are "off." CM is an implication of the conjunction of MOD* and Indication.

4. The Argument. Here is the basic outline of our argument for CM:

1. For all X_i and X_j in $\mathbf{V}$ ($i \neq j$) such that X_i does not cause X_j, $P_{[Z_i=\text{on}]}(X_j) = P_{[Z_i=\text{off}]}(X_j)$ (this is the contrapositive of MOD*).
2. $\forall$ X_j in $\mathbf{V}$ distinct from X_i, if $P_{[Z_i=\text{on}]}(X_j) = P_{[Z_i=\text{off}]}(X_j)$, then $X_j \perp X_i/pa_i$ (this is the contrapositive of the condition that we called "Indication" above).
3. If X_i does not cause X_j, then $X_j \perp X_i/pa_i$ (CM) (from 1 and 2).

We assume that 1 is relatively uncontroversial and that those who, like Cartwright, reject CM ought to oppose 2 instead. What then can be said on behalf of 2?

Suppose that the full set of parents of X_i are held fixed (that is, conditioned on) and assume that U_i' (the combined effect of the omitted causes for X_i) has no causal connections to any variables in $\mathbf{V}$ apart from those that follow from its being a direct cause of X_i. Then U_i' behaves like an intervention (in the sense we have defined) on X_i. Fixing pa_i, U_i' is an exogenous source of variation in X_i that counts as an intervention Thus,

4. If $P_{[Z_i=\text{on}]}(X_j) = P_{[Z_i=\text{off}]}(X_j)$, then $X_j \perp U_i'/pa_i$.

If one also assumes

5. If $X_j \perp U_i'/pa_i$, then $X_j \perp X_i/pa_i$.

one may derive 2. Under determinism or pseudo-indeterminism, 5 is trivial. Since X_i is a deterministic function of pa_i and U_i', if X_j and U_i' are independent conditional on pa_i, it follows that X_j and X_i are independent conditional on pa_i. Under indeterminism 5 is more controversial. We think that those who, like Cartwright, reject 2 are more likely to object to 5 rather than 4.

Here is a heuristic argument that (5) should hold under indeterminism: By hypothesis, there are two sources of variation for X_i (or $P(X_i)$): pa_i and U_i'. If one conditions on pa_i, one removes it as a source of variation. The only remaining source of variation in X_i (or $P(X_i)$) is U_i'. So the only way a correlation could arise between X_i and X_j (conditional on pa_i) is if

https://doi.org/10.1086/425235 Published online by Cambridge University Press

U_i' is correlated with X_j (conditional on pa_i). Thus if $U_i' \perp X_j/pa_i$, it should be the case that $X_i \perp X_j/pa_i$.

This heuristic argument rests on the assumption that a correlation between X_i and X_j will not arise "spontaneously"—that is, in the absence of some varying cause of X_i that is itself correlated with X_j. In the indeterministic case, where the values of X_i and X_j are not determined by their causes, this "no spontaneous correlation" assumption rules out the possibility that the variation in the values of these two variables for a fixed value of pa_i is coordinated, even though interventions on X_i are independent of X_j.

As an illustration, suppose there are two samples of radioactive material, S_1 and S_2. By bombarding S_1 with neutrons (C), one can increase the probability that the atoms in S_1 decay, but aside from this, decay in both samples is an entirely random process that depends only on the characteristics of the material in S_1 and S_2. Suppose that (C) can take one of two values—on or off—indexed by a random variable W and that one can control P(W)—that is, one can introduce or remove the neutron source. Let P(X_1/pa_1) and P(X_2/pa_2) be the probabilities for decay of an individual atom in S_1 and S_2 respectively, conditional on the characteristics of those atoms, which are summarized by pa_1 and pa_2. 5 claims that if the random variable W is independent of X_2 conditional on pa_1, then X_1 will be independent of X_2 conditional on pa_1. This condition would fail if, as we vary P(W), W remains independent of X_2 conditional on pa_1 (behavior that, given that W behaves like an intervention with respect to X_1, we would ordinarily take to show that X_1 does not cause X_2) but it is also true that X_1 is not independent of X_2 conditional on pa_1. If 5 fails in this way, the condition that we called determinism of probabilities will also fail—the joint distribution P(X_1, X_2) will depend not just on the probability of X_1 conditional on the full set of its causes (P(X_1/pa_1, W)), the probability of X_2 conditional on the full set of its causes P(X_2/pa_2) and the probability of each set of these causes. Instead there will be an additional dependence in the joint distribution, reflecting the correlation between X_1 and X_2, that is not derivable from the probabilities just described (and which indeed seems to come from thin air).

These remarks do not show that violations of (5) are impossible or that they are ruled out by an uncontroversial part of "our concept" of causation. Someone who accepts MOD* yet is convinced that violations of CM are possible has the option of denying (5) and the Indication condition (2). Nonetheless our discussion identifies just what those who reject CM are committed to denying and highlights commonplace expectations about causal relationships that are violated when CM fails to hold.

5. Cartwright's Chemical Factory. Nancy Cartwright (e.g., 2002) asks us

https://doi.org/10.1086/425235 Published online by Cambridge University Press

to imagine a chemical factory that produces two chemicals A and B by an indeterministic process P. P is a common cause of the production of A and B and the only such cause, aside from the possibility of interventions that influence the production of the chemicals directly. In particular, the production of A does not cause the production of B or vice-versa. Let X and Y be indicator variables taking the values 0 or 1 depending on whether A and B are or are not produced. Assume the state of P can also be represented by an indicator variable C. As Cartwright describes the case, $P(X/C) = P(Y/C) = 0.8$ but $P(X/C.Y) = 1$, in violation of CM.

Let Z_x and Z_y be intervention variables with respect to X and Y. In the expanded variable set including Z_x and Z_y, $pa_x = \{C, Z_x\}$, and $pa_y = \{C, Z_y\}$. Suppose $Z_x =$ on. One possibility is that X and Y remain dependent under this intervention. If so, it will be extremely tempting to conclude that X causes Y in accord with MOD* but contrary to Cartwright's description of the example. If X causes Y, there is of course no violation of CM.

Although Cartwright (2002) criticizes MOD*, we think she *requires* it if her example is to be even a prima facie counterexample to CM. Only by adopting MOD* and a corresponding interventionist interpretation of causation can one tell a consistent story about her example. Accordingly, in what follows, we will assume in accord with MOD* that when $Z_x =$ on, X and Y are independent conditional on C. However, when $Z_x =$ off, X and Y are dependent conditional on C.

Under this construal, MOD* holds but Indication (premise 2) fails, as of course does CM. Although the distribution of Y is the same when the intervention variable for X is on as when it is off (since, by hypothesis, X and Y are no longer correlated when the intervention is "on"), it is not true that $X \perp Y/C$. Moreover, of the two premises—4 and 5—that jointly imply Indication, it is, as suggested above, 5 that fails rather than 4. Construing U_x' as an intervention variable, U_x' is independent of Y, conditional on C for both the on and off values of U_x', as 4 requires, but 5 is violated.

What is odd about the example is not that certain variables are independent or conditionally independent for some values and not for others. There is nothing perplexing about this. Instead what is odd is that conditional independence relations hold or fail to hold for the same values of variables depending on *how* those values are produced. When the values of X and Y are caused by C, and the intervention variables Z_x and Z_y are "off," X and Y are conditionally dependent. When X and Y have those very same values but those values are instead caused by the "on" values of Z_x or Z_y, X and Y are conditionally independent. There is no systematic connection between the conditional independence relations that hold when interventions are on and the relations that hold when inter-

https://doi.org/10.1086/425235 Published online by Cambridge University Press

ventions are off. This violates standard expectations, described in Section 3, about the connection between the behavior of causes in experimental and nonexperimental situations.

6. Conclusion. In her discussions of CM, Cartwright claims not just that violations of CM are logically possible but that such violations "are to all appearances widespread" even in the macroscopic, non-quantum-mechanical realm. (Her discussion makes it clear that the violations she imagines are over and above any that may be due to the defeating conditions described in Section 1 above.) This is a surprising claim. The chemical factory is imaginary and Cartwright has not produced any real-life macroscopic examples in which CM fails. This suggests (we do not claim that it conclusively establishes) that violations of CM violate widely shared expectations about the relationship between causes, interventions, and probabilities. In this essay we have suggested what some of these expectations are. People do not expect spontaneous correlations, and they do expect that there will be systematic relationships between (in)dependence relationships when intervention variables are off and when they are on, so that they can use each kind of information to learn about the other.[1] As in the case of MOD*, these expectations may not be logically inviolable but they seem to be highly reliable in actual application, and they seem to be required if one is able to learn about causal relationships from nonexperimental evidence in the absence of detailed background information.

REFERENCES

Cartwright, Nancy (2002), "Against Modularity, the Causal Markov Condition and Any Link between the Two: Comments on Hausman and Woodward", *British Journal for the Philosophy of Science* 53: 411–453.

Gopnik, Alison, Clark Glymour, David Sobel, Laura E. Schulz, Tamar Kushnir, and David Danks (2004), "A Theory of Causal Learning in Children: Causal Maps and Bayes Nets", *Psychological Review* 111(1): 3–32.

Hausman, Daniel, and James Woodward (1999), "Independence, Invariance, and the Causal Markov Condition", *British Journal for the Philosophy of Science* 50: 521–583.

Hitchcock, Christopher (2001), "A Tale of Two Effects", *The Philosophical Review* 110: 361–396.

Spirtes, Peter, Clark Glymour, and Richard Scheines (2000), *Causation, Prediction, and Search,* 2nd ed. New York: MIT Press.

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

1. These claims are testable in principle: put people in situations in which such expectations are violated and see how they react. If Cartwright is correct, they will not find such violations surprising or puzzling, since they are common in real life. For some preliminary evidence (involving young children) to the contrary, see Gopnik, Glymour, Sobel, Schulz, Kushnir, and Danks 2004.

https://doi.org/10.1086/425235 Published online by Cambridge University Press