From Metaphysics to Method:

Comments on Manipulability and the Causal Markov Condition

Nancy Cartwright

LSE and UCSD

ABSTRACT

Daniel Hausman and James Woodward claim to prove that the causal Markov condition, so important to Bayes-nets methods for causal inference, is the 'flip side' of an important metaphysical fact about causation – that causes can be used to manipulate their effects. This paper disagrees. First, the premise of their proof does not demand that causes can be used to manipulate their effects but rather that if a relation passes a certain specific kind of test, it is causal. Second, the proof is invalid. Third, the kind of testability they require can easily be had without the causal Markov condition.

1 *Introduction*

2 *Earlier views: manipulability v testability*

3 *Increasingly weaker theses*

4 *The proof is invalid*

5 *MOD* is implausible*

1 Introduction

Metaphysics and methodology should go hand-in-hand. Metaphysics tells us what something is; methodology, how to find out about it. Our methods must be justified by showing that they are indeed a good way to find out about the thing under study, given what it is. Conversely, if our metaphysical account does not tie in with our best methods for finding out, we should be suspicious of our metaphysics.

Daniel Hausman and James Woodward try to forge just such a connection in their work on causation. They claim that the central characterizing feature of causation has to do with manipulability and invariance under intervention. They then use this to defend the causal Markov condition (*CMC*), which is a key assumption in the powerful Bayes-nets methods for causal inference. In their own words, '…the view that causes can in principle be used to control their effects lends support to the causal Markov condition…' (Hausman and Woodward [2003], ms p. 1). This is an important project and, to my mind, a model of the kind of thing we should be trying to do. Their first attempt to prove a link between manipulability and CMC (Hausman and Woodward [1999]) had a number of problems however (see Cartwright [2002]). Unfortunately, so too does their latest

attempt, 'Modularity and the Causal Markov Condition' (hereafter *M&CMC*) (Hausman and Woodward [2003]).

Although the connection they picture is just the kind we need between metaphysics and method, this particular link just isn't there. The first reason is that the premise they start from has nothing to do with the fact that causes can be used to control their effects. Instead it, at best,[1] lays out a sufficient condition for inferring a causal relation in ideal experimental tests; and taking such a condition as part of the metaphysics of causality, as central to the very idea of causality, smacks too much of operationalism. The most blatant of the problems with their project in M&CMC, however, is that the proof is not valid, at least under what seems to me the most natural reading of it. On a second reading, the premise is blatantly false and on a third, the proof is again invalid.

I shall explain the problems with the proof after a review of the switch they have made in their work from taking control, or manipulability, to taking a sufficient condition for inferring a causal relation in an experimental test as the starting point. The final discussion will focus on cases of probabilistic causality. When causes can act probabilistically, CMC will be violated in any case where causes produce by-products in tandem with their main effects. Hausman and Woodward maintain that causes can't do that. I shall defend my view that there is nothing to stop them from doing so.

First a definition and some notation. The causal Markov condition is formulated relative to a population Φ, a set of random variables $\mathcal{V}$ on that population, a set of random variables $\mathcal{U}$ representing omitted causes of features represented in $\mathcal{V}$ sufficient in combination with the variables in $\mathcal{V}$ to fix the values (or, for indeterministic cases, the

chances) of every variable in $\mathcal{V}$, a probability measure P over $\mathcal{V} + \mathcal{U}$, and a directed graph G of the causal relations among the features represented by variables in $\mathcal{V} + \mathcal{U}$[2]:

> **CMC**: Φ,$\mathcal{V}$,$\mathcal{U}$,P,G satisfy CMC iff for all X_i,X_j, $i \neq j$, in $\mathcal{V}$, if X_i does not cause X_j, then X_i and X_j are probabilistically independent conditional on pa_i (i.e. $X_i \perp X_j/pa_i$),

where *pa_i* is the set of direct causes, or parents, of X_i relative to $\mathcal{V}$ and G .

As to notation, throughout I shall use *$Xc \rightarrow Y$* to mean that X causes Y and *X c= f(...)* to indicate that the factors in the function on the right-hand side cause those on the left and that the functional equality holds, where in both cases generic-level as opposed to singular causation is intended. I shall denote members of $\mathcal{V}$ by *X_i* or *Y_i*, values of variables by lower case versions of the letter representing the variable, and a member of $\mathcal{U}$ that causes X_i by *U_{ij}*. Following Hausman and Woodward, I shall use $U_i{}'$ to represent the net causal effect on X_i of a minimal set of omitted causes of X_i that in combination with the parents of X_i are sufficient to fix the value (or for indeterministic cases, the chance) of X_i. *$A \perp B/C$* means that A is probabilistically independent of B conditional on C.

Hausman and Woodward treat CMC for purely deterministic causality and for purely probabilistic causality in one fell swoop. I shall divide my discussion to focus on different aspects of the proof.

2 Earlier views: manipulability v testability

Hausman and Woodward have long defended the idea that *modularity* is a characterizing feature of causality and this term appears in the title of the paper with the new proof – 'Modularity and the Causal Markov Condition'. In much of their earlier work modularity was intimately connected with claims that at the heart of causation is the idea that something is a cause just when it can be used to control the putative effect. I want here to review the earlier ideas to make clear that their new proof does not link CMC with modularity as we first saw them talking about it. Nor do they claim so when they write down their central premise in M&CMC. Still, it is easy to be misled since they retain the earlier language as well as a number of the earlier slogans, such as the one quoted in section 1 above claiming that the fact that causes can be used to control their effects supports CMC.

I shall not start by defining modularity because I think some of the arguments in their earlier work, including the paper where the earlier proof appears, speak to a somewhat different thesis than the one they formally state as **MOD** in the earlier proof.[3] Rather I shall describe two motivations for modularity we can find in their work, motivations that lead to different conditions.

1. *Manipulability*: It is essential to causality[4] that causes can be used to manipulate their effects.[5] So (roughly)…

(Xc➔Y) ➔ there is some way (they call it an 'intervention' or sometimes a 'manipulation') to change X so that Y changes.

Both 'intervention' and 'manipulation' suggest human agency and indeed for many philosophers that has been important. This, however, is not part of Hausman and Woodward's programme. What they require is not that a cause be manipulable by us in the right ways but merely that it be possible that the cause *vary* in the right ways, whether we vary it or not. This is a theme familiar from the literature on natural experiments – i.e., situations in which one factor varies naturally, without our help, in just the right way to count as a test for causality. Hausman and Woodward are explicit that human agency is not required. Nevertheless 'manipulation' and 'intervention' are the words they regularly use rather than a more neutral description in terms of variation. So we must be careful to focus on the definitions themselves and not the labels.

Even with this understanding of what 'manipulability' means, the condition seems far too strong. *If* a cause can vary in the right way, then (for the most part[6]) we can expect its effects to change in train. But there is no guarantee that such variation is always possible.[7]

2. *Testability*: In discussing the chemical factory example described below in Section 9, Hausman and Woodward take it to be an advantage of their view that it allows one 'to disentangle different possibilities concerning the causal structure of the situation. If one thinks of the example as one in which the effects cannot, even in principle, be separately interfered with, the example does not really have a common cause structure, but is rather

one in which a single mechanism links [the putative common cause] to X [one putative effect] and Y [the other putative effect]' (Hausman and Woodward [2003], ms p.14). Thus, it seems, the case would not have a common cause structure if there were no interventions possible on X and Y, at least in principle, that would allow us to determine that it does, and to do so by showing that X does not cause Y nor the reverse. This suggests that they want to require that for every possible causal connection, Xc → Y, there should be (at least in principle) an intervention on X that would show whether it holds or not.

This suggestion is supported by the kinds of arguments about causal mechanisms that they repeatedly offer in defence of modularity. Each causal principle is to represent a separate mechanism for the production of some given effect X. But there is no separate mechanism for X unless it is possible to intervene on X without changing any other causal principles. So again, it looks as if for every possible effect X there must be a possible intervention, and presumably this intervention should leave P(Y) unchanged if X does not cause Y since the intervention is supposed to have no effect on any other mechanisms, either to add, subtract or change them. Again this demand is tantamount to a condition that each possible causal claim be testable, and testable by what I shall call 'experiment': intervene and see what happens.

So, the requirement of *testability by experiment* provides a second distinct way to formulate a modularity requirement. Using the notation and formulations of M&CMC, we have:

For all $\Phi, \mathcal{V}, \mathcal{U}, P, G$, and X_i in $\mathcal{V}$, there is a Z_i in $\mathcal{V} + \mathcal{U}$ such that Z_i is an

intervention on X_i and $\forall X_j \neg (X_i c \rightarrow X_j)$ iff $P(X_j/Zi=on) = P(X_j/Zi=off)$.[8]

This requires explanation.

Z_i in $\mathcal{V} + \mathcal{U}$ is **an intervention** on X_i relative to $\Phi, \mathcal{V}, \mathcal{U}, G$ iff

(i) Z_i causes X_i on G

(ii) Z_i is not caused on G by any of the other variables in $\mathcal{V} + \mathcal{U}$

(iii) Z_i does not on G cause any members of $\mathcal{U}$ and has no causes in common with

any members of $\mathcal{U}$ or other Z′s on G

(iv) For all X_j, $j \neq i$, if Z_i or any cause of Z_i causes X_j on G, then it does so only via

a path passing through Z_i and X_i first

(v) If X_i is deterministically caused on G,P, 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 other members

of $\mathcal{V} + \mathcal{U}$. If X_i is indeterministically caused, then for some range of values of Z_i,

$\mathbf{z_i}^*$, if $Z_i = z_i^*$ in $\mathbf{z_i}^*$, $P(X_i) = P_i^*$ regardless of the values of any other members of

$\mathcal{V} + \mathcal{U}$. For other values of Z_i, X_i or $P(X_i)$ is a function of pa_i and members of $\mathcal{U}$.

The values in $\mathbf{z_i}^*$ are designated as the *on* values for Z_i. So the condition says that for

every variable X_i there always is an intervention and that the probability of any other

variable X_j changes when X_i is intervened on iff X_i causes X_j.

What is important to notice is that testability is stronger than manipulability in two

ways:

- Testability requires that there exists an 'intervener'/'manipulation' for *every* factor, not just for causes.
- Under testability, manipulating a cause changes its effects; but also manipulating *non-causes* of a factor does *not* change it.

I think that Hausman and Woodward's views on causal mechanisms and interventions commit themselves to something like testability by experiment. As an attempt to motivate modularity, testability also has the advantage over manipulability that it has the arrows of implication going the right way. Manipulability says that if $X \rightarrow Y$ then intervening on X changes Y. The assumption they call *MOD**, which is the premise in their argument for CMC, instead requires that if $\neg(Xc \rightarrow Y)$ then intervening on X does not change Y, which does follow from testability by experiment. To the extent that I am right, this is extremely restrictive: it not only requires that causal relations, in order to be causal, must each be ascertainable by us, but moreover that they be ascertainable by one specific method among the many that we use (like various 'mark methods'). This is operationalism pushed beyond its limits.

These two considerations lead me to

Conclusion 1: Modularity in the form of either the manipulation or the testability thesis is too strong a condition to characterize causality.

3 Increasingly weaker theses

For the latest proof Hausman and Woodward do not start from testability but from a far weaker assumption about the metaphysics of causality. Why? Testability tells us that for every $\mathcal{V}$ and every X_i in $\mathcal{V}$ *there is* an intervention Z_i for X_i and Z_i changes X_j iff $X_ic{\rightarrow}X_j$.[9] This is too strong for at least two reasons:

> There isn't always such a Z_i. (This is my explanation; they don't themselves say this.)

> The 'iff' is too strong. If X_i both causes and prevents X_j then X_j need not change as X_i does. (They do say this, though, as I argue in section 4, I do not think they need to.)

What they propose instead in M&CMC is this: ***MOD**** 'says that when X_i does not cause X_j, then the probability distribution of X_j is unchanged when there is an intervention with respect to X_i' (Hausman and Woodward [2003], ms p. 3). So, roughly, (for a more precise statement, see section 5)

if Z_i is an intervention for X_i *then* $(X_j$ or $P(X_j)$ changes under $Z_i) \rightarrow (X_ic{\rightarrow}X_j)$

We should note that this gives up on

➢ The claim that the possibility of full testability is necessary for the applicability of causal concepts

➢ The claim that it must be possible to use a cause to manipulate its effects

and it does so in two ways

- It is no longer necessary that an intervention on X_i exist in order for X_i to cause some other factor.

- It does not require that manipulating a cause changes the effect but rather that if X_i *does not* cause X_j then manipulating X_i will *not* change X_j.

MOD* (or rather some more precise version of it as I discuss below) is the premise in Hausman and Woodward's new proof of CMC. So it seems they do not link *manipulability* with the causal Markov condition, but at best only a claim *about one test that can guarantee that a causal relation holds.* Nor do they deny this: in M&CMC they conclude 'The causal Markov condition is a doppelganger for *invariance to intervention'*(Hausman and Woodward [2003], ms p 7, my italics). Still, they call the section with the proof, 'Causation and *Manipulation'* (my italics) and begin it with

When X causes Y and one can intervene to change the value of X, one can use one's knowledge of the causal relation to influence the value of Y…This is an extremely important feature of causation. One way to formulate a connection between causation and manipulability… is to say that if an intervention with

respect to X_i changes the probability distribution of some other variable X_j, then X_i causes X_j. (Hausman and Woodward [2003], ms p. 2.)

Given that their central premise is MOD*, their proof may connect something about manipulability with causation but

> ➤ The proof in M&CMC does not connect the claim 'Causes *can always* be manipulated to affect their effects' with causation
> ➤ Nor the weaker claim '*If* a cause can be manipulated (in the right way), the effect is changed'.

That's because MOD* says that if we manipulate a factor that is *not* a cause of another, the other does *not* change.

From this consideration and others in this section and the last I draw

Conclusion 2: The premise (MOD*) in their proof is not *manipulability* but at best one **test** that, **if** it can be applied and **if** it is passed, can guarantee that one factor causes another.

4 The proof is invalid

There are two points we need to beware of –

> ➢ The earlier Hausman and Woodward proof used the strong premise that for every X_i in $\mathcal{V}$ *there is* an intervention Z_i and manipulating Z_i leaves X_j unaffected if $\neg(X_i \text{ c}\rightarrow X_j)$. The earlier proof didn't work. The new proof has a *weaker* premise. How can it work?

> ➢ One would think that whether the probability of a non-effect of X_i, X_j, is left the same under an intervention on X_i will depend on whether the intervention on X_i is *probabilistically dependent on* any causes of X_j. Such dependencies are often prohibited by definition of 'intervention'. But not so for Hausman and Woodward. How can they get by without this?

Let me recall a well-known result:

> For any $\mathcal{V}$, $\{U_i'\}$ are independent of each other in all combinations $\rightarrow$ CMC.

So when are the U''s independent? Here is one common hypothesis:

> **CM1:** factors that are not causally connected are independent in all combinations. (X is causally connected with Y iff Xc$\rightarrow$Y or the reverse or they have a common cause.)

If $\mathcal{V}$ is causally sufficient, the U_{ij}'s will be causally unconnected and hence given CM1 independent in all combinations. This rules out 'brute correlations' that have no causal explanation, like

1. Elliott Sober's case of bread prices in Britain and sea levels in Venice (Sober [2001]).
2. Any case with time trends.

In their new proof Hausman and Woodward do **not** assume CM1. So how do they rule out probabilistic dependencies that are incompatible with CMC? They think they can do so using MOD* plus two other assumptions, where more precisely stated

$\Phi,\mathcal{V},\mathcal{U},$P,G satisfy ***MOD*** iff for every X_i in $\mathcal{V}$ and every intervention Z_i in $\mathcal{V}+\mathcal{U}$ on X_i, $Z_i \perp X_j$ for any X_j such that $\neg(X_i c \rightarrow X_j)$.

Notice that as I have written it, MOD* is a condition that a system might satisfy, not a claim. So too is CMC. I shall be concerned about what claims Hausman and Woodward want to assert. One claim that many favour is that any representative causally sufficient system, $\Phi,\mathcal{V},\mathcal{U},$P,G, satisfies CMC. Hausman and Woodward say, 'We shall show that MOD* …[and some other assumptions] … imply CMC.' (Hausman and Woodward [2003], ms p. 3) The most natural reading of this is that any system that satisfies MOD* plus the other assumptions satisfies CMC, and this is what I shall suppose they mean.

The two other assumptions for the case of determinism are

A. '[A]ll the variables in V are distinct, … we are dealing with the right

variables, and … selection bias and other sources of unrepresentativeness…are

absent' (Hausman and Woodward [2003], ms p. 2)

B. $\mathcal{V}$ is causally sufficient.

So,

H-W claim₁: for any Φ,G,$\mathcal{V}$,$\mathcal{U}$,P, *if* A., B. and MOD* hold for Φ,G,$\mathcal{V}$,$\mathcal{U}$,P, *then*

CMC holds for Φ,G,$\mathcal{V}$,$\mathcal{U}$,P.

Here is how their proof proceeds:

- Define 'intervention' so that interventions on X_i are causally unconnected with U_j'
  if X_i does not cause X_j

- Show that in a certain subpopulation – the subpopulation where pa_i is fixed – U_i'
  satisfies the definition of an intervention.

- Use MOD* to claim that in this population $U_i' \perp X_j$; i.e. $U_i' \perp X_j/pa_i$

- It follows, they say, that $X_i \perp X_j/pa_i$

But the proof must be invalid since there are cases that satisfy the premises but where

CMC fails. Consider Graph 1 for some population Φ that satisfies A., B. and MOD* and

for which X_i c= $a_iY + U_i'$ ($a_i \neq 0$) and for which U_1 and U_2 are dependent conditional on

Y. This system is inconsistent with CMC. (A dotted line indicates probabilistic

dependence; a dotted line with a *y* through it, dependence conditional on Y; a dotted line with a slash through it, independence. In this example it makes no difference whether U_1' $\perp U_2'$.)[10]

Graph 1
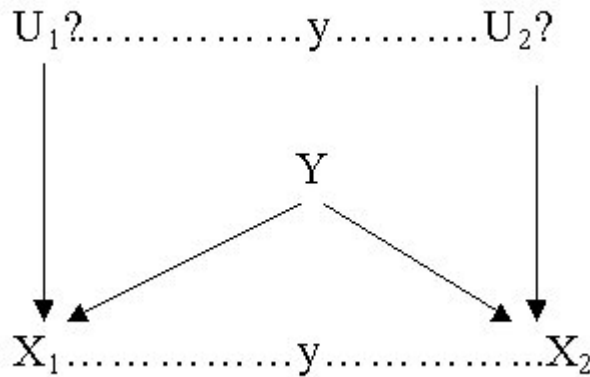
In Graph 1 MOD* is satisfied: there is no factor in $\mathcal{V}+\mathcal{U}$ that sets the value of X_i, hence no intervention, so MOD* holds vacuously.

It is worth rehearsing just why the proof is invalid. The theorem we wish to prove says that if a given population Φ and a given probability P^{11} satisfy MOD* then *that* population and *that* probability satisfy CMC (for X_i, X_j such that $\neg(X_i \, c{\rightarrow}X_j)$). Hausman and Woodward's argument establishes that if MOD* for some *different* population Φ' and *different* probability measure P', then CMC holds for Φ. (The population Φ' is the subpopulation in which $p\alpha_i$ is fixed and $P'(\ldots) = P(\ldots|p\alpha_i)$.
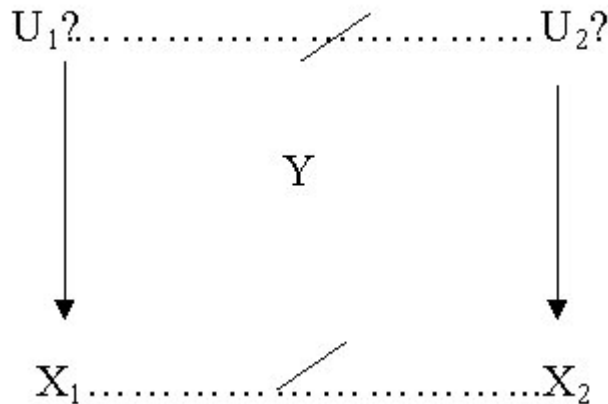
We can thus draw

Conclusion 4: The proof is invalid under the most natural reading of Hausman and Woodward's claim.

5 MOD* is implausible

We can also use Graph 1 to illustrate how strange the condition MOD* is, independent of its connection with CMC. Compare Graph 1 and Graph 2.

Graph 2

For Graph 2 suppose, as with Graph 1, that X_i c= $a_i Y_i + U_i'$ and A. and B. are satisfied. But for Graph 2 imagine that $a_i = 0$ and that $U_1' \perp U_2'$ as MOD* requires. So, MOD* allows U_1 and U_2 to be probabilistically dependent in Graph 1 but prohibits it in Graph 2.

That seems to require a completely ad hoc distinction between the two cases. Suppose we start with a situation appropriately represented by Graph 1, with U_1' and U_2' probabilistically dependent. Consider a situation identical with this except that Y's influence on X_1 and X_2 is just slightly less (i.e. a_1 and a_2 are slightly smaller). MOD*

17

does not prohibit this new situation either. Now consider a series of situations in each one of which Y's influence on the X's is smaller than in the one before. Still MOD* does not prohibit the U''s from being dependent. This is true no matter how small Y's influence on the X's becomes, so long as it has any influence at all. But as soon as the influence disappears altogether ($a_1 = 0 = a_2$), suddenly under MOD* the U''s must be independent. What is responsible for this sudden jump?

We may even suppose that the diminutions of Y's influence occur across time in the very same physical system. Gradually Y's powers to influence X_1 and X_2 give out. What would ensure, when Y's influence finally disappears altogether, that suddenly U_1' becomes independent of U_2'? I see nothing that could.

Here is an example (or rather, a caricature of an example). Suppose Elliott Sober is correct that bread prices in England are probabilistically dependent on Venetian sea levels. We can suppose that the real levels of these two variables in combination with the measurement apparatuses employed (call this combination U_1 for sea levels and U_2 for bread prices) are each a central cause of the respective measured values of the levels (X_1 and X_2); presumably so too will be the skill of the persons taking the measurements. For the sake of an example let us suppose that there is one team of experts that make both such measurements and that every ten years more and more automated technology is introduced in both places so that gradually the results depend less and less on the skills of the measurement team(Y). We can suppose that U_1 and U_2 are probabilistically dependent because by hypothesis bread prices and sea levels are dependent. This is consistent with MOD* so long as skills matter. But as soon as the measurement process becomes fully automated and the skills of the team have no influence on the measured

values, suddenly bread prices and sea levels, which were dependent until then, must become independent if MOD* is to be satisfied. I don't see why this kind of thing should happen.

Of course if we assume CM1, bread prices and sea levels will not be dependent in the first place. But recall that Hausman and Woodward do not assume CM1. And that is all to the good given their overall programme because, given CM1 and their other assumptions, CMC follows without assuming MOD*, so no argument is at hand that MOD* supports CMC. MOD* is supposed to replace CM1 and provide an independent basis for CMC. Even if the proof were valid, I do not think that this would be a very sensible basis since, as I have just been arguing

Conclusion 5: MOD* is highly implausible unless dependencies between causally unconnected quantities are already ruled out in the first place.

6 Two alternative claims and their defects

Let us try some other formulations of Hausman and Woodward's claim to see if they fare better. For their proof they need MOD* to hold in the specific population in which the parents of X_i take a fixed value. Perhaps then they intend that MOD* should hold in *every* population and hence in the requisite one:

H-W claim₂: *if* MOD* holds for every $\Phi, \mathcal{V}, \mathcal{U}, G, P$ such that $\mathcal{V}$ is causally

sufficient relative to G and P, *then* CMC holds as well for every $\Phi, \mathcal{V}, \mathcal{U}, G, P$ such

that $\mathcal{V}$ is causally sufficient.

Given the antecedent, it is true that for any population, in the subpopulation where the

parents of X_i take fixed values, $U_i' \perp X_j$; i.e. for every population, $U_i' \perp X_j / pa_i$. The

consequent then follows that CMC holds for every population. Graph 1 is no longer a

counterexample, since by inspection we can see that there is a population – the

subpopulation of Φ picked out by fixing a value for Y – for which MOD* is violated; this

is ruled out by the antecedent of the reformulated claim.

But the antecedent for this formulation is altogether too strong: it does not hold for a

vast array of perfectly ordinary situations, including a host of ones in which CMC is

satisfied. Consider, for example, a population Φ with probability measure P in which

(where causes on the right-hand-side)

$Y \leftrightarrow X_1$ or X_2

$X_1 \leftrightarrow U_1'$;

$X_2 \leftrightarrow U_2'$;
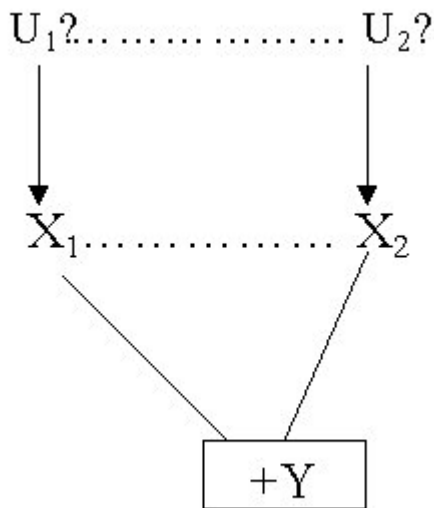
$U_1' \perp U_2'$;

$\neg(X_1 c \rightarrow X_2)$

$P(U_1') = P(U_1'/U_2') = P(U_1'/U_2' \& X_2 \& Y) = r \neq 1.$

where all the variables are dichotomous. For this population CMC holds.

Consider next a second population, Φ' – the subpopulation of Φ picked out by $+Y$. In this subpopulation U_1' sets the value of X_1, but $P(U_1'/\neg X_2\&+Y) = 1 \neq P(U_1'/X_2\&+Y)$. So $\neg(U_1' \perp X_2)$ in Φ', as illustrated in Graph 3.

Graph 3 for subpopulation $+Y$:

Or look again at Graph 2 and consider the subpopulation in which $\neg X_1 \vee X_2$. In this population U_1' is still an intervention on X_1 and X_1 still does not cause X_2, yet $P(X_2/U_1'=\text{on}) = 1 \neq P(X_2/U_1'=\text{off})$.

It is, however, almost certain that Hausman and Woodward do not wish to formulate their claim in this way. After all, the populations in my examples are unrepresentative relative to the larger populations from which they are drawn, and we see by condition A. that in their proof of CMC they assume that 'selection bias and other sources of

unrepresentativeness' are absent. Certainly my subpopulations suffer from 'selection bias'. So let us try instead

> *H-W claim₃*: *if* MOD* holds for every Φ,$\mathcal{V}$,$\mathcal{U}$,G,P such that $\mathcal{V}$ is causally
>
> sufficient relative and Φ is representative, *then* CMC holds as well for every
>
> Φ,$\mathcal{V}$,$\mathcal{U}$,G,P such that $\mathcal{V}$ is causally sufficient and Φ is representative.

The antecedent in this formulation is more plausible. But it undermines the argument that Hausman and Woodward wish to make in establishing the consequent. The subpopulations selected by fixing values of pa_i are themselves unrepresentative, and it is just these populations in which MOD* must hold if CMC is to be deducible in the manner they suggest. There is a central unresolved issue about how to define 'selection bias' and 'unrepresentative'. I myself think that it is very difficult to do for purposes of defending CMC in general. In this case in particular I see no promise for defining it in a way that is not ad hoc and yet counts all unrepresentative subpopulations as biased except those selected by pa_i for each X_i is any variable set we may wish to consider.

I am thus led to

> Conclusion 6: Of the two alternative plausible readings, the first claim has a
>
> blatantly false premise and the second has no valid argument to support it.

7 A true claim and a valid argument

A more direct approach would be to formulate the thesis to say explicitly what is required for Hausman and Woodward's proof:

H-W claim$_4$: for every $\Phi, \mathcal{V}, \mathcal{U}, G, P$, *if*

(i) for all X_i and all assignments of values, pa_{ik}, to the parents of X_i in G, ($\Phi_{ik}, \mathcal{V}, \mathcal{U}, G(pa_{ik}), P(pa_{ik})$ satisfies MOD*)

(ii) $P(pa_{ik})(--) = P(--/pa_{ik})$ and $G(pa_{ik}) = G$

(iii) $\mathcal{V}$ is causally sufficient

then $\Phi, \mathcal{V}, \mathcal{U}, G, P$ satisfies CMC,

where $G(pa_{ik})$ is a graph of the causal relations over $\mathcal{U} + \mathcal{V}$ in the subpopulation of Φ in which the parents of X_i take the values pa_{ik}, and $P(pa_{ik})(--)$ is the probability distribution over $\mathcal{U} + \mathcal{V}$ in that same subpopulation.

H-W claim$_4$ is true and the argument that Hausman and Woodward give in M&CMC shows that it is valid. But it does not gain Hausman and Woodward what they want – a route from manipulability/testability to CMC, for three reasons:

1. Claim$_1$ – the claim Hausman and Woodward seem to make, that any representative causally sufficient system that satisfies MOD* also satisfies CMC – is an interesting and surprising claim. Claim$_4$ is not. It tells us that it a very special set of unrepresentative

populations, all subpopulations of Φ, satisfy MOD* then Φ will satisfy CMC. Now this may seem to be some gain, but I am afraid that it is very little, and it is certainly not the gain I had hoped for in connecting metaphysics and methods. For we have little reason for accepting that the premise should be true of any given population than we would for expecting CMC in the first place. We could think that these very special populations satisfy MOD* because all populations do, and they do because manipulability is essential to causality. We have seen that does not work because MOD* rests on testability not manipulability.

So what if we were to suppose that testability is the form MOD* itself is essential to causality? This still gets us nowhere in the proof for it takes us back to Claim$_3$, and we have seen the problems with the premise in Claim$_3$. It is false that MOD* holds for all populations: the unrepresentative picked out by +Y in Graph 3 shows that. But if the premise is restricted to all representative populations, the argument does not go through. For a valid proof we need to suppose that MOD* applies for just the right special set of unrepresentative populations. I do not find any independent reason for that in Hausman and Woodward's discussions.

2. The problem pointed out in section 6 still arises. $\mathcal{V}$ is causally sufficient but we do not presume from this that the U′′'s are independent. Nor do we suppose CM1 to ensure they are independent. That is, they are not independent because they are causally unconnected – that it seems is not enough. But when we add that they set the values for quantities represented in $\mathcal{V}$, that is enough. But why?

3. The claim does not after all connect testability with CMC. Rather, it lays down very strong constraints on the populations, variable sets and graphs for which CMC is derived, and these constraints are strong enough to ensure both testability and CMC. This is exactly the same kind of problem that beset their earlier proof. We have a set of constraints *C*; C implies testability and C implies CMC. Of course by logic then, in C, testability implies CMC. But that is because in C, anything implies CMC. It is the constraints that imply CMC, not testability. In this case the constraints are conditions (i) – (iii) in the antecedent of H-W claim$_4$.

But isn't constraint (i) itself an assertion of testability, and the inclusion of constraint (i) is essential to the truth of claim$_4$, as we all admit? No, constraint (i) is not a reasonable assertion of testability: it guarantees testability, but is itself stronger and stronger in just the way necessary to guarantee CMC.

Here is what I would take instead to be a reasonable statement of testability:

$\mathcal{V}$ is *c-testable* in Φ relative to $\mathcal{U}$,G,P iff for all X_i in $\mathcal{V}$, there is an intervention Z_i

in $\mathcal{V}+\mathcal{U}$ such that for all X_j [$\neg(X_i c \rightarrow X_j) \rightarrow$ P(X_j/Zi=on) = P(X_j/Zi=off)].

I call this *c*-testability to stress that it is only one kind of testability – the kind we identify with a controlled test. As discussed in section 4, I myself would want to make the condition on the probabilities both necessary and sufficient for testability; but I do not do so here in order to stay as close as possible to Hausman and Woodward's formulations.

Notice how c-testability for Φ,$\mathcal{V}$,$\mathcal{U}$,G,P differs from MOD*. In the first place, c-testability requires that *there be* an intervention for every variable. On the other hand, it

does not require that everything that counts as an intervention on Hausman and

Woodward's definition should satisfy the independence assumption, merely that each

variable *has* an intervention that does so. Hence nothing about c-testability automatically

forces the U''s to satisfy the requisite independence assumption. This is for the reasons I

have rehearsed. In Hausman and Woodward's scheme, we do not assume that a factor's

being causally unconnected with others in the right ways is sufficient for guaranteeing the

independence assumption; adding that that factor sets the value for a variable in $\mathcal{V}$ does

not seem to add any reason for it to do so. On the other hand, if there is such a factor for

each variable, then any hypothesis about one variable in $\mathcal{V}$ causing another can be

tested.[12]

    The point now is that MOD* (i.e. (i)), (ii) and (iii) guarantee c-testability as well as

CMC. But c-testability in combination with (ii) and (iii) does not guarantee MOD*, nor

CMC. That's because testability does not require that intervention be via a U' – it just

requires there be some intervention for each variable, and that is compatible with the U''s

not being mutually independent. Graph 4 shows a particularly simple case:

Graph 4
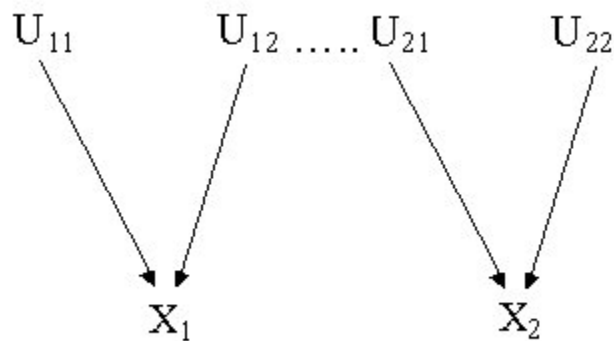
The equations for the population Φ in Graph 4 are

$$X_1 \text{ c} = U_{11} \vee U_{12}$$

$$X_2 \text{ c} = U_{21} \vee U_{22}$$

with $P(i_{11}U_{11}, i_{12}U_{12}, i_{21}U_{21}, i_{22}U_{22}) = P(i_{11}U_{11})P(i_{21}U_{21})P(i_{22}U_{22})$ if $i_{12} = i_{21}$, and $= 0$ otherwise, where $i_{jk} = +, \neg$ and $P(i_{12}U_{12}) = P(i_{21}U_{21})$.

In this case U_{11} is an intervention on X_1: conditions (i) – (iv) in the definition of intervention are met by inspection and if U_1 occurs – call that 'on' – , X_1 occurs no matter what values other variables take, so (v) is met as well. Similarly U_{22} is an intervention on X_2. Also, $U_{11} \vee U_{12}$ is an intervention on X_1 and $U_{21} \vee U_{22}$ is an intervention on X_2. $P(X_2/U_{11} \text{ is on}) = P(X_2/U_{11} \text{ is off})$ and $P(X_1/U_{22} \text{ is on}) = P(X_1/U_{22} \text{ is off})$. But

$P(X_2/U_{11} \vee U_{12} = \text{on}) \neq P(X_2/ U_{11} \vee U_{12} = \text{off})$ and $P(X_1/U_{21} \vee U_{22} = \text{on}) \neq P(X_1/U_{21} \vee U_{22})$.

So $\mathcal{V} = \{X_1, X_2\}$ is *c-testable* in Φ relative to $\mathcal{U} = \{U_{ij}\}, G, P$. Conditions (ii) and (iii) of the antecedent of H-W claim₄ are met as well. But condition (i) of that claim is not met and correlatively, CMC fails. C-testability obtains without the strong assumption needed for the true H-W claim and without CMC.

The claim I have formulated as H-W claim₄ is the only one I have been able to construct that makes their basic argument valid. If I am right that that is the only claim supported by their argument, then…

Conclusion 7: Hausman and Woodward can, using their basic ideas, produce a

true claim and a valid argument. But their argument does not show that testability

implies CMC; rather the constraints they need imply both testability and CMC;

without these constraints, c-testability does not imply CMC.

8 Indeterminism

So far I have discussed only the deterministic case. For indeterminism we need more

because in the probabilistic case a cause may produce a product and a by-product – i.e.

two effects in correlation – and in this case the causal Markov condition will be violated.

I have suggested for instance that a factory might produce an unwanted pollutant as a

side-effect during a purely probabilistic process that produces a desired chemical. In my

comments on Hausman and Woodward's proof I represented this example thus:

$$X_1 \ c = \alpha_1 Y + U_1$$

$$X_2 \ c = \alpha_2 Y + U_2$$

$$P(+\alpha_1) = .8 = P(+\alpha_2)$$

Here Y is the presence of the chemical factory process; X_1, the presence of the chemical;

X_2 the presence of the pollutant; α_1 and α_2 the operation of the chemical factory process

to produce the chemical and the pollutant respectively;[13] and '[U_1] and [U_2] each satisfy

28

the requirements of an intervention' (Cartwright [2002], p 436).[14] Since the U's satisfy

the definition of an intervention, $U_1 \perp X_2$ and $U_2 \perp X_1$, unconditionally and conditional

on Y.

In the example there is a l00% correlation between the presence of the chemical and

that of the pollutant and this correlation remains even when we condition on Y. The

reason for the correlation is that Y produces the two in tandem; it produces one if and

only if it produces the other (though any other correlation between 0 and 1 could be

possible as well). The correlation need not confuse us about what is going on. Since the

U's satisfy the criteria of an intervention, it is easy to test that the chemical is not causing

the pollutant, nor the reverse; and supposing that Y can be intervened on as well, it is

easy to test that the chemical process is causing both.

Hausman and Woodward maintain that this kind of case is impossible, at least at the

macrolevel. The issue is about $P(\alpha_1 \, \alpha_2)$. Can it, for instance, equal $P(\alpha_1)$, so that the

pollutant is a byproduct of the chemical – it is produced iff the chemical is produced? If

causation must be deterministic, this can easily happen but then CMC will not be violated

because all the relevant probabilities will equal one. But we had best not assume that

causality must be deterministic or we won't be able to say that what causes us to see the

stars is the emission of photons that occurred on them long ago. So what happens when

causation is probabilistic?

Hausman and Woodward maintain that it is impossible in this case for a cause to

produce its effects together – it must produce one effect independently of the other. They

argue that this is assumed on all standard accounts of causation. I do not agree. What

kinds of things do we expect of causation in our various standard accounts? Here are a

few: a)Causes should make their effects happen. Y does that for both X_1 and X_2. b) In the nice cases where all probabilistic dependencies can be derived from the causal laws operating, MOD* should be satisfied. And it is. c) In many situations if we put a mark on the cause we expect to find a trace of the mark on the effect. There is no reason to think that we cannot mark Y and find a trace later on both X_1 and X_2. So causation in this case has a great many of the features we expect of it.

If causes can produce their effects in tandem, CMC is violated. To prove CMC, Hausman and Woodward rule this possibility out directly with a premise they call '*no spontaneous correlation*':

for every $\Phi, \mathcal{V}, \mathcal{U}, G, P$ and for every $X_j \, \varepsilon \, \mathcal{V}$ distinct from X_i, if $X_j \perp U_i'$, then $X_j \perp$ X_i/pa_i.

I of course reject this premise. I also think the name may be misleading. The correlations that remain between X_1 and X_2 given Y's occurrence do not arise 'spontaneously' in the same sense in which time trends do or Sober's correlations between Venetian sea levels and British bread prices. They arise from the occurrence of a cause and the way it operates.

This brings us to one of the nice features of Hausman and Woodward's proof. They make very clear that even for causally sufficient variable sets, CMC could be violated for two different reasons: 'brute' dependencies not following from the causal principles governing the system as with time trends and bread prices and those due to causes producing their effects in tandem. They then offer separate cures for each: MOD* for the

first, no spontaneous correlations for the second. This is a strong point about their proof –

this distinction is clearly drawn and the separate problems are ruled out by separate

premises. As they intended, it makes it easy to see where disagreements lie. I clearly

reject the second of these premises.

What about the first? Here I take issue with Hausman and Woodward's discussion of

my view. They spend a great deal of effort in reconstructing the factory example exactly

as I presented it in my comments on their first proof. They then say, '…to the extent to

which Cartwright is unwilling to commit herself to specific claims about what would

happen under various interventions …, it seems to us she has not clearly specified the

causal structure of the example' (Hausman and Woodward [2003], ms p 14).[15] But it is

clear from the formulation what happens: intervene by manipulating U_1 and X_1 changes

because U_1 causes X_1; X_2 and $P(X_2)$ do not change because U_2 and Y cause X_2 and since

U_1 is an intervention, changes in it are supposed not to change U_2 and Y since they are

not effects of X_1; $P(X_2/Y)$ does not change because $U_1 \perp X_2$, unconditionally and

conditional on Y; and of course $P(X_2/X_1)$ does change.

Hausman and Woodward also say 'Cartwright's case that the chemical factory

example is a genuine counterexample to [CMC] seems most plausible if one accepts

MOD* …' (Hausman and Woodward 2003, ms p. 14), suggesting by this and other

remarks that I do not. To the contrary, I accept MOD* for a vast array of cases[16] and I

built the chemical factory formulation to satisfy it. As they say, we must be assuming

MOD* or something like it every time we draw a causal conclusion from a controlled

experiment.

They also take issue with me for accepting in the case of the chemical factory that 'It should make no difference to the value of [X_1] whether we set [X_2] [by intervention] or observe [X_2] once we set the parents of [X_1] [i.e., once we set Y by intervention]' (Hausman and Woodward [2003], ms p. 19, fn 11) while rejecting their claim called *PM2* as it applies in the chemical factory case; PM2: $P(X_1/\text{set-}Y\&X_2) = P(X_1/\text{set-}Y\&\text{set-}X_2)$. But it is right to accept the first for the chemical factory example and reject the second.

Imagine an occasion on which we set Y so Y must occur. Y occurs. On this occasion Y produces X_1 and thus, since Y produces X_1 iff Y produces X_2, X_2 occurs. If we also on this occasion intervene on X_2 to make X_2 occur, X_2 will still occur – it will be overdetermined – and so will X_1 occur. So whether we intervene on X_2 will make no difference to the value of X_1. Imagine on the other hand that Y does not produce X_2, so X_1 does not occur on this occasion. If we were to produce X_2 by intervening, that won't make Y suddenly produce X_1 so X_1 will still not occur. Again, whether we intervene on X_2 will make no difference to the value of X_1.

But the claim about probabilities doesn't follow from the claim about values and is indeed false. The conditional probabilities of X_1 change although the values never do for the usual reason. Imagine Y is set. Then when the intervention is off, all X_2 occasions will be X_1 occasions. But among the set-X_2 occasions, only 80% will be X_1 occasions; that's true just because no $\neg X_1$ occasion ever turns into an X_1 occasion just by turning the occasion from a $\neg X_2$ one into an X_2 one.

They also say that I cannot endorse the first claim and accept the arrow-breaking interpretation of intervention that they offer in their new proof and that I suppose in my chemical factory case. But that's a mistake too. Perhaps Hausman and Woodward think

that intervening on X_2 will interfere with Y's operations, but obviously that shouldn't be the case for an intervention. Setting $U_2 = 1$ should leave Y's operations unaffected. (Here we see some of the complications in defining 'intervention' – obviously in cases of probabilistic causality we want to ensure that an intervention on one variable doesn't interfere with whether another would or would not produce its result on any occasion.).

In their discussion of product/by-product cases, Hausman and Woodward argue that 'the explication of causal claims in terms of what would happen under various hypothetical interventions does provide …an independent purchase [on the content of causal claims]' (Hausman and Woodward [2003], ms p. 14). I agree that it does – so too do all the other theories of causation on offer and all the other methods (like the mark method) that we use to test for causality. But even if we took theirs as the central purchase, it does not help the case for CMC nor provide support for the no-spontaneous-correlation premise since MOD* can be readily satisfied in cases where causes produce their effects in tandem.[17]

So I draw

> Conclusion 8: Product/by-product cases that violate CMC can be ruled out by a
> specially designed premise but that does not show much. And it is no help in
> establishing a route from testability to CMC.

9 Overall conclusion

The route from manipulability/testability to CMC isn't there. CMC is not a reflection of any important metaphysical facts about causation. And anyway, those putative facts about causation are not facts!

Acknowledgements

Department of Philosophy, Logic and Scientific Method

London School of Economics and Political Science

Houghton Street

London WC2A 2AE, England

Email: philcent@lse.ac.uk

and

Philosophy Department, 0119, UCSD

9500 Gilman Drive

La Jolla, CA 92093-0119, USA

Email: ncartwright@ucsd.edu

References

Cartwright, N. [2001]: 'Modularity: It Can – and Generally Does – Fail', in M.C. Galavotti, P. Suppes and D. Constantini (eds), *Stochastic Dependence and Causality*, Stanford: CSLI Publications.

Cartwright, N. [2002]: 'Against Modularity, the Causal Markov Condition and any Link between the Two' *British Journal for the Philosophy of Science*, **53**, pp.411-53.

Hausman, D. and Woodward, J. [1999]: 'Independence, Invariance, and the Causal Markov Condition', *British Journal for the Philosophy of Science,* **50**, pp. 521-83.

Hausman, D. and Woodward, J. [2003]: 'Modularity and the Causal Markov Condition', forthcoming in *British Journal for the Philosophy of Science.*

Sober, E. [2001] **'**Venetian Sea Levels, British Bread Prices, and the Principle of the Common Cause' *British Journal for the Philosophy of Science,* **52**, pp. 331-46.

Spirtes, P., Glymour, C., and Scheines, R. [1993]: *Causation, Prediction and Search*, New York: Springer-Verlag.

Footnotes

[1] See discussion in Section 6 for why I say 'at best'.

[2] Usually it is said that the graph is over $\mathcal{V}$, but often in practice U's appear on the graphs as exogeneous causes. This is particularly important for Hausman and Woodward since in their proof the interventions on variables in $\mathcal{V}$ – which are supposed to be causes of the quantities represented by those variables – will be members of $\mathcal{U}$. There is also the question of whether various concepts like *intervention* are defined relative to a graph or to 'reality'. Since concepts central to CMC are defined relative to a graph, I think it is best to define all the concepts relative to a graph. (The alternative is at best very messy and certainly impossible without resorting to the concept of a 'correct' graph.)

[3] Cf. my discussion of MOD versus MOD# (Cartwright [2002], p. 417).

[4] Hausman claims that his work is intended to provide a boundary condition for the applicability of causal concepts. In this case my remark here should read, 'It is essential to the correct application of "cause" that …'. Perhaps thereafter we need always to read 'Xc $\rightarrow$ Y' as 'It is correct to say that X causes Y.'

[5] Cf Woodward's principal claims/definitions in Woodward [2003], **TC**, **DC** and **M**. All state as necessary and sufficient condition for X to be some particular kind of cause of Y, that 'there is a possible intervention on X that will change Y.' Again, in [2003], p. 114, Woodward explains that 'these conditions tell us what must be true of the relationship between X and Y if X causes Y…' It is in this very same paragraph, however, that we see a withdrawal from the strong claims recorded in **TC**, **DC** and **M**, that causality requires that there *be* interventions: ' When we engage in causal inference regarding the effects of

X in a situation in which there is no variable that satisfies all the conditions for an intervention variable with respect to X … we should think of ourselves as trying to determine what would happen in an ideal hypothetical experiment in which X is manipulated in such a way that the conditions in [the definition of intervention] *are* satisfied.' [2003, p 114, ital. original]

[6] See caveat in next section.

[7] For a more extended discussion see Cartwright [2001].

[8] Hausman and Woodward do not use the conditional probability, presumably because they do not wish to assume that the interventions themselves have a probability measure. But they need a measure over interventions to discuss MOD* since interventions are supposed to be probabilistically independent of various variables in $\mathcal{V}$, so I assume throughout that there is a measure over $\mathcal{V} + \mathcal{U}$.

[9] They still maintain this thesis in places in M&CMC and still sometimes conflate it with the weaker MOD*. See footnote 13 below.

[10] In many treatments the situation pictured in Graph 1 is ruled out by CM1. But recall that Hausman and Woodward do not assume CM1.

[11] For brevity I here repress the other quantifiers and the other assumptions.

[12] We should make special note of this last as well, for it is a very strong notion of testability – we want to be able to test every single causal hypothesis about the variables in $\mathcal{V}$.

[13] We need not be distracted about the issue of whether or not when an effect follows the occurrence of a purely probabilistic cause we should think that there is an additional event of the cause's 'firing' or 'producing' the effect. If we do not want to admit these

kinds of events, we can take the α's to be mere notational devices that allow us to represent causal claims-cum-probability distributions as equations.

[14] The requirements for an intervention are slightly different in the new paper from any versions in the old. For the definition in the new paper I am not quite sure how they envisage writing equations where some of the U's are interventions. Perhaps $Y_c = \delta_Z (\Sigma a_i X_i + U_y) + y^* Z$ for some chosen value y^* of Y, where $\delta_Z = 1$ when $Z (\varepsilon \, \mathcal{U}) = 0$ (ie, Z = off) and $\delta_Z = 0$ when $Z = 1$. The exact formulation doesn't matter though, since I began my formulation with a perfectly standard deterministic case where the U's satisfied the requirements for an intervention, whatever Hausman and Woodward wanted these requirements to be, then simply changed the operation of the factory from one that produced the chemical and the pollutant deterministically to one that produced them probabilistically, leaving intact from the previous deterministic case any alternative factors that can intervene and create the chemical or the pollutant independently of the action of any other causes.

[15] Sometimes I think Hausman and Woodward conflate the issue of whether there are interventions (as defined in any of the ways they propose) that can set the values (or probabilities) of the chemical and pollutant independently of what other causes for them are doing with the question of whether it is possible to stop Y itself from causing X_1 without stopping it from causing X_2. The formulation I gave is explicit about the first – which is what matters for MOD* and for tests of whether, for example, the chemical causes the pollutant or not (i.e., in their language, for 'disentangling' the common cause explanation of the correlation between chemical and pollutant from a direct cause

account), but my formulation is silent about the second. The answer would presumably

differ from one case to another, depending on the facts of the situation.

[16] Though not all cases. I think brute correlations may well occur in many situations; we

want to be sure they aren't happening whenever we draw causal conclusions from

correlations.

[17] Hausman and Woodward also, in passing, try to defend the view that it should be

possible to manipulate each factor separately – that is, that intervention is always

possible. They do so by attacking my claim that equations that provide information about

a full set of causes need not also provide information about what can and cannot be

manipulated separately. Their argument is just their argument in favour of MOD* -- 'in

the absence of modularity there will be changes in the values of variables under

interventions on other variables that are not reflected in the causal claims expressed in the

system of equations.' (Hausman and Woodward [2003], ms p. 13). This argument is

invalid since the premise supports MOD*, which states what happens **if** intervention

occurs, but the conclusion is that intervention is always possible.