Chapter 15

Causality, Invariance, and Policy

Nancy Cartwright[1]

Introduction

This chapter has five aims:

1. To explain the puzzling methodology of an important econometric study of health and status.

2. To note the widespread use of *invariance* in both economic and philosophical studies of causality to guarantee that causal knowledge can be used, as we have always supposed it can be, to predict the effects of manipulations.

3. To argue that the kind of invariance seen widely in economic methodology succeeds at this job whereas a standard kind of invariance now popular in philosophy cannot.

4. To question the special role of causal knowledge with respect to predictions about the effects of manipulations once the importance of adding on invariance is recognized.

5. To draw the despairing conclusion that both causation and invariance are poor tools for predicting the outcomes of policy and technology and to pose the challenge: what can we offer that works better?

1. A Puzzling Study of Health and Status

It seems being poor is not good for your health. Consider the following remarkable observation

> Travel from the southeast of downtown Washington DC to Montgomery county in Maryland. For each mile travelled life expectancy rises by about a year and a half. There is a twenty-year gap between poor blacks at one end of the journey and rich whites at the other" [1] .

This striking quote is from an eminent epidemiologist, Michael Marmot, who argues that there is a social gradient in health: The higher your status the better your health; and the phenomenon is widespread, observed in the highly unequal United States, in more equitable Scandinavia, and even in the illustrious British civil service.

But how does Marmot know that it is *status* that is the cause of the health differences and not for example the reverse? I want to describe one attempt to answer this—a study by Adams, McFadden, et. al., a group of prominent economists (including a Nobel prize winner) who use econometric techniques to investigate the causal relations between socio-economic status and health among elderly Americans. What matters for our discussion here is the way the study tests for causality, in particular the fact that it uses two different tests.

The first test of the Adams, McFadden, et. al.[2] study looks for

Correlations between health and status, holding fixed other postulated causes of status.

[1] Marmot, 2004 , p.2
[2] Adams, P., Hurd, M., McFadden, D., Merrill, A., Ribeiro, T., 2003.

The second looks for

- Invariance of the estimated correlation across the sample period. Is the correlation that obtains in one period the same as that in another?

My puzzle was, 'Why two tests?' My first thoughts were that the authors are cautious. They offer two tests for the same claim interpreted in the same way: Low status causes poor health. This suggestion has an initial plausibility. The first test is of a kind widely used throughout the social sciences; it is just what we would expect under a *Suppes—*or *Granger*–style (Suppes 1970; Granger 1969, 424–438) theory of probabilistic causality. As for the second, invariance is the central characterizing feature of causality under a number of contemporary accounts in both economics and philosophy. Since any single test is likely to be flawed, the cautious scientist will aim for the convergence of results across different tests for the same thing.

If this were the aim though, the specific strategy employed in the study would be a mistake. On a variety of current invariance theories of causality it is easy to produce scenarios on which Suppes-style causality holds, but the requisite invariance does not, for the very reasons that economists from Mill to Lucas, including the founders of econometrics, have stressed: The underlying arrangements that give rise to economic regularities are often not stable across time and can be highly sensitive to interventions. So the two tests must be testing for different things, perhaps two different kinds of causality.

I now think we should interpret the use of two separate tests differently. The first is a genuine test for causation.[2] The second is a test to see whether the causal relations confirmed to occur in one period by the first test continue by the same test to be confirmed to hold in a second period. The authors themselves claim that the second

serves as a weak test for the usefulness of the estimated correlation for policy

prediction. We should like to know

> Will the correlations that occur in the data set (and thus the causal relations
>
> they indicate) be invariant across the envisaged policy changes?

> It is a small indication in that direction that they are invariant at least across

the period of the data.[3]

Following their lead, I think we should interpret the first as a test for causation

and the second as a step toward showing that the established relation is useful for

predicting outcomes of proposed policy. This interpretation gives rise to the central

question of this paper: What makes causal relations especially useful for predicting

the outcomes of future policy and technology?

2. Invariance in Economics and Philosophy

This two-step process is not peculiar to the Adams, McFadden et. al. study. It is to be

found in many other places in the current literature on causality, notably in the

accounts of James Woodward (2003), Judea Pearl (2000), the Glymour-Spirtes group

(1993), and in David Hendry's (2000) work. In each of these, it plays the same role as

in the Adams, McFadden, et. al. study. Each of these provides an account—a different

account—of what makes a set of relations *causal* relations. Causal knowledge is

valuable because of its importance for policy and planning; we suppose some kind of

tight connection between causal knowledge and the ability to predict the results of

manipulations. In all these accounts, it is some kind of *invariance* assumption that

secures this connection.

On the philosophy side, I shall focus on Pearl (2000), Woodward (2003), and the Glymour-Spirtes group (1993), both because invariance is an explicit demand in their links between causal knowledge and policy prediction and because the link is seldom made in other accounts. It is easy to see this point by looking first at an account of causation where we might have supposed the link to be immediate, David Lewis's counterfactual account (Lewis 1993 [1973]). For Lewis, C causes E just in case C had not occurred, E would not have occurred, where the change from C to $\neg C$ is supposed to occur by miracle; that is, nothing changes except C and whatever is causally consequent on that. Suppose then that we know with certainty that C causes E in this sense. What does that tell us about the effects on E should we manipulate C? Nothing—unless we are in a position to perform miracles. No inferences about strategies can be drawn from the fact that C causes E on Lewis's account without making additional assumptions.

This is exactly what both Woodward and Pearl do. Both add the assumption—the *modularity* assumption—that miracle-like changes are always possible with respect to any factor that can genuinely be counted a cause. Both Woodward and Pearl discuss only systems in which the processes connecting causes with their effects are discrete: there is always one last set of causal factors (the 'direct' causes) that operate just before the effect is produced. This provides them with an analogue of Lewis's assumption that the miracle happens at the last instant, which avoids a host of counterexamples and inconsistencies. They suppose that C causes E only if the law connecting C with the last set of factors that produce it can be replaced with a new law that dictates $\neg C$ while nothing else changes that is not causally consequent from that. So knowing that C causes E tells us at least this about manipulation: it is actually

possible for *C* to change and only *C,* and if that happens the requisite change in *E* will follow.

To see this, let us look at Woodward's work. I concentrate on him because he is probably the most vocal champion of invariance. Both Pearl and Woodward focus on systems of linear equations of a familiar sort, which I call *linear deterministic causal systems with probability measures*. The same form is also compatible with the more general Glymour-Spirtes scheme.

A linear deterministic causal system with probabilities looks like this:

$x_1 \; c= u_1$

$x_n \; c= \Sigma a_{ni} x_i + u_n$

$\text{Prob}(u_1, \ldots, u_n) = \ldots$

where the *u*'s represent quantities not caused by any of the *x*'s, and the symbol *c=* means that the left- and right-hand side are equal and that the factors on the right are the direct causes of those on the left. Different theories of causality place a variety of different constraints on $\text{Prob}(u_1, \ldots, u_n)$. In characterizing a linear causal system with probabilities, I take a minimalist stand and include none of these constraints. The system is defined by its triangular form, which reflects a number of usual assumptions about causality, for example, that causality is irreflexive and asymmetric.

For Woodward, two demands must be fulfilled for equations like these to be properly labeled "causal."

- *Level invariance:* The equation must remain invariant under any changes
  on right-hand-side variables 'by intervention'. EDITOR: it is essential that
  the quotes stay here. They indicate that this is a special usage to be
  cautious of – so-called 'scare quotes'(An intervention on a factor changes
  the law linking that factor with its direct causes to a law that sets that factor

at some specified value, with no other changes than those causally consequent on that.)

- *Modularity:* There must be some way to change the other causal relations in a system that leaves any genuine causal relation invariant.

Within a linear deterministic causal system, if we assume that any functionally true association derives from the basic causal principles of the system, it can be shown that being level invariant is a sufficient condition for a functionally true association to be one of the basic causal principles (Cartwright, 2007). So level invariance can be seen as a representation of the triangular structure of a causal system and the underlying facts about causality that it reflects.

What then of modularity? Both Woodward and Pearl demand not only that a system of causal equations be triangular in form but also that it be modular. Why do they build this additional demand into their characterization of causality? The effect of this requirement is that each variable in a system[4] can be changed (by changing the law that governs it) as if by miracle, without changing anything else except the effects of that variable. What justifies this as a condition on causality? Woodward and Pearl both give the same reason as Adams, McFadden, et. al:[5] this addition allows us to use the relation in question to predict what will happen under manipulation. That is, I take it, why Woodward calls his account of causality indifferently an "invariance" account and a "manipulability" account.

The special kind of miracle-like interventions envisaged by Lewis, Pearl, and Woodward are important for manipulability in Spirtes, Glymour, and Scheines (1993) as well. Spirtes, Glymour, and Scheines (1993) have a "manipulation theorem" that tells how to calculate facts about the new probabilities that occur after one of these

special interventions from facts about the probabilities and causal relations that obtain before (Spirtes, Glymour & Scheines 1993, 75–81). But they are more cautious in their claims than Woodward and Pearl, for they do not assume modularity—that is, they do not assume that such interventions are always possible. Rather, they show how to calculate what would follow were such an intervention to occur.

3. A More Useful Kind of Invariance

Modularity thus secures a sure connection between causality and predictability under manipulation. But how satisfying is this connection? In fact it will allow us to use a given causal relation for very few policy manipulations. That is because the kind of manipulations under which it guarantees invariance—and hence predictability from the laws of the system—are very special. They are just the kinds of "surgical incisions" that we would demand in a controlled experiment, and these are very unlike real policy changes.

First, in policy cases we have little guarantee that causal processes are discrete; and even where they are we seldom are in control of the direct causes of a factor we consider manipulating to produce some desired effects. Second, when we do manipulate some factor we generally find ourselves changing far more than that single factor and its direct consequences. We usually end up changing a number of other factors relevant to the effect and very often we change the very principles by which these factors operate as well.

As I noted in Section 1, this is a problem that economists have been sensitive too from Mill (1994 [1884]) through the founders of econometrics (especially Haavelmo (1997 [1994]) and Frisch (1997 [1938])) to Lucas (1981 [1976]). Perhaps that is why it is in economics that we find accounts that connect causation with more

realistic kinds of manipulations. Consider David Hendry (2000), who sometimes suggests that causes are super*X,* where super*X* is a combination of *X* and invariance. Hendry's most developed example involves weak exogeneity and invariance.

Weak Exogeneity:<

Given $P(Y\&X,\beta U\gamma) = P(Y/x,\beta)P(x,\gamma)$, x is weakly exogenous to a vector of outcomes Y if the parameters γ of the marginal distribution have no cross-restraints with the parameters β of the conditional distribution.

Weak exogeneity is a nice characteristic because it ensures that the marginal distribution can be ignored in estimating the conditional distribution. But it is not essential. If we envisage changing x to control the probability of Y, it is the conditional distribution itself that matters for our predictions independent of how we can learn about it. I follow Hendry in illustrating with a case where the variable (x) we envisage manipulating is weakly exogenous to the vector of outcome variables (Y) we care about because it makes very clear the importance of the additional invariance assumption.

Suppose we think of changing the distribution of x in order to affect the distribution of Y. It may seem that we can predict the outcome from the formula for the conditional distribution. But that is not so. Changing β changes the joint probability distribution, and there is nothing that ensures that the new distribution will still be the same. In the original distribution, β and γ may have no dependencies but that does not show what happens if the distribution is changed. So Hendry adds a constraint and defines: x causes Y if and only if the parameters of $P(Y/x)$ stay fixed as

we vary the parameters of the distribution of *x*. In the case where *x* is weakly

exogenous to *Y,* this gives us the notion of *super-(weak)exogeneity.*

> *Super-(weak)exogeneity*
>
> Suppose *x* is weakly exogenous for *Y*. Then *x* is *super-(weak)exogenous*
>
> relative to a proposed intervention (say a change in γ) if the parameters of
>
> interest (say β) for $P(Y/x,\beta)$ do not vary under the intervention.

I am going to discuss Kevin Hoover's account of causality. I note here that it

shares with Hendry's an important advantage vis-à-vis strategy over the

Pearl/Woodward/Spirtes, Glymour, and Scheines approach. For Hendry, causation is

characterized relative to a proposed intervention as that intervention will actually

occur. His definition of causality requires that the relation we should like to use to

predict the outcomes of our proposed manipulations be invariant under exactly those

manipulations. The others insist, instead, on the invariance of the relation under

highly artificial manipulations, manipulations that might occur in a controlled

experiment to test for a causal connection but would hardly ever be ones we envisage

for a real application, either in setting policy or in building a device or an institution.

Unlike the earlier philosophical accounts, the accounts of causality that Hendry and

Hoover offer show why causal knowledge is good for policy prediction, as we think it

is, whereas these others do not show this.

Despite its nice connection with policy prediction, there is a difficulty with

Hendry's account, however; it doesn't seem to be an account of causality at all. That's

because of Hendry's focus on the conditional probability. A factor *x* we consider

manipulating to affect *Y* will do so given Hendry's invariance assumption just in case

Y is probabilistically dependent on *x*. But it is one of the truisms of causal theory that probabilistic dependence (correlation) is not causation.

It is easy to see these points by looking at a case with dichotomous variables. By the laws of probability

$$P(E) = P(E/C)P(C) + P(E/\neg C)P(\neg C).$$

In order to increase the probability of *E*, we consider a manipulation that takes the probability *P* to a new *P*,' where *P*' has an increased probability of *C*. Under Hendry's invariance assumption, $P(E/\pm C)$ is to stay fixed. So

$$P'(E) = P(E/C)P'(C) + P(E/\neg C)P'(\neg C).$$

So the strategy works just in case $P(E/C) > P(E/\neg C)$. From this I conclude that the *X* in Hendry's account (Causality (relative to *G*) = super*X* (relative to *G*) (i.e., *X* + invariance of *X* under *G*)) EDITOR yiu've taken out the quote marks here….but the phrase needs something more than commas to indicate that I am here reiterating his account. And anyway, you didn't even have a closing comma. On reflection I think the parenteses are the best solution can easily be probabilistic relevance. But that's odd because we all believe that probabilistic relevance is not causation and adding on that the relevance relations are invariant under some envisaged manipulation does not seem to make it so.

There is a similar difficulty with Hoover. Hoover distinguishes between parameters (represented by Greek letters, *α, β, γ,* . . .), which we control 'directly' EDITOR again 'scare quotes are required here and again I don't know if you want to

use single or double quotes. But one or the other is needed, and variables (represented by Latin letters $x, y, \ldots$), which we do not. He takes the notion of direct control by us to be primitive in his account and uses it to define causal relations between quantities we cannot directly control. The account is restricted to quantities whose values can be fixed, albeit indirectly, by manipulations we can perform. Let Par_z stand for the set of parameters that determine z. Hoover's definitions require that

Hoover causation: x causes y iff Par_x is a proper subset of Par_y.

So x (Hoover) causes y iff anything we can do to fix the value of x partially fixes the value of y but not the reverse. So Hoover's characterization ensures that knowing causal relations allow us to predict the results of manipulations we might perform. But as with Hendry this characterization will sometimes count as causes factors that would usually be counted as mere correlates. Consider for instance this simple arrangement:
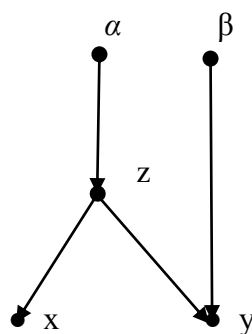
<Figure 15.1 about here>

Figure 1

Here the arrows are meant to represent relations that count as causal by whatever is your favourite other characterization of causality; we can for instance imagine building a machine with mechanical connections that fits the model in Figure 1.

We are thus faced with a trade-off. Hendry and Hoover connect causal knowledge with the predictability of the results of real manipulations, but they do not seem to be real causes; whereas Woodward, Pearl and Spirtes, Glymour and Scheines seem to connect real causes with manipulations, but they are not real manipulations.

4. Causality: What Is the Use of It?

Setting Hoover aside for the moment, there is another problem raised by the discussion of the approaches in the last section. Whether we start with real causation or some other relation and whether we end up with the ability to predict what happens under realistic or under miracle-like manipulations, what makes for the connection between the two is invariance. Woodward defines a causal relation as one that is invariant under miracle-like manipulations on right-hand-side variables. So clearly a Woodward-causal relation will predict accurately what happens under those kinds of exotic manipulations. The same is true for Pearl.[6] Hendry defines a causal relation as one that is invariant under various more realistic manipulations. So clearly a Hendry-causal relation will predict accurately what happens under those more realistic manipulations.

The logic is simple. We have an association. We assume it to be invariant under a particular kind of manipulation. So we are able to use that association to predict what happens under the specified kind of manipulation.[7] This logic works no

matter whether the starting association is causal or not. Hendry's proposal is a case in point.

What good is causation then? It is generally supposed that there is some special connection between causation and policy prediction. Knowing the causal relation between two variables is supposed somehow to put us in a better position to predict what happens when we manipulate the first than simply knowing some arbitrary 'spurious' relation between them. But that does not seem to be the case.

Perhaps, despite my qualms in section 3), modularity is the answer after all. Both Woodward and Pearl insist that having the triangular form of a causal system is not enough to make a set of associations causal (with which I agree). There must in addition, they maintain, be some miracle-like manipulation possible on every variable in the system. Perhaps it is this very fact that makes causal knowledge so much more useful in general than knowledge of 'mere correlation.'

First, I do not think the claim is true. On the modularity thesis a relation is not causal unless there is some way to manipulate the cause, no matter how many other earmarks of causation the relation has. Nor is it enough that we be able to manipulate the cause, which may be hard enough; it must be possible to manipulate it in a very specific way—surgically, as if by miracle. I do not see any reason for believing this,[8] other than to satisfy the demand that causes should connect with strategies.[9]

My second worry is that the proposal does not do the job it is supposed to: It does not show what is special about causal relations over spurious ones. The modularity solution maintains that if a relation is causal, then there is always some manipulation of the cause that leaves that relation invariant, albeit an exotic manipulation. Suppose that is true. Then it is equally true for the spurious relation between joint effects of a common cause: There is always some manipulation on the

first that leaves the relation between them invariant. Simply use the miracle-like manipulation hypothesized to be always available on the common cause to manipulate the first of the two joint effects. (For instance, in Figure 1 jiggle α to manipulate x by manipulating z.) This will change the second of the joint effects as well and leave the spurious relation between them invariant. Clearly, this manipulation will not itself be miracle-like on the first of the two related factors. But if the hope was to argue that causal relations are special because there is always at least some manipulation that leaves them invariant, miracle-like manipulations seem to have no special place.

This is indeed my third worry. Miracle-like manipulations of the kind under consideration are great for finding out about causal relations since they are the kinds of manipulations we would wish to make in a controlled experiment. But, as I urged in Section 3, they are not the manipulations we envisage in policy and technology. Yet we do think knowing causal relations is especially useful for planning. The possibility of invariance under miracle-like manipulations does not account for this.

An alternative to the modularity thesis that could explain the practical usefulness of causal knowledge would be to argue that causal relations are more likely to be stable than are mere correlations. This might be supported by my own observations that we often build both devices and institutions with shields to protect the structural arrangements from disruptions (Cartwright, 2007; 1989). But I cannot see how to use this to support the distinction I am looking for. Once a shield has been put in place to protect the internal arrangements of a structure, then causal relations and mere correlations may be equally stable. For instance, imagine that a shield is built around x,y,z, and β and their causal connections in Figure 1, allowing only the influence of α to penetrate. Then the spurious relation between x and y will be just as stable as that between z or β and y.

Conversely, one of the special worries in policy that we have noted is that causal relations are often not all that stable under manipulation. Not only is this the core of the famous Lucas critique; it was central to the philosophy of John Stuart Mill. It was, for instance, an important support for Mill's opposition to the subjugation of women. Mill admitted that, under the contemporary structure, putting women into positions of authority might well not produce good outcomes. But that, he maintained, would most probably change if the institutions of society changed to provide women the education and opportunity that would allow them to develop and exercise their native capacity for independent and creative thought (Guillin 2006).

In the end, the claim that causal relations are in themselves more stable than spurious ones seems too vague and too weak to serve as a defense of the vast effort we put into trying to secure causal knowledge. What seems true is that knowing causal relations is hugely useful for planning and prediction whenever we can add on the assumption that they will be stable. Nor is there anything wrong with an account of causation that needs to add this on. As Kevin Kelly[10] has pointed out, it is equally true of theories of mechanics (indeed, any theory for that matter) that they need what (following Wilfrid Sellars) I call *theory-exit assumptions* if they are to be put to use.

The problem is that this does not seem to distinguish between causal and noncausal relations. A simple kind of Humean view does better: If a causal relation is a universal association—it always holds whenever the cause occurs—then, clearly, causal relations are sure predictors. But this is not the case with any of our contemporary theories of causality. It seems that causal relations will provide secure predictions about what happens under manipulations just in case they are invariant under those manipulations. But so, too, will noncausal relations. Why then do we take causal knowledge to be so much more useful than knowledge of other relations?

My worries here are not that we can find no difference between stable causal relations and other stable relations vis-à-vis manipulation. Consider a situation in which neither x nor y occurs but in which the principle "x causes y" holds. In this case, if we can change from not ¬x to x, leaving the principle that x causes y unchanged, then we can ensure not only that y changes but also that it is x that changes it. This contrasts with the relation between x and y in Figure 2. There if we change from ¬x to x, y changes as well. But the change from ¬x to x will not cause the change in y. We do not change y *through* changing x. This difference may be important ontologically. But it goes no way toward explaining why causal knowledge should be especially useful for prediction and control, worth buying at great cost.

5. Are Both Invariance and Causality Red Herrings?

Sandra Mitchell (2003) points out that any true claim can be useful. Suppose we follow her lead. Perhaps causation and invariance are not the best keys to good prediction; they certainly are not the most direct. The simplest claim that will allow us to predict what will happen under manipulation is a true claim that describes just that.

This is in essence what Hoover-causation consists of. Causal claims, in Hoover's use of the term *causal,* describe what will happen under the manipulations we can perform. His account thus has the advantage over the other accounts discussed here. With Hendry it focuses on the kinds of manipulations we might actually carry out and does not restrict itself to miracle-like interventions. But it is more general. Hendry-causation secures prediction when a relation that predicts the outcome under current arrangements continues to occur under the proposed manipulation. Hoover-causation looks instead directly for information about what will happen given the

manipulation. It does not depend on associations from the past continuing to hold across interventions.

Of course there is a sense in which this advantage is illusory. For nothing about Hoover-causation suggests how we are to come up with a Hoover-causal claim. But it can point us to an important lesson. Mill taught that economics cannot be an inductive science. Economic arrangements shift regularly in ways we generally cannot predict, and recent economics has made a point of how much more likely this is when interventions occur. Accounts that rely on invariance run just counter to Mill's cautions. Induction is what they offer, though with an explicit admission of Mill's worries, namely, use the associations of the past for future predictions, but use them only when they will continue to hold. We need something better.

Hendry himself is attentive to the fact that existing economic relations cannot be relied on to hold under manipulation. When it comes to forecasting, the use of causal models—even very accurate ones—can be dangerous, he warns, and for the very reasons that worried Mill: The arrangements correctly described in a causal model at one time are not likely to stay fixed across time. In his recent forecasting work Hendry develops a number of alternative modeling strategies that can be shown to give more accurate predictions across time if certain specified kinds of changes are occurring (Hendry & Mizon 2005; Andrews & Stock 2005).

What is surprising is that Hendry urges that these models may be good for forecasting but not for planning. Presumably that is because he imagines that the kinds of changes generally envisaged in planning are not the kinds that his strategies for forecasting deal with. What, then, do we do for planning? What kind of evidence will support policy and technology predictions and how is it to be marshalled and evaluated? That is the challenge, and it is an especially pressing one now that the call

everywhere is for evidence-based policy. As methodologists we need to offer good counsel about just what counts as evidence when predictions about the effects of interventions are at stake, and about how to use that evidence. I do not think we have enough to say.

My conclusions in this chapter about the usefulness of causal knowledge are unfortunately negative. First, surprisingly, causation (at least under conceptions of it of the kind discussed here) seems irrelevant for reliable prediction in policy and technology planning. Causation without invariance will not do the job, and any invariant relation will provide reliable predictions regardless of whether it is causal. Second, invariance may be a good tool but, as Mill taught, it is altogether too rare and too unpredictable to do much for us. If we need to rely on invariance, we will not get very far, and the focus on it may distract attention from the fundamental challenge: to develop and understand methods—generally applicable methods—for evaluating policy and technology predictions.

In Sum

Overall my discussion raises a disturbing question. Causal knowledge is hard won. We spend a great deal of effort to achieve it. But what is the use of it once we have it? Invariance fares little better since it can generally not be relied on in economic policy considerations. What can we offer that is better?

Notes

2. Because of my pluralist views about causality, I would more accurately say "a genuine test for one significant kind of causal relation."

3. Probably we shall really want to be exporting the conclusion to another population as well, not just to the same population under different policies, and that is clearly an even stronger move.

4. That is, any variable that appears as an effect in a law in the system of laws.

5. Actually, he gives the same reason—causes must be usable to manipulate their effects—for both level invariance and for modularity. I cite it only for modularity because level invariance does not provide manipulability unless modularity is added, and I, at any rate, have an alternative defense of level invariance.

6. Recall, Pearl demands of any equation in the causal system that it be invariant when the laws that determine the direct causes in that equation are replaced by laws that set the values of those direct causes at any arbitrary value.

7. A similar claim is true for Glymour, Spirtes, and Scheines, though their scheme is more complicated and provides more complicated inferences. (This is true, too, for Pearl when it comes to his full counterfactual account.) They begin with a mixed set of causal and probabilistic claims and then tell how to calculate what happens to various probabilities under certain miracle-like manipulations—but only supposing that the relations that support the calculation are invariant.

8. I argue for this more extensively in Cartwright (2001)

9. This indeed is how Pearl does defend it.

10. Kelly raised this point at my Center for Philosophy of Science lecture at the University of Pittsburgh, "Where Is the Theory in Our Theories of Causality," March 2006.

References

Adams, P., Hurd, M., McFadden, D., Merrill, A., Ribeiro, T.,(2003) 'Healthy, wealthy, and wise? Tests for direct causal paths.', *Journal of Econometrics*, Volume 112, Issue 1: 3-56

Andrews, D.W.K. and Stock, J.H. (2005) "Inference with Weak Instruments," Cowles Foundation Discussion Paper No. 1530, 2005. Also in *Advances in Economics and Econometrics, Theory and Applications: Ninth World Congress of the Econometric Society*, Vol. III, (2006) R. Blundell, W.K. Newey and T. Persson (eds). Cambridge, UK: Cambridge University Press,

Cartwright, N. (2007) . *Hunting Causes and Using Them.* Cambridge, England: Cambridge University Press.

Cartwright, Nancy D. (2006) 'Where is the Theory in our 'Theories' of Causality?' *Journal of Philosophy*, Vol. CIII, No. 2: 55-66. Also in Nancy Cartwright's *Hunting Causes and Using Them* (2007).

.Cartwright, N. (2001) "Modularity: It Can—and Generally Does—Fail," in *Stochastic Causality (CSLI Lecture Notes)*, M.C. Galavotti, P. Suppes, & D. Costantini (eds), University of Chicago Press.

Cartwright, N. (1989). *Nature's Capacities and their Measurement.* Oxford: Clarendon Press.

Frisch, R. (1997 [1938]). "Autonomy of Economic Relations." In D.F. Hendry & M.S. Morgan, Eds., *The Foundations of Econometric Analysis.* Cambridge, England: Cambridge University Press.

Granger, C.W.J. (1969). "Testing for Causality and Feedback." *Econometrica* 37: 424–438.

Guillin, V. (2006). "Auguste Comte and John Stuart Mill on Sexual Equality: Historical, Methodological and Philosophical Issues." Ph.D. thesis, LSE

Haavelmo, T. (1997 [1944]). "The Probability Approach in Econometrics." In D.F. Hendry & M.S. Morgan, Eds., *The Foundations of Econometric Analysis.* Cambridge, England: Cambridge University Press.

Hendry, D. F., & Mizon, G. E. (2005). "Forecasting in the Presence of Structural Breaks and Policy Regime Shifts." In D. Andrews & J. Stock, Eds., *Identification and Inference in Econometric Models: Essays in Honor of Thomas Rothenberg.* Cambridge, England: Cambridge University Press.

Hendry, D. F. (2000). *Econometrics: Alchemy or Science?* Oxford: Oxford University Press.

Lewis, D. (1993 [1973]). "Causation." In E. Sosa & M. Tooley, Eds., *Causation.* Oxford: Oxford University Press.

Lucas, R. E. (1981 [1976]). "Econometric Policy Evaluation: A Critique." In R.E. Lucas, Ed., *Studies in Business Cycle Theory.* Oxford: Basil Blackwell.

Marmot, M. (2004). *Status Syndrome—How your social standing directly affects your health and life expectancy.* London: Bloomsbury.

Mill, J. S. (1994 [1844]). "On the Definition and Method of Political Economy." In D. Hausman, Ed., *The Philosophy of Economics.* Cambridge, England: Cambridge University Press.

Mitchell, S. D. (2003). *Biological Complexity and Integrative Pluralism.* Cambridge, England: Cambridge University Press.[3]

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

Spirtes, P., Glymour, C. & Scheines, R. (1993). *Causation, Prediction, and Search.* New York: Springer.

Suppes, P. (1970). *A Probabilistic Theory of Causality.* Amsterdam: North-Holland.

Woodward, J. (2003). *Making Things Happen.* New York: Oxford University Press.