Causal Instrumental Variables and Interventions

Julian Reiss†‡

The aim of this paper is to introduce the instrumental variables technique to the discussion about causal inference in econometrics. I show that it may lead to causally incorrect conclusions unless some fairly strong causal background assumptions are made, assumptions which are usually left implicit by econometricians. These assumptions are very similar to, albeit not identical with, James Woodward's definition of an 'intervention'. I discuss similarities and differences of the two points of view and argue that—understood as a practical method of causal inference—the set presented here is superior.

1. Introduction. The interest in the topic of causality has been revived in recent years, not only in philosophy but also in empirical disciplines such as epidemiology, empirical sociology and econometrics. Especially econometrics has seen a number of developments from within its own field (compare, for instance, the contributions made by Hendry, LeRoy, Heckman, Hoover) and applications of methods developed in other fields (e.g., the Bayes'-nets methods). One purpose of this paper is to introduce yet another alternative to the discussion: the instrumental variables technique. A curious fact about that technique is that although its use is fairly widespread as a method of causal inference in applied econometrics, its official methodology (as introduced in econometrics textbooks) is insufficient to guarantee the validity of causal claims that are established on the basis of instrumental variables. On the other hand, in practical applications, econometricians often (implicitly or explicitly) make at least

†To contact the author, please write to: Departamento de Lógica y Filosofía de la Ciencia, Universidad Complutense de Madrid, 28040 Madrid and Centre for Philosophy of Natural and Social Science, London School of Economics, Houghton St, London WC2A 2AE; e-mail: jreiss@filos.ucm.es.

‡Work on this paper was conducted under the CPNSS research project Causality: Metaphysics and Methods. I am very grateful to the AHRB for funding. Many thanks to Nancy Cartwright for valuable comments on an earlier draft.

Philosophy of Science, 72 (December 2005) pp. 964–976. 0031-8248/2005/7205-0026$10.00
Copyright 2005 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/508953 Published online by Cambridge University Press

some of the assumptions that are needed to ensure the validity of the claims.

The negative claim I want to make in this paper is that the instrumental variables technique can lead to false causal conclusions—if the official methodology is applied blindly. This is easily demonstrated by means of a counterexample which satisfies the assumptions required according to that methodology. The positive claim is that under a suitable set of assumptions the instrumental variables technique is a defensible method of causal inference from observational data. The set of assumptions that I use contains a number of very general principles as well as specific conjectures about the system under investigation. The specific conjectures very similar to (but not identical with) James Woodward's definition of an intervention. Thus, the positive claim of the paper can be interpreted as showing that if an instrument satisfies a set of assumptions similar to the definition of a Woodward-intervention, causal inference on its basis is valid. And this is true even where the intervention is not an intervention in the often used sense of the word of "intentional interference by a human agent". The causal structure of the system has to be of a particular kind, but it does not matter whether it is human agents in general, or the experimenter in particular, who intervene.

That's the good news. The bad news is that the assumptions made are very strong indeed. Except a few suggestions that there is good reason to believe that they are at least sometimes satisfied I do not attempt to defend any of them here. Nor do I defend the choice of these particular assumptions except by showing that they do the intended job. Specifically, I do not claim that my preferred set is the only set sufficient for the rendering the instrumental variable technique valid. But I do want to make clear that some set or other is required; and that among the assumptions there will be at least some that are causal in nature and at least one that carries us from probabilities to causes.

2. Can We Get Causes from Statistics? "Correlation is not causation" is a well-rehearsed slogan. The purpose of this section is to argue that prior knowledge about causal relations is not only necessary to identify causal parameters from statistics but that it is also often easier to have than knowledge about correlations.

The reason to re-rehearse the slogan is that econometrics textbooks tend to avoid explicitly causal language in their descriptions of estimation procedures. For example, in simultaneous equations, variables divide into endogenous and exogenous variables, and that means that they are either determined within the model or outside it. Importantly, without qualification, the "determined" can be read both functionally or causally. Or sometimes a model is called "structural" if its form is given by the un-

https://doi.org/10.1086/508953 Published online by Cambridge University Press

derlying *theory*. But no word is lost on whether the theory specifies functional or causal relations. To give a third example, the error term in a regression model is often said to represent "stochastic disturbances" or "shocks" as well as omitted factors, measurement error and other "influences". Again, it is unclear whether these terms carry a causal meaning or not.

The neglect of causality in econometrics has its roots deep in the origins of modern statistics itself. Thus Karl Pearson, famous for being one of the founders of modern statistics, is at least equally famous for his attacks against causal language and his substitution of the Humean concepts of association and correlation for causal concepts (see, for instance, Pearson 1911). For him, causal language was the language of the past; the language of the new era is that of association and correlation.

This sentiment carried over to early econometrics. For example, in one of his volumes on business cycles, Wesley Mitchell muses:

> In the progress of knowledge, causal explanations are commonly an early stage in the advance toward analytic description. The more complete the theory of any subject becomes in content, the more mathematical in form, the less it invokes causation. (Mitchell 1927, 55, quoted from Hammond 1996, 10)

Surveying more recent contributions (a growing number of exceptions notwithstanding), one still gets the impression that, in order to be scientific, causal language has to be avoided, while the Humean language of correlation and association counts as scientific.[1] In particular, many textbooks contain 'recipes' for econometric inference that give the impression that econometrics can proceed without causal background assumptions.

But strangely, a large part of theoretical economics asks questions that are causal in nature. Many historical controversies, for example, can be understood as controversies about the causal role certain aggregates play in the economy. Famous examples include the causal role of money in determining other aggregates such as income or prices or the interest rate or the causal role of aggregate demand. Methods tailored to measure the association between variables only seem ill-suited to bear any light on such controversies.

Suppose, then, we use the standard techniques of econometrics to answer particularly *causal* questions. Say we are interested in the question

1. See also Pearl 1997, who makes a similar remark about statistics in general.

https://doi.org/10.1086/508953 Published online by Cambridge University Press

of whether money (X) causes income (Y). We can expect the estimator $\hat{\beta}$ of the causal effect of X on Y of the standard regression model[2]

$$Y = \alpha + \beta X + \varepsilon \tag{1}$$

to be biased for a number of different reasons. First, it is possible that X is measured with error, and that the measurement error is correlated with Y's error term ε. That X is measured with error is more than likely; it has been a matter of dispute what the 'right' definition of money should be for about as long as economists thought about the quantity theory. The construction of the money variable has also been one of the major criticisms of Milton Friedman's epic studies of the role of money in the economy. Second, the relationship can be confounded by an unobserved variable. (1) is simplified in so far as the vector of known confounders (e.g., a time trend) has been omitted as a regressor on the right hand side. However, not all potential confounders can be measured, and thus there might always be residual correlation between X and ε. Third, there may be feedback effects from income to money. This is also a point that came up in the debate about Friedman's analysis of the role of money. Critics tended to understand the monetarist position to assert that money is the only cause of economic activity, but Friedman emphasized time and again that he thinks of money merely as one cause (the principal cause, however), and that it is well possible, especially in the short run, that causality runs from income to money.

For any of these three reasons, X might be correlated with ε, which biases the estimator for β. A standard technique to solve that problem is so-called instrumental variables estimation. According to many econometrics textbooks (see e.g., Greene 2000), a variable Z is an instrument with respect to (X, Y) if and only if the following conditions are satisfied:

IV-1. Z is correlated with X.

IV-2. Z is uncorrelated with ε, the error term for Y.

If such a variable can be found, an estimator of the parameter of interest can be constructed as follows:

$$\hat{\beta} = \text{corr}(Z, Y)/\text{corr}(Z, X). \tag{2}$$

If these are the only assumptions made about the system of interest, the definition of an instrumental variable is too weak for a causal interpretation of the estimand. To see that most clearly, let us proceed in three

2. The hat signifies that this is the *estimator* of the true parameter rather than the true parameter itself.

https://doi.org/10.1086/508953 Published online by Cambridge University Press

stages with increasingly stronger assumptions made about the relations between the variables.

There are two steps involved in the inference from correlations to causes. The first is the inference from sample correlations to the probabilities of the underlying populations. The second is from probabilities to causes. In all that follows, I take the first step to be unproblematic. That is, I take all measured correlations to be representing the true population probabilities. Although the two steps may be interrelated in practice and it is disputable whether the step from correlations to probabilities must be taken prior to inferring from probabilities to causes, analytically, the two steps can be separated, and discussion only the second step makes for a much clearer argument.

Now, without *any* further assumptions, the above pattern of correlations is compatible with any causal structure imaginable. The least principle one needs for causal inference from correlations is an inductive principle which connects probabilities with causes. This is implicit in all econometric work but it is hardly ever addressed explicitly.[3] The reason is simple: If we allow for 'brute correlations', i.e., correlations that cannot be accounted for in terms of causal relations, then any attempt to infer causal relations from statistics—no matter how strong the additional assumptions about the system of interest—is futile. Any residual correlation could then always be 'brute'. An almost identical remark make Daniel Hausman and James Woodward:

> So when X_j [the putative effect variable] "wiggles" in response to an intervention with respect to X_i [the putative cause variable] (or in general without any change in its parents) one has a covariation between X_i and X_j that cannot be explained by X_j causing X_i or by a common cause of X_i and X_j [this follows from their definition of an intervention]. If there are no restrictions on the ways that a correlation between X_i and X_j might arise, one could go no further. But if one takes **CM1** [their version of the principle that connects causality and correlation] for granted, then one can conclude that X_i causes X_j (Hausman and Woodward 2004, 151).

The particular version of the Reichenbach principle I am going to assume is the following:

> **RP**. Any two variables A and B are correlated if and only if either (a) A causes B, (b) B causes A, (c) a common cause C causes both A and B, or (d) any combination of (a)–(c).

3. Obvious exceptions include Pearl's 2000 discussion of econometrics.

https://doi.org/10.1086/508953 Published online by Cambridge University Press

Let us call two variables that are related by any of the causal structures (a)–(d) *causally connected* (and *causally unconnected* in case they fall under none of the structures). I do not want to argue for the truth or plausibility of the particular version of the principle here. But let me stress again that we need some principle of the kind, and since investigations in economics typically face causally complex systems, we cannot, for example, exclude a priori the possibility of simultaneous causation or any other combination of the simpler structures (a)–(c).

The second general assumption we must make is the transitivity of causal relations. Singular causal relations are not always transitive (cf. the examples in McDermott 1995) but at the generic level, which I am talking about here, transitivity seems a safe assumption. Thus:

> **T**. For any three variables A, B, and C, if A causes B and B causes C, then A causes C.

I will restrict the analysis to linear equations of the general form $X_j = \alpha_{j1}X_{j1} + \alpha_{j2}X_{j2} + \cdots + \alpha_{jn}X_{jn} + \varepsilon_j$. For these, I define the concept of "functional correctness of a structural equation" as follows:

> **FC**. A structural equation $X_j = f(X_{j1}, X_{j2}, \ldots, X_{jn}, \varepsilon_j)$ is functionally correct if and only if it represents the true functional (but not necessarily causal) relations among its variables.

On the basis of these, let us examine what further assumptions we must make about the system of interest in order for the instrumental variables technique to yield causally correct conclusions. I will proceed in three stages, each with different causal assumptions added to the general assumptions.

Stage 1: No assumptions added.

If no causal assumptions are added to **RP**, **T**, and **FC**, one can easily show that the instrumental variables technique can yield causally incorrect conclusions. The system consists of two functionally correct equations:

$$Y = \alpha X + \varepsilon, \tag{3}$$

$$X = \beta Z + \varepsilon. \tag{4}$$

According to the instrumental variables technique under these assumptions, the claim is that if Y is correlated with Z then X causes Y. However, it is possible that the correlation between X and Y arises from other causal

https://doi.org/10.1086/508953 Published online by Cambridge University Press

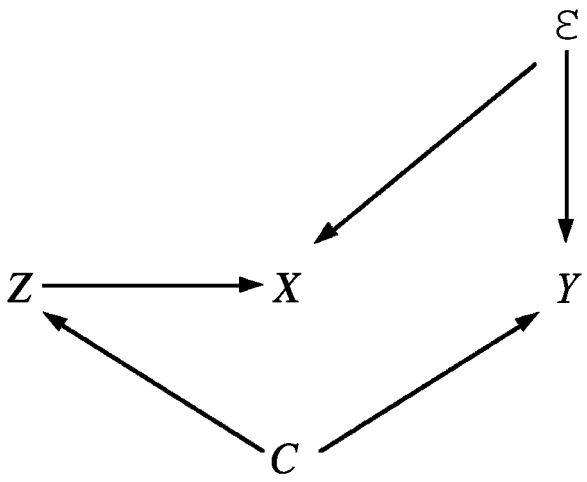


Figure 1. A counterexample to **IV** under **RP**, **T**, and **FC**.

relations. Consider the structure in Figure 1.[4] Nothing in the assumptions prevents this situation. In addition to (3) and (4), we now have two further functionally correct equations:

$$Z = \gamma C + \mu \tag{5}$$

and

$$Y = \delta C + \nu. \tag{6}$$

If $\beta > 0$ (which is implied by **RP**), then Z and X are correlated and thus **IV-1** is satisfied. Suppose further that $\gamma > 0$ and $\delta > 0$ (which, too, is implied by **RP**), which implies that Z and Y are also correlated. This means that the test condition is fulfilled. Now, depending on the various

4. I could show the same result with an even simpler structure in which the instrument is a common cause of X and Y. But since in this case, it is so obvious that the correlation between Z and Y is not due to a causal influence of X to Y, I prefer to illustrate the claim with this slightly more complicated structure.

https://doi.org/10.1086/508953 Published online by Cambridge University Press

parameters, there are cases in which **IV-2** is satisfied: Z and ε may or may not be correlated as can be seen from the following derivation:

$$\begin{aligned} Z &= (\gamma/\delta)Y - \gamma\nu/\delta + \mu \\ &\qquad\qquad [\text{substitute } C \text{ from (6) into (5)}] \\ &= (\gamma\alpha\beta/\delta)Z + (\gamma/\delta)(1+\alpha)\varepsilon - \gamma\nu/\delta + \mu \\ &\qquad\qquad [\text{substitute for } Y \text{ from (3) and (4)}] \\ &= \zeta(1+\alpha)\varepsilon - \zeta\nu + (\zeta\delta/\gamma)\mu \\ &\qquad\qquad [\text{define } \zeta = 1/(\delta/\gamma - \alpha\beta) \text{ and rearrange}]. \end{aligned}$$

Therefore,

$$E(Z\varepsilon) = \zeta(1+\alpha)E(\varepsilon^2) - \zeta E(\varepsilon\nu) + (\zeta\delta/\gamma)E(\varepsilon\mu) \tag{7}$$

(multiply by ε and take expectations) which may or may not be zero. This means that there are parameterizations under which **IV-1**, **IV-2**, and the other assumptions are satisfied, but X does not cause Y.

It seems unlikely that an exact cancellation such as this (which makes (7) equal to zero) should occur, but it is not impossible. One way to ensure that exact cancellations do not obtain that has been offered in the literature is to assume that correlations are stable under parameter changes (Pearl 2000, 48). In this example, we see that the lack of correlation between Z and ε obtains only under a very specific parameterization.

One assumption that would make the procedure stronger is that the error terms in an equation represent the net effect of *all other causes*. The above counterexample could not obtain because there could not be a cause of Y, C, which is not represented by ε. That the technique yields causally correct conclusions in general I show in stage 2.

> Stage 2: ε represents the net effect of all other causes on Y (except X and any cause that influences Y through X or lies on the causal route from X to Y).

If Z and Y are correlated, **RP** tells us that they must be causally connected. Since there is a third variable, X, I will first show that Z and Y cannot be causally connected on a route that excludes X.

Can Z and Y be causally connected on a route that excludes X? Under the given assumptions, Z cannot cause Y directly or via a route that excludes X because its influence would be represented by ε, which we have excluded. Y cannot cause Z unmediated by X either, since by transitivity ε would cause Z and by **RP** they would be correlated. Nor can there be a common cause (that causes Z and Y on a route that excludes

https://doi.org/10.1086/508953 Published online by Cambridge University Press

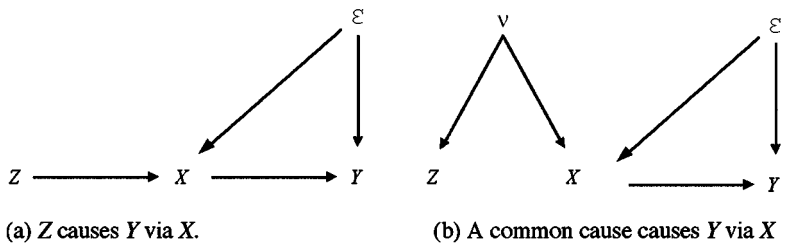


Figure 2. Two valid instruments to identify ε. (a) Z causes Y via X. (b) A common cause causes Y via X.

X) because its influence would be represented by ε. Hence, the causal connection between Z and Y must run through X.

Hence there remain three possibilities: Y causes Z through X, X is a common cause of Y (or a common cause causes Z and Y through X) and Z and Z causes Y through X. Y cannot cause Z through X because, as before, that would imply that ε is correlated with Z. A common cause now may either be between X and Y or between X and Z. If it is between X and Y (which means that its influence is represented by ε), then X must cause Z and thus ε causes Z by transitivity and would be correlated with it. However, there may be a common cause between Z and X. In this case Z and Y would be correlated if only if X causes Y, and thus Z would be a valid instrument.

The last remaining possibility is that Z causes Y via X. Again, in this case Z and Y would be correlated if and only if X causes Y, and Z is a valid instrument. The two causal structures in which Z is a valid instrument are depicted in Figure 2.

Thus, under the above assumptions, the instrumental variables technique makes causally correct inferences. The difficulty with the technique is, however, that it is, taken literally, unoperationable. By its very nature, the error term is ε unobservable. Hence it is not possible to test statistically whether Z is or is not uncorrelated with the error term. This a problem in particular because even slight correlations can severely bias the estimator (see Pearl 1993, Bartels 1991).

The irony is that one motivation to use the instrumental variables technique is exactly that there are unobservable common causes between X and Y. Now, if that is so, it seems hard to see how an instrument in the sense of a variable that satisfies **IV-1** and **IV-2** should be found (if a variable is not measurable, a fortiori its correlation with another variable is not measurable).

Stage 3: Assuming Z is a "causal instrumental variable."

https://doi.org/10.1086/508953 Published online by Cambridge University Press

In practice, econometricians identify an instrument on the basis of causal background knowledge, which often derives from an institutional and/or historical analysis of the case at hand. Statisticians sometimes judge background assumption of this kind inadmissible, "subjective."[5] But apart from the fact that the whole approach would be ineffective unless these background assumptions (in addition to the general assumptions) are made, there are two further justifications for making the causal assumptions explicit. First, the structure of assumptions shows that the system Z-X-Y is equivalent to an experimental set-up in which Z is used as an "intervention variable" to test the hypothesis whether X causes Y. Let us call a variable Z a "causal instrumental variable" with respect to (X, Y) if and only if **RP**, **T** and **FC** as well as the following assumptions are satisfied:

CIV-1. Z causes X.

CIV-2. Z causes Y if at all only through X (i.e., not directly or via some other variable).

CIV-3. Z and Y do not have causes in common (except those that might cause Y via Z and X).

It can be proved relatively easily that these assumptions are sufficient for the claim "If a variable Z is a causal instrument and it is correlated with a putative effect Y, then the putative cause X actually causes Y"[6] **CIV-1–3** make the similarity with a controlled experiment very easily visible. A randomised clinical trial (RTC) is a paradigmatic example for a controlled experiment. In an ideal[7] RTC, the random allocation is the causal instrument Z. It (and only it) causes whether or not a subject will receive treatment (X). It does not have a causal influence on recovery (Y) which is not mediated by X. And the allocation itself is not caused by anything that also causes recovery. It is in fact **RP** which ensures that randomization is successful (if the random allocation 'happened' to produce biased groups, the resulting correlation of a cause of recovery with the allocation would have to have a causal explanation, which would violate either **CIV-2** or **CIV-3**).

3. Interventions and Causal Instruments. Conditions **CIV-1**–**3** are very similar to James Woodward's definition of an intervention. According to

5. Even Judea Pearl once seems to have thought this, see his 1993.

6. I have done so in Reiss, forthcoming, Chapter 7. Cf. also Cartwright, forthcoming, who shows that the causal assumptions one must make in order to render Herbert Simon's method of causal inference valid are very similar.

7. I qualify RTC here in order to avoid discussing problems with compliance etc.

https://doi.org/10.1086/508953 Published online by Cambridge University Press

Woodward, a variable I is an intervention variable for X with respect to Y if and only if the following conditions are satisfied:

> I1. I causes X.
>
> I2. I acts as a switch for all the other variables that cause X. That is, certain values of I are such that when I attains those values, X ceases to depend on the values of the other variables that cause X and instead depends only on the value taken by I.
>
> I3. Any directed path from I to Y goes through X. That is, I does not directly cause Y and is not a cause of any causes of Y that are distinct from X except, of course, for those causes of Y, if any, that are built into the I-X-Y connection itself; that is, except for (a) any causes of Y that are effects of X (i.e., variables that are causally between I and X and have no effect on Y independently of X.
>
> I4. I is (statistically) independent of any variable Z that causes Y and that is on a directed path that does not go through X (Woodward 2003, 98).

The differences concern I2, which does not exist in my characterization of a causal instrumental variable, and I4, which differs from **CIV-3** unless further assumptions are fulfilled. In this section I want to discuss the differences between a causal instrument and an intervention with respect to their relevance for econometrics.

I2 demands that an intervention breaks all causal laws that have X as an effect. This condition ensures that any variation in Y that follows the intervention is due to a causal influence of X on Y rather than another cause of X, which happened to 'fire' at the same time as the intervention (e.g., a common cause of X and Y). For practical purposes, I2 appears to be very strong indeed. It presupposes, for example, perfect compliance in a RCT. If the concept of an intervention is not only supposed to illuminate the *concept* of cause but also to help testing for causal relations, the condition is too strong, surely for econometrics applications. Consider, for example, Joshua Angrist's study of the effect of veteran status on civil earnings (Angrist 1990). In this study, Angrist exploits the lottery that assigned a number between 1 and 365 to the birth dates of white men born between 1950 and 1952 randomly. Men from each year were drafted up to a threshold number, depending on the manpower needs of the Defense Department. He uses the random number as an instrument. But the random number raises the probability of being drafted by only 10 to 15 percent (see, for instance, his Table 2, 321). In no way does it stop other causes of the conscription operating.

In the characterization of a causal instrument, the equivalent job is

https://doi.org/10.1086/508953 Published online by Cambridge University Press

done by assuming the Reichenbach principle. If a causal instrument Z is correlated with Y, there can be no other cause of X that is correlated with Z. Suppose there is a common cause V of X and Y. If V was correlated with Z, that would imply that either V causes Z, Z causes V or there is a common cause which causes Z and V. If V causes Z and if there is a common cause, **CIV-3** is violated; if Z causes V, then **CIV-2** is violated.

As mentioned in the previous section, I do not want to defend **RP** as a universal principle. Clearly, there are many counterexamples. However, let me say that controlling for violations of **RP** (and subsequently assuming the principle in the controlled data) is more consistent with econometric practice than Woodward's assumption I2. Elliott Sober, famously, pointed out that the correlation between British bread prices and Venetian Sea levels constitutes a violation of **RP** (e.g., Sober 2001). As Kevin Hoover (2003) remarked, however, econometricians would not regard the correlation between the two variables as indicative of a causal connection. Rather, they would regard these two time-series as non-stationary and use co-integration as the appropriate measure of probabilistic dependence in this case. Thus only if the two series were co-integrated, they would start looking for a causal connection.

Similar things can be said about other types of 'brute' correlations. For example, minimum wage legislation can be used as an instrument to test the claim whether an increase in the minimum wage causes employment to change (see, for instance, Card and Krueger 1995). In such studies the effect of economic conditions on employment is carefully controlled for. Importantly, this happens independently of whether or not the wage bill is passed in response to favourable economic conditions or merely accidentally at the same time. An accidental correlation between the two events would constitute a violation of **RP**. But prudent econometricians do not let that happen.

The other difference between an intervention and a causal instrument concerns the difference between I4 and **CIV-3**. This difference, too, can be explained by the assumption of **RP** in the definition of a causal instrument. In fact, under **RP**, the two conditions are equivalent. Without **RP**, Woodward must assume that I is probabilistically independent of other causes of C that are not on the route I-X-Y, because otherwise the variation in Y that follows the intervention could be due to C rather than the causal effect of X on Y. But again, I want to claim that **CIV-3** in conjunction with controlling for violations of **RP** is more in line with econometric practice than is I4. In practice, the assumptions about the absence of a probabilistic dependence are made on the basis of causal considerations plus (at least, implicitly), **RP**. In few cases, thus, I4 can be assumed without grounding it in **CIV-3** and **RP**.

My conclusion is therefore that testing claims using causal instruments

https://doi.org/10.1086/508953 Published online by Cambridge University Press

is a more practical affair than testing them using interventions. There may be a good reason why this is so. Woodward explicitly states that he regards the aim of his analysis to be an illumination of the concept of cause and of the truth conditions for causal claims rather than a provision of feasible *tests* (e.g., 95). I have nothing to say here about the metaphysics of causation. But I do agree that as a test for causal claims, Woodward's conditions fare less well than the econometricians' conditions for causal instruments.

REFERENCES

Angrist, Joshua (1990), "Lifetime Earnings and the Vietnam Era Draft Lottery: Evidence from Social Security Administrative Records", *American Economic Review* 80: 313–336.

Bartels, Larry (1991), "Instrumental and 'Quasi-Instrumental' Variables", *American Journal of Political Science* 35: 777–800.

Card, David, and Alan Krueger (1995), *Myth and Measurement: The New Economics of the Minimum Wage*. Princeton, NJ: Princeton University Press.

Cartwright, Nancy (forthcoming), "Causal Inference à la Herbert Simon: A Primer", in Nancy Cartwright, *Hunting Causes—and Using Them*. Cambridge: Cambridge University Press.

Greene, William (2000), *Econometric Analysis*. Upper Saddle River, NJ: Prentice-Hall.

Hammond, J. Daniel (1996), *Theory and Measurement: Causality Issues in Milton Friedman's Economics*. Cambridge: Cambridge University Press.

Hausman, Daniel, and James Woodward (2004), "Modularity and the Causal Markov Condition: A Restatement", *British Journal for the Philosophy of Science* 55: 147–162.

Hoover, Kevin (2003), "Nonstationary Time Series, Cointegration, and the Principle of the Common Cause", *British Journal for the Philosophy of Science* 54: 527–551.

McDermott, Michael (1995), "Redundant Causation", *British Journal for the Philosophy of Science* 46: 523–544.

Mitchell, Wesley (1927), *Business Cycles: The Problem and Its Setting*. New York: National Bureau of Economic Research.

Pearl, Judea (1993), "Mediating Instrumental Variables", *Statistical Science* 8: 266–273.

——— (1997), "The New Challenge: From a Century of Statistics to an Age of Causation", manuscript, University of California, Los Angeles.

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

Pearson, Karl (1911), *The Grammar of Science*. London: Walter Scott.

Reiss, Julian (forthcoming), *Error in Economics: Toward a More Evidence-Based Methodology*. London: Routledge.

Sober, Elliott (2001), "Venetian Sea Levels, British Bread Prices, and the Principle of the Common Cause", *British Journal for the Philosophy of Science* 52: 331–346.

Woodward, James (2003), *Making Things Happen*. Oxford: Oxford University Press.

https://doi.org/10.1086/508953 Published online by Cambridge University Press