Southern Africa Labour and Development Research Unit

Econometric methods and Reichenbach's principle

by

Sean Muller

Working Paper Series
Number 85

About the Author(s) and Acknowledgments

I am grateful to Martin Wittenberg for comments on an earlier draft of this paper. Comments from four anonymous referees were valuable in the process of clarifying various aspects of the arguments. The paper also benefited from presentation at a SALDRU seminar at the University of Cape Town. All views and errors remain my own.

Recommended citation

Muller, S. (2012). Econometric methods and Reichenbach's principle. A Southern Africa Labour and Development Research Unit Working Paper Number 85. Cape Town: SALDRU, University of Cape Town

ISBN: 978-1-920517-26-7

Working Papers can be downloaded in Adobe Acrobat format from www.saldru.uct.ac.za.
Printed copies of Working Papers are available for R15.00 each plus vat and postage charges.

Orders may be directed to:
The Administrative Officer, SALDRU, University of Cape Town, Private Bag, Rondebosch, 7701,
Tel: (021) 650 5696, Fax: (021) 650 5697, Email: brenda.adams@uct.ac.za

Econometric methods and Reichenbach's principle[1]

Reichenbach's 'principle of the common cause' is a foundational assumption of some important recent contributions to quantitative social science methodology but no similar principle appears in econometrics. Reiss (2005) has argued that the principle is necessary for instrumental variables methods in econometrics, and Pearl (2009) builds a framework using it that he proposes as a means of resolving an important methodological dispute among econometricians. We aim to show, through analysis of the main problem instrumental variables methods are used to resolve, that the relationship of the principle to econometric methods is more nuanced than implied by previous work, but nevertheless may make a valuable contribution to the coherence and validity of existing methods.

Keywords: Reichenbach's principle, econometrics, causality.

1 Introduction: Reichenbach's principle and microeconometrics

Reichenbach's self-titled 'principle of the common cause' is concerned with the explanation of improbable coincidences; "If an improbable coincidence has occurred, there must exist a common cause" (Reichenbach 1991, 157). Determined by frequency of occurrence, one might represent an improbable coincidence in

[1]I am grateful to Martin Wittenberg for comments on an earlier draft of this paper. Comments from four anonymous referees were valuable in the process of clarifying various aspects of the arguments. The paper also benefited from presentation at a SALDRU seminar at the University of Cape Town. All views and errors remain my own.

probability terms as: $P(A \wedge B) > P(A)P(B)$.[2] When Reichenbach refers to a common cause, C, 'explaining' this coincidence he means: $P(A \wedge B|C) = P(A|C)P(B|C)$ (Reichenbach 1991, 159). In short, conditional on the common cause the two events are statistically independent. The implicit assumption is that, by virtue of temporal simultaneity of A and B, neither event causes the other. Hence 'Reichenbach's principle' (henceforth RP) is often formulated as: 'given a statistically significant correlation between two events, either one event is the cause of the other, or they share a common cause (or some combination of these)'.

Subsequent analysis has suggested the principle is not true in general; there exist cases in which sigificant correlations between variables cannot be attributed to a causal relationship. Arntzenius (1992, 2010) provides an overview of the merits of RP, including counterexamples. As regards physics these concern quantum phenomena and laws of coexistence, while problems possibly relevant to social sciences concern time-series processes or deterministic systems. The validity of the principle therefore appears to be domain-specific. As regards social science, RP appears to clash with a popular mantra among economists (and others) that "correlation does not imply causation". However, the immediate tension is superficial: the mantra states that correlation between two variables need not imply that one causes the other, which is consistent with Reichenbach's suggestion that correlation may arise from a common cause. The difference is in emphasis: the former suggests causal inference may proceed from correlations, while the latter

[2]Where $P(A \wedge B)$ is the probability of A and B both occuring. We will use $P(A|C)$ to represent the probability of A given that C is known to have occurred.

emphasises caution in doing this. This critical distinction is reflected in method. For instance, one might contrast Spirtes, Glymour, and Scheines (2000)'s use of RP as part of an axiomatic foundation for a generic, algorithmic approach to establishing causal relations within cross-sectional datasets, with a general suspicion of simple correlations in applied work in social science.

The modern literature on microeconometrics - the development and application of statistical methods for empirical analysis of microeconomic issues - is primarily concerned with empirical identification of plausibly unconfounded effects of one variable of interest on another, sometimes referred to as 'the identification problem'. However, limits to observational data on economic systems are such that resolving this problem by statistically controlling for all possible confounding factors is seen as unlikely (Wooldridge 2002a, 3-4). One solution, increasingly presented using the counterfactual-based 'Rubin Causal Model' - see Angrist, Imbens, and Rubin (1996) - is to utilise a source of 'exogenous' variation in the explanatory variable of interest. That variation can be constructed - as in the case of a randomised control trial (RCT) - or the result of a 'natural experiment' that provides 'serendipitous randomization' (see DiNardo (2008), or Rosenzweig and Wolpin (2000) on "natural 'natural experiments'"). Where randomization has not occurred, researchers may use 'quasi-random' variation in which there is a source of variation that is not strictly random but is 'plausibly exogenous' under certain assumptions or conditions.[3] A variable representing the source of such variation

[3] As Rosenzweig and Wolpin (2000) put it, "This approach essentially *assumes* that some components of nonexperimental data are random"(our emphasis, Rosenzweig and Wolpin 2000, 827).

is one form of an 'instrumental variable', a formal definition of which is provided below.

Although the willingness to explicitly connect identification with claims about causal relationships has varied over the history of econometrics, current confidence in methods like those above is such that Angrist and Pischke (2009) frame reluctance to do this as characteristic of a statistician rather than an econometrician.[4] This appetite for causal claims has, however, not been accompanied by engagement with issues identified by philosophers. For instance, Holland (1986)'s explication of the Rubin Causal Model notes that the notion of causality underlying that approach can be encapsulated by the mantra 'no causation without manipulation' (Holland 1986, 959), highlighting the relevance of manipulation accounts of causality such as Woodward (2003), and the interventionist account developed in statistics by Pearl (2009). That in turn indicates the relevance of Nancy Cartwright's *criticisms* of such accounts and claims regarding randomised trials; see for instance Cartwright (2007) and Cartwright (2010), respectively. While that work has received some recognition in intra-disciplinary debates regarding the merits of RCT-based empirical analysis - most notably Deaton (2008, 2010) - other authors and philosophical concerns per se are absent. One notable lacuna for philosophers of science is the wholesale omission of Reichenbach's principle.

[4]For this and other reasons we therefore disagree with Boland (2010)'s claim - in a review of Cartwright (2007) - that causality is not an issue in modern economics.

RP's relevance for econometrics has received some philosophical attention in relation to a counterexample proposed by Sober (2001), see for instance Hoover (2001, 2003), Steel (2003), Reiss (2007) and Hoover (2009). As Sobel's counterexample relates to processes with 'similar laws of evolution' (Arntzenius 1992) generating non-causal correlations across time, those contributions are focused on *macro*econometrics. While macroeconometrics is mostly concerned with datasets containing observations over time ('large T, small N' in econometricians' parlance), microeconometrics focuses on single, or repeated, cross-sections ('large N, small T'). This distinction is not absolute - there have been analogous developments in both areas, as made explicit in Heckman (2000)'s valuable history - but it represents an important sub-division within the discipline that, if ignored, could lead to confusion and inappropriate conclusions or emphasis.

Our insistence on such distinctions reflects a broader position: demonstrating the conceptual and practical importance of philosophical issues for econometrics requires a full appreciation of the existing rationale for, and substance of, existing methods. Further to the macro-micro distinction, one should note the two main, rhetorically conflictual, stances on microeconometric method at present.[5] The first stance, associated with researchers Deaton (2008) named 'randomistas', emphasises finding or creating sources of random or quasi-random variation. Some have claimed that this experimentalist approach frees empirical analysis from reliance on strong a priori assumptions based on economic theory with dubious foundations. That has been strongly contested by the opposing group - for instance Keane

[5] At least three journal issues - *Journal of Econometrics*(2010, 156(1)), *Journal of Economic Perspectives*(2010, 24(2)) and *Journal of Economic Literature*(2010, 48(2)) - have dealt with aspects of this debate.

(2005, 2010b,a)'s work disputing that RCTs enable 'atheoretical econometrics' - who we may refer to as 'structural econometricians'. The preferred method of these researchers begins with specification of a so-called 'structural model' that, in particular, makes explicit assumptions about the decision-making behaviour of economic agents and therefore draws on, or even develops, economic theory. This *potentially* allows estimation of causal parameters with *observational* data, but only if the structural model is, in an appropriate sense, correctly specified.

One need not subscribe strictly to either view; estimation of structural models may benefit from randomised experiments, and estimated effects of RCTs could have greater external validity if behaviour of economic agents is accounted for. An *ideal* RCT suffices for valid causal inference of a certain kind, but actual RCTs in economics may deviate from the ideal in various ways - particularly selection into treatment groups, variation in 'compliance' within those groups and links between these issues and economic theory - thereby providing traction for the arguments of structural econometricians. While the rhetorical gap has narrowed recently, it remains to be seen how this is reflected in the proportion of studies using the different methods.

An important point regarding both approaches is that if direct randomization is not possible, researchers may instead randomise change in a variable that *affects the explanatory variable* and proceed by using instrumental variable methods. In short, these methods present economists with the possibility of causal inference from experimental *and* observational data. This brings us to Reiss (2005, 2008)'s argument that RP is necessary for, and possibly implicit in, economists' instrumental variables method. Given that method's importance in the discipline this is a weighty claim, and our first contribution will be to argue that it is premised on a mistaken understanding of the logic of causal inference in econometric methods. However, we will argue that RP does provide important methodological insights into instrumental variable methods.

Our second contribution concerns an important branch of technical work from outside economics that engages with micro-level quantitative empirical methods in social science: the work by Spirtes et al. (2000) and Pearl (2009) on causal graphs. While a significant component of that work concerns development of algorithms for identifying causal relationships in datasets (observational or otherwise), this is not especially relevant for our purposes here. The more important aspect is that these approaches are explicitly premised on Reichenbach's principle. Arntzenius (1992), for instance, has stated - in reference to the work of Spirtes et al. - that "[the common cause principle] appears to be an indispensable part of the *best* method for inferring causal structure from statistical data in the social sciences" (emphasis added, Arntzenius 1992, 234). Furthermore, regarding the dispute between randomistas and structural econometricians, Pearl argues that the conceptual framework for causal inference he develops offers "a simple and precise unification of these two antagonistic and narrowly focused schools of econometric research" (Pearl 2009, 379). While we agree with a number of Pearl's arguments, the breadth of this claim is such that it cannot be assessed here. Instead, we focus on the importance of RP as a key axiom of his framework. Using a specific example - economists' approach to correlations among explanatory variables in regressions, known as 'multicollinearity' - we aim to show that Pearl's proposal, by virtue of assuming RP, has implications for economics beyond those he discusses. As we explain below, multicollinearity is the obverse of the instrumental variables case, and therefore complements our discussion of Reiss.

The remainder of the paper is structured as follows. Section 1 extends the disciplinary background above with a short technical introduction to regression and instrumental variables methods; economists, or philosophers already familiar with econometrics, will likely want to glance over this section. Section 2 examines the work by Reiss (2005, 2008), and shows that at base it is premised on a misunderstanding of economists' stated rationale for instrumental variable solutions to the identification problem. Nevertheless, we argue that Reiss's concerns are somewhat vindicated by economists' choice and justification of instruments. Finally, section 3 considers the related example of collinearity between explanatory variables and, in particular, the implication of RP for the interpretation of regression coefficients. Our objective is not to argue for or against adoption of the principle, but rather to give an idea of what is at stake and to contest and clarify some of the extant analysis in the hope that this may contribute to a more satisfying understanding of the implications of Reichenbach's principle for econometrics than has been the case to date.

Regression, instrumental variables and causal inference

For the analysis that follows, we provide a basic introduction to economists' approach to regression and instrumental variables analysis. Woodward (1988) gives what one might call a 'traditional' overview of the formalities of regression methods directed at philosophers, while our discussion relies more on the presentations by Manski (1991) and Wooldridge (2002a). Though we do not directly address Woodward's concerns and conclusions regarding regression as a means of causal inference, our analyses are broadly compatible.[6]

A univariate ('simple') 'regression' refers to representation of the mean of one random variable, y, conditional on another random variable, x, as a function of the latter variable. By conditionality we mean: what is the average value of y given that x takes on some specific value x_0? We can write this as: $E(y|x = x_0)$. A mean regression expresses the *variable* $E(y|x)$, representing all values taken by y, as a function of x: $E(y|x) = f(x, \beta)$, where beta represents the parameters of $f(\cdot)$. There is a clear asymmetry of *interest*, such that x is referred to as the 'explanatory' variable and y as the 'dependent' variable. In the case of *multivariate* (or 'multiple') regression, we instead have vectors $\mathbf{x}$ and $\boldsymbol{\beta}$ representing multiple explanatory variables and associated parameters. For instance, if we assume the

[6] A valuable additional reference is Morgan and Winship (2007)'s book, which provides an overview of regression and graph-based methods within a detailed discussion of causal inference based on counterfactuals. This includes a brief introduction to work at the frontier of econometrics by Heckman and coauthors that attempts to integrate the approaches of the structural and randomista camps - see for instance Heckman and Vytlacil (2007a,b).

function to be linear in the parameters and variables, we can write:[7]

$$E(y|\mathbf{x}) = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \cdots + \beta_k x_k \tag{1}$$

Any representation where $f(\cdot)$ is *linear in the parameters* and $k > 1$ is known as a 'multiple linear regression'. We need not be interested only in the mean of the conditional distribution. Other properties like the median or variance may be of interest (see Manski (1991)), where the relevant regression would replace $E(y|x)$ with the property of interest, say $var(y|x)$. Furthermore, there is no general reason to assume linearity of $f(\cdot)$. The strength and nature of the assumptions determines the methods available for empirical estimation. Methods for the case where no specific form for $f(\cdot)$ is assumed are often referred to as 'nonparametric' methods.[8] Nevertheless, empirical work in the social sciences, including economics, remains dominated by use of mean regressions and for ease of exposition we will focus on these, a linear *in the parameters* functional form and only briefly mention one form of estimation (the 'method of moments').

So far we have said nothing of causality, nor is it necessary to do so. Regression may be descriptive and the imposed asymmetry need not represent anything

[7]Linearity in the parameters allows $E(y|x)$ to be a non-linear function of the explanatory variables. E.g. We could have $E(y|\mathbf{x}) = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_2^2$.

[8]As Manski notes, this terminology is somewhat misleading: "Use of the term nonparametric to mean that the parameter space is a space of functions is an illogical but firmly entrenched semantic convention"(Manski 2007, 34).

besides the researcher's interest. However, as noted, modern econometric analysis is interested in more than mere description of associations between variables. Consequently, it usually begins with a *structural* equation of the dependent variable of interest (y) that is explicitly or implicitly causal.[9] A note on terminology: *structural equations* may be based on explicit structural *models* but need not be. And conceptually there are two categories of structural equations: those used in discussions of method that are, for that purpose, 'true', and equations that represent *hypotheses* about the underlying structure. The latter are the starting point for empirical analysis, while the former are the starting point for methodological analysis, and the *estimated* equation could look very different from either.

To get to an 'estimable' equation one can use the fact that it is always possible to decompose the dependent variable in 'error form' (Wooldridge 2002a, 18) as:

$$y = E(y|\mathbf{x}) + u, \tag{2}$$

where as a matter of *definition*: $E(u|\mathbf{x}) = 0$. This implies, in particular, that u is uncorrelated with all explanatory variables and any function thereof. If we then assume equation (1) to be true, we can write the structural equation:

$$y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \cdots \beta_k x_k + u \tag{3}$$

[9]The *implicit* assumption of causality, whether in the presentation of methods or their application, is usually observed in the interpretation or use of the estimates from the subsequent empirical analysis, as noted by Woodward (1988).

Given: the properties of u, which now reflect assumptions about the correctness of (1) as a representation of the underlying structure (correct functional form and explanatory variables); no linear dependence between any of the explanatory variables; and all variables observable, the parameter vector $\boldsymbol{\beta}$ is said to be 'identified'. Specifically, "$\boldsymbol{\beta}$ can be written in terms of population moments in observable variables"(Wooldridge 2002a, 53). Using the vector of explanatory variables $\mathbf{x}$, we can write:[10]

$$\beta = E[\mathbf{x}^T\mathbf{x}]^{-1}E(\mathbf{x}^T y) \tag{4}$$

In the usual case researchers do not have data on the entire population but rather a random sample with N observations. Methods of moments estimation is based on replacing the 'population' means in (4) with their sample analogues. Doing this gives (5), known as the ordinary least squares (OLS) estimator.[11]

$$\hat{\beta} = \left(\frac{1}{N}\sum_{i=1}^{N}\mathbf{x}_i^T\mathbf{x}_i\right)^{-1}\left(\frac{1}{N}\sum_{i=1}^{N}\mathbf{x}_i^T y_i\right) \tag{5}$$

When economists refer to 'running a regression' they are typically referring to this final process of estimating parameters of an equation based on some pre-

[10]Rewrite (3) as $y = \mathbf{x}\beta + u$, premultiply both sides by the transposed vector $\mathbf{x}^T$, apply the expectations operator to both sides and solve for $\boldsymbol{\beta}$ (noting that $E[\mathbf{x}^T\mathbf{u}] = 0$).

[11]The reason for the name is that $\hat{\beta}$ can be shown to be the solution of a loss minimisation problem where the chosen loss function is the square of the deviations from the prediction - see Manski (1991).

specified structural conditional expectation. Whether that estimates causal parameters depends on how accurately the estimable equation represents the true structure. There are various obstacles to identification, most of which imply correlation between the error term and one or more of the explanatory variables *in the equation used for estimation* - also referred to as an 'endogeneity problem'. This leads to bias in estimates: $E(\hat{\beta}) \neq \beta$. In what follows we will use one important form of endogeneity known as 'omitted variables bias' to illustrate the instrumental variables method.

As noted by Woodward (1988, 259), a key decision in specifying an estimable regression equation is determining which variables should be included. Economists take the view that mistakenly including a variable not in the true structural equation is typically less problematic than excluding a relevant variable. This could be due to a flaw in economists' a priori theory, or because a given variable is not empirically observable. A popular example of the latter is individuals' intrinsic ability where a researcher is interested in the effect of education on earnings. Ability is hypothesised to affect educational attainment *and* affect earnings directly, but is unobservable and consequently acts as a confounding factor.

To illustrate the general case, assume (6) is the true structural equation, where q represents one or more variables that will be omitted from the final estimated equation. We can rewrite this as an estimable equation like (4) but with error term $v = u + \gamma q$. For v to satisfy the same conditions as u - allowing least-squares estimation of $\boldsymbol{\beta}$ - q may not be correlated with any elements of $\mathbf{x}$.

$$E(y|\mathbf{x}, q) = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \cdots \beta_k x_k + \gamma q \quad (6)$$

In the event of correlation with q one can show that the estimated parameter $\hat{\beta}_j$ will be biased in a direction dependent on the signs of $corr(q, x_j)$ and β_j. A primary motivation of experiments that randomise realisations of x_j is precisely that - in the ideal case - this will sever any structural connections between x_j and q.[12] It is this method that is closest to the philosophical assumption of modularity and definition of causation through manipulation advocated by Hausman and Woodward (1999) and Holland (1986) among others. For simplicity, however, our analysis will focus on the so-called 'instrumental variables' (henceforth, IV) solution to the omitted variable problem using *observational* data.

The *theoretical* IV solution is to utilise a variable (z) *not* in the true structural equation, but (conditionally) correlated with the endogenous variable (here x_j) and *not* correlated with the omitted variable(s) (here q). The latter two requirements are often stated formally as:

IV1 $corr(z, x_j|\mathbf{x_{\neg j}}) \neq 0$

IV2 $corr(z, v) = 0$

Where $\mathbf{x_{\neg j}}$ is a vector containing all explanatory variables *except* x_j. The first condition, less often formalised, implies that the instrument be 'redundant' in ex-

[12]There is some disagreement about how this severance of relationships should be represented; see Pearl (2009, 376-377).

plaining variation in the dependent variable given the other explanatory variables:

IV3 $E(y|\mathbf{x}, q, z) = E(y|\mathbf{x}, q)$

Figure 1 illustrates this scenario using causal graphs.[13]

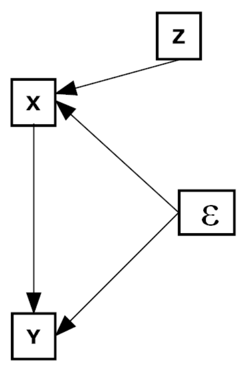

Figure 1: The standard instrumental variables scenario

Given IV1-IV3 one can write an estimable equation with an error term satisfying the same conditions as in the standard regression case so that $\boldsymbol{\beta}$ is identified as:

$$\beta = E[\mathbf{z}^T \mathbf{x}]^{-1} E(\mathbf{z}^T y) \tag{7}$$

Where $\mathbf{z}$ is the $\mathbf{x}$ vector *including* z and *excluding* x_j. To estimate this one could substitute the sample analogues as before. A key point is that while *IV1* and

[13] This is identical in structure to a graph by Pearl (Figure 7.8(a) 2009, 248).

IV3 can be tested empirically, *IV2* cannot because it concerns the unobservable term u.[14] Consequently, economists rely on qualitative 'stories' supporting the validity of a given instrument. As Murray (2006) notes, "all instruments arrive on the scene with a dark cloud of invalidity hanging overhead. This cloud never goes entirely away, but researchers should chase away as much of the cloud as they can"(Murray 2006, 114).[15]

Besides the details, what this abbreviated discussion should make clear is that microeconometric methods for non-experimental data proceed from specific, ex ante assumptions about the true underlying structural relationships. As we will see, this is key to understanding the strengths and weaknesses of these methods.

2 Instrumental variable methods do not require Reichenbach's principle

If philosophical analyses yield genuinely important insights into econometric methods, there should be demonstrable implications for empirical analysis, and these are likely to be necessary to interest practitioners in philosophical work. Some au-

[14] Indeed *IV1* is a serious concern in the empirical literature because of theoretical results showing the negative consequences of 'weak instruments' (small value for $corr(z, x_j|\mathbf{x_{\neg j}})$).

[15] As another example, in a widely-cited survey article, Angrist and Krueger (2001, 73) state: "In our view, good instruments often come from detailed knowledge of the economic mechanism and institutions determining the regressor of interest".

thors have, admirably, attempted to grasp this nettle - Cartwright (1999), Hoover (2001) and Reiss (2008) being just a few notable examples. To do this successfully is of course very hard, partly because it requires a detailed knowledge of both literatures, including implicit assumptions in economic methodology against which philosophical claims are to be measured. Below we address our first concern relating to existing work on RP and econometrics, namely the claim by Reiss (2005) that instrumental variables analysis requires RP, and in doing so we also seek to clarify a few potential misunderstandings regarding econometric methods.

Reiss (2005)'s basic argument - also in Reiss (2008, 126-145) - is that the IV logic is flawed because the two key criteria typically formalised - IV1 and IV2 - can be satisfied without identifying a genuine causal relationship between y and x. Given this, Reiss proposes three additional sets of assumptions that would justify causal claims based on instrumental variables and divides these into 'stages' of analysis. We state these now for reference:

Stage 1 Assume Reichenbach's principle ($\mathbf{RP}$), causal transitivity ($\mathbf{T}$) and 'functional correctness' ($\mathbf{FC}$);[16]

Stage 2 Assume the structural error term includes all causes of the dependent variable not specified explicitly in the structural equation, unless those causes

[16] Reiss (2005, 969) defines functional correctness as: "A structural equation is functionally correct if and only if it represents the true functional (but not necessarily causal) relations among its variables".

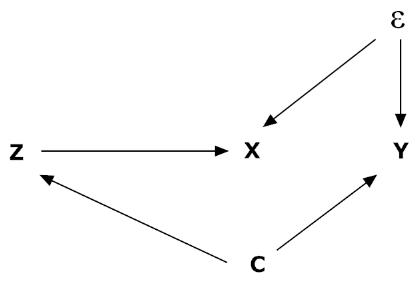

Figure 2: Reiss (2005)'s counterexample

work through a specified variable;

Stage 3 Assume that the instrument is a 'causal instrumental variable'.

Following Stage 1, Reiss proposes the causal system represented graphically in Figure 2 as a counterexample to the claim that IV1 and IV2 suffice to identify the coefficient on X in the structural equation.

That Reiss's stage 2 assumption is the *standard* way of interpreting the error term in a *hypothesised* structural equation, and the corresponding estimated equation, already gives cause for concern. Econometric textbooks often make this interpretation explicit - see for instance (Greene 2003, 8), it is recognised in Woodward (1988, 261) and Pearl also endorses a conceptual understanding of such error terms as representing omitted factors since it is a useful guide "when building, evaluating and thinking about causal models" (Pearl 2009, 162-163).

A second issue is that Reiss is interested in *causes* whereas one can describe the IV logic without reference to causality per se. As Reiss notes, "many textbooks contain 'recipes' for econometric inference that give the impression that econometrics can proceed without causal background assumptions" (Reiss 2005, 966). In fact, econometrics can proceed without such assumptions, but - as per Cartwright's mantra: 'no causes in, no causes out' - not if the interest is in *causal* inference. Nevertheless, there is substance to Reiss's concern as interpretations in the discipline of 'structural' equations have not always been consistent or explicit. The asymmetry of regression equations is often implicitly causal even though it need not be.[17]

The primary problem with Reiss' analysis, however, is the neglect of additional implicit, or definitional, assumptions in addition to the explicitly stated IV1 and IV2.[18] Two assumptions implicit in (6) are: That q and each x_i, $i = 1 \ldots k$, have independent explanatory power for y; and, that no other observed or unobserved variable associated with any of the *specified* explanatory variables will have independent explanatory power for y.[19]

[17]See for instance the comment by Fennell (2007) in the context of systems of equations.

[18]The word 'definitional' is intended to indicate that these assumptions are not 'implicit' in the sense of being wholly unstated (as Reiss suggests later in his paper). Rather, they follow directly from initial definitions of the problem, such as specification of the structural equations.

[19]The second assumption could be strengthened by requiring that the structural equation specify all causes, but this is not necessary for causal inference.

This oversight manifests in the counterexample in Figure 2. In that system, the variables X and Y have a common cause (ϵ) so that their correlation is 'spurious', while Y and Z also have a common cause (C) and Z is a cause of X. Recall the question in omitted variables analysis: is the coefficient on a variable that *genuinely belongs in the structural equation* empirically biased because of correlation with another relevant, but omitted, variable? By contrast, Figure 2 represents a scenario where a researcher seeks to instrument for an *irrelevant* variable in the estimable equation using one that is correlated with *another* omitted factor (C). It is true that this is disastrous for causal inference. However, that is due to a failure of the economist's 'extra-statistical' knowledge, which manifests - through a hypothesised structural equation - in an estimable equation that does not achieve identification of the parameter(s) of interest. It is not a failure of the method per se.

To be specific: the hypothesised cause (X) has, in fact, no causal role, and the common cause structure the $Z \leftarrow C \rightarrow Y$ fork implies is disallowed by the definition of a structural equation as made explicit in IV3.[20] This oversight is related to misunderstanding the properties of error terms in (true) structural and estimated regression equations. Indeed, as Reiss notes, if "the error terms in an equation

[20]The fact that an instrument cannot belong in the true structural equation is an assumption made clear in a number of texts, e.g. Wooldridge (2002b, 517) and Pearl (Figure 7.8 (d), 2009, 248). Reiss actually *proposes* this later in that paper (see condition CIV-2, Reiss (2005, 973)) as one of a set of assumptions that would justify the econometrician's approach, but this assumption clearly *is* made both in theory and in practice.

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." (Reiss 2005, 971).

'Extra-statistical' assumptions and Reichenbach's principle

These points should further emphasise our earlier statement regarding the extent to which microeconometricians currently rely on ex ante assumptions abut causal structure for causal inference. As Woodward (1988, 1995) has noted, "These...assumptions are commonly described as "a priori" or "extrastatistical," where what this means is not that they are non-empirical or incapable of being tested, but rather that they are not inferred just from the statistical data at hand, but rather have at least in part some other rationale or justification" (Woodward 1988, 259).[21] The source of Reiss's problem may be a subtle confusion of (micro)economists' methods of causal discovery with Pearl (2009) and Spirtes et al. (2000)'s algorithmic approach to 'hunting causes'. When it comes to non-experimental data, economists rely heavily on a priori 'extra-statistical' assumptions based on theory or some kind of professional intuition. Graph-based algorithms, by contrast, show that under a set of core, generic assumptions - like RP - causal structure can, to some extent, be inferred from purely statistical information. This is, in part, what Cartwright takes umbrage with in arguing that the notion of causation is *not* generic, and one cannot draw causal conclusions

[21] We exclude Woodward's reference to these as 'causal' assumptions since that remains a moot point.

without substantive assumptions about causal structure. As relates to common cause assumptions like RP, the analysis of Arntzenius (1992) would appear to indirectly support Cartwright's stance, since he demonstrates that there exist convincing counterexamples to RP, and yet that these are sufficiently domain-specific for there to exist systems of interest under which it is a valid assumption.

The difference in emphasis of the two approaches is directly connected to the necessity of RP. What Reiss has done is construct a system where the two basic correlative relationships in instrumental variable analysis are satisfied. The axioms $\mathbf{RP}$, $\mathbf{T}$ and $\mathbf{FC}$ render this a causal system with particular properties. A counterexample is then constructed to show that such correlations can exist without supporting the conclusions of the instrumental variables method. If economists applied IV methods in a mechanical fashion based on sample statistics to obtain supposedly causal parameters the example might be justified. However, economists tend to be uninterested in algorithmic approaches, as can been seen in the response - or lack thereof - to Spirtes et al. (2000), Pearl (2009) and in macroeconometrics to the general-to-specific modelling algorithm of Hendry and Krolzig (2005).[22] While Spirtes et al. (2000) and Pearl (2009) place the burden of causal structure on RP and other generic assumptions, economists rely on specific theoretical 'knowledge'. Consequently, while there has been little work on this point, if anything it is the lack of a sound methodological foundation for *theoretical* development, rather

[22] The actual merit of algorithmic approaches in social science is a separate, contentious issue; see the contributions to McKim and Turner (1997) and the discussion of 'automated discovery' by Glymour (2004).

than empirical method, that may turn out to be the Achilles heel of causal inference in economics. And it has been in part the dissatisfaction with that aspect of the discipline that led many empirical researchers to methods based on experiments that, initially, promised less reliance on a priori theoretical assumptions. The work by Keane (2010a) against the possibility of 'atheoretical' econometrics, Heckman and Vytlacil (2007b) and others, has undermined that hope and whether the basic concern is assuaged or manifests in some other form remains to be seen.

Instruments as causes

A final aspect of Reiss' analysis that merits additional consideration is the third stage in which he advocates for an explicitly causal interpretation of IVs. Although in principle one could write a *structural* equation for an endogenous variable with the instrument on the right-hand side, in expositions of the method the focus is only on statistical properties of instruments. Consider a popular alternative statement of IV1. Given data from a sample of a population, we can always estimate a linear regression regardless of whether it will have any causal meaning. Similarly, for a hypothetical population we can always write a linear projection of one variable on a set of other variables. In our case the linear projection of interest is:

$$x_j = \gamma_0 + \boldsymbol{\gamma}\mathbf{x}_{\neg\mathbf{j}} + \theta z + \varphi \tag{8}$$

By the *definition* of a linear projection, the parameters are such that the error term has the same properties as the error term in the structural equation: $E[\varphi|\mathbf{z}] = 0$. The key difference is that in the structural equation the error has these properties by *assumption* (it is assumed to reflect the causal structure), whereas in the linear projection case it is by *construction*. IV1 is equivalent to requiring $\theta \neq 0$ in (8). Consequently, a common empirical test of IV1 is to estimate such a regression equation and test the hypothesis that $\theta = 0$. However, unlike for the original structural equation, no causal foundation is provided or implied by (8). To the contrary, as Wooldridge emphasises: "there is nothing necessarily structural about [the] equation" (Wooldridge 2002a, 84). This equivocation about causal structure in relation to instruments seems at odds with the basic logic of causal inference in econometrics explained in section 1.[23] Absent *any* commitment by economists to the causal character of relationships between the endogenous, confounding and instrumental variables, we might ask what difference assuming RP would make.

RP implies that if an instrument satisfies IV1 it could be a cause or effect of the endogenous variable, they could share a common cause, or some combination of these. Assume for simplicity that the confounding factor is a common cause of x and y - as seems to be the case for most illustrative examples in economics. Then causal transitivity rules-out the instrument being caused by x, since that would imply correlation between z and the omitted factor, violating IV2. The original contribution of Reiss's paper may be the point that if economists fail to address the (potentially) causal origins of endogeneity, they cannot convincingly make a

[23]It may be useful for some readers to note that within the discipline it is known that implicit in such presentations in the econometrics literature is that if y was a cause of any explanatory variables - i.e. there was 'simultaneous causation' - then an explicit *system* of equations would be required. While such systems are popular examples in the philosophical literature, partly to engage with arguments of early econometricians and other scholars, they are emphasised less in modern presentations of econometric methods.

causal case for instrument validity. For instance, if economists accept $\mathbf{RP}$ and $\mathbf{T}$ then, as a methodological point, where the endogenous explanatory variable shares the omitted variable as a common cause with y valid instruments *must*:

1. Be causes of the endogenous variable, or;[24]

2. Share a common cause with the endogenous variable.

As an example, consider Murray (2006)'s illustration of his discussion of IV methods with well-known work by Levitt (1997, 2002) that attempts to estimate the causal effect of changes in police numbers on crime. Because police numbers and crime rates could have common causes, Levitt uses two separate instrumentation strategies: first, he uses local election dates as an instrument for police numbers, arguing that police numbers increase before elections - the instrument is effectively posited as a direct cause; second, he uses the number of firefighters as an IV arguing that this will change along with police numbers due to budgetary changes - the instrument shares a common cause with the endogenous variable.

That most instruments and the 'stories' told to support them in the literature follow this logic from common cause endogeneity to causal instruments would seem to support Reiss's position and arguably make it a worthwhile addition to the standard textbook account. That suggestion is likely to be resisted by economists for

[24]Reiss makes an argument somewhat along these lines, including a discussion of a related, intervention-based, approach by Woodward (2003) - see Reiss (2005, 972-974) - but it is afflicted by some of the misunderstandings already described.

a number of reasons, of which we mention two here that have actually been made in response to the basic arguments of the present paper. First, that the conclusion regarding the causal relationship between instruments and other variables is self-evident; it is simply *obvious* that a valid instrument cannot be caused by the endogenous variable. A second response has been that no papers in the extant literature come to mind that use any instruments other than the two sorts described above, *ergo* the insight is valueless since it will not change empirical practice.

The appropriate response to the second point is that it supports Reiss's claim that economists require RP and T to justify their methods, since there appear to be no implicit or explicit assumptions within the discipline that explain this state of affairs. The more general point is: if a causal relationship underlies the endogeneity problem, then only instruments with certain kinds of causal relations to the endogenous variable can satisfy IV1-IV3. In response to the claim that the above implications of RP are 'obvious', we may challenge the sceptical econometrician to explain why a variable satisfying the formal requirements, but caused by the endogenous variable, cannot be a valid instrument. We suggest it is unlikely any answer will avoid assumptions about causal relationships, which currently seem to be implicit in economists' determination of the 'plausibility' of a given qualitative justification for instrument validity. If this is the case, the question of interest becomes what such assumptions might imply for other areas of the discipline were they to be made explicit. That is the question to which we now turn.

3 Multicollinearity and the interpretation of regression coefficients

To consider the possible import of RP, we focus on a specific example: the way in which economists deal with significant (empirical) correlations between *explanatory* variables, known as 'multicollinearity'. As noted by Angrist and Pischke (2009): "The importance of...[omitted variables bias]...[is] that if you claim an absence of omitted variables bias, then typically you're also saying that the regression you've got is the one you want. And the regression you want usually has a causal interpretation." (Angrist and Pischke 2009, 62). Multicollinearity is the alternative empirical scenario to the one addressed by IV methods, where the confounding variable is *observed* and hence can be included as a covariate in a multiple regression. Indeed, it should be clear from section 1 that empirical collinearity is simply a logical consequence of including a covariate that is genuinely necessary for identification. Consequently, consideration of this issue provides a natural extension of our arguments above.

Contrary to this seemingly obvious perspective, many textbooks treat collinearity as arising from *spurious* correlation and go as far as asserting that it is simply a *sample* (rather than 'population') problem to be resolved by more, and better, data.[25] A similar attitude is evident in Blanchard (1987)'s statement that: "Multicollinearity is God's will, not a problem with [ordinary least squares] or statistical

[25]No particular justification for this assertion is provided, though in principle it should be testable; standard statistical tests can be used to examine the likelihood of a given correlation being due to chance. Yet such commentary typically makes no mention of examining the presumption of spuriousness.

techniques in general" (Blanchard 1987, 449). However, he advocates the use of further *theoretical* assumptions to resolve the problem rather than additional data and our view is in line with that position. In particular, given Reichenbach's principle one *cannot* simply dismiss statistically significant correlations as happenstance; theoretical assumptions are required that preclude these correlations from representing causal relationships, or preclude their relevance for estimation of the parameters of interest, and those in turn require some foundation.

To the extent that multicollinearity has drawn any sustained attention within economics, the focus has been on perfect collinearity: for two variables this simply means a correlation between them equal to one; with multiple variables it means that one variable can be written as a linear combination of the others.[26] For a brief period, there was some concern about the effects of even lesser correlations on the validity of estimates from a standard least-squares regression - see Farrar and Glauber (1967) and Mansfield and Helms (1982) - but the modern consensus is that provided the collinearity is not perfect, or close to perfect, there is essentially no problem. The basis for this is a simple proof that, under the standard regression assumptions, correlation *per se* does not affect the desirable properties of the least-squares estimator.

Does this result change if the correlation is due to causation? Acceptance of Reichenbach's principle necessitates that question, and it may seem possible that

[26]In matrix representations the assumption of no perfect collinearity is clearly stated and known as 'the rank condition'.

the result could change. However, if we represent the causal relationships using linear equations, it is fairly straightforward to show an absence of 'bias' *per se* in the estimated coefficients. In short: *regardless* of causal relationships between explanatory variables, provided all confounding causes are included, the estimated coefficients remain unbiased where the assumption of linearity in relations holds.[27] This seems like a reassuring result for econometricians.

The result is somewhat misleading, however, since under RP the parameters estimated are *conceptually* different to those from scenarios where covariates are uncorrelated. To be specific, consider the causal systems represented in Figure 3. Figures 3.i and 3.ii illustrate two possible causal systems under $\mathbf{RP}$ if we have significant correlation between covariates.[28] Assume the econometrician is interested in the effect of X on Y, and conditions on C to avoid possible omitted variable bias. The arrow between C and X follows from empirical correlation between these variables, and assuming $\mathbf{RP}$. The total effect of X on Y in Figure 3.i is β^*, while the total effect of C is equal to its direct effect and indirect effect ($\beta^* \times \beta$).

Absent some basis for thinking C causes X - like temporal order for instance - collinearity could instead imply a system like Figure 3.ii. Comparing the two

[27]The relevant derivations are available from the author, but the result should be unsurprising given that the presence of causal relationships does not alter the statistical results.

[28]The case of a common cause is omitted since it will suffice for us to demonstrate that at least one system may exist that would require a reinterpretation of estimated coefficients.

systems reveals the problem for interpretation: in 3.i the direct effect of X is the same as its 'total effect' (equal to all direct and indirect effects), whereas in 3.ii there is a separate indirect effect that has been screened-off. Interestingly, in the inclusion of covariates to mitigate or avoid bias it is not uncommon for economists to justify inclusion by an ex post reduction in the magnitude of the estimated coefficient on the covariate of interest. While in Figure 3.i the reduction occurs because a confounding factor is correctly controlled for, Figure 3.ii shows that such a reduction could occur due to screening-off a portion of X's total effect.

The problem, then, is that economic method does not demand causal assumptions regarding omitted variables or correlated explanatory variables. This leaves open *all* empirical possibilities (within the limits of the original structural equation) and can therefore can lead to inconsistencies in empirical work. For example, researchers have often interpreted estimated parameters of explanatory regressors *symmetrically* at the same time as dismissing multicollinearity as unproblematic. Even where 'controls' are included because of hypothesised connections to a particular explanatory variable of interest, the coefficients on all variables are typically interpreted in the same way - which clearly makes no sense in causal systems like those illustrated. While it is true that all estimated coefficients will represent *direct effects* - often called 'partial effects' in the econometric literature - for many purposes (e.g. policy advice) it matters whether the direct effect is equivalent to the total effect or not. Either can be coherently referred to as 'the causal effect of X on Y', depending on what assumptions are made about causal intermediaries.

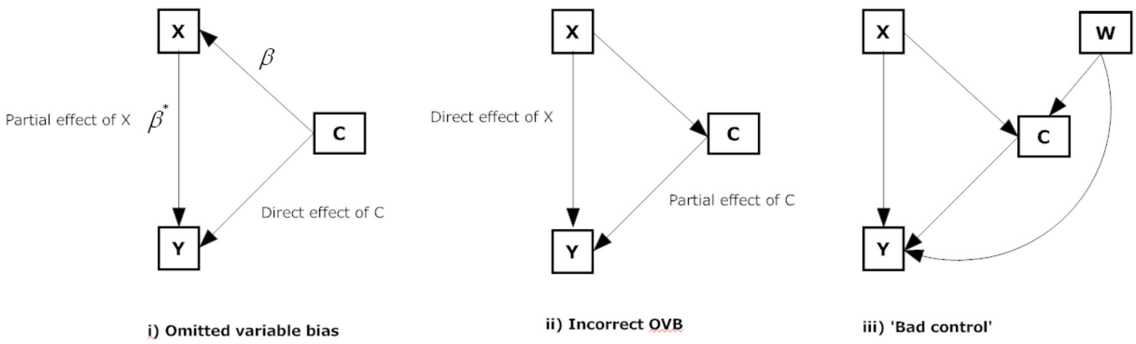

Figure 3: Implications of multicollinearity under Reichenbach's principle[29]

To avoid some of the conceptual mistakes discussed in section 2, it is important to reiterate that the above argument does not *necessarily* contradict the internal validity of the econometric method described. Strictly, our claim is that *given* $\mathbf{RP}$, economists' approach to multicollinearity is flawed. With that caveat in mind, it appears hard to construct a non-arbitrary formulation of the control-based method that does not suffer from the problem of interpretation identified above, while also allowing causal interpretations of regression estimates. As with the reluctance to commit to a causal representation of instruments, economists have taken the view that it is unneccessary to consider the implications of a causal aspect to collinearity. But they have not in any way *ruled-out* causal origins of such correlations. The nature of these assumptions appears to represent an inclination to terminate consideration of causal issues in a seemingly ad hoc manner, perhaps a vestigial trait of the causal agnosticism mentioned in earlier sections.

One notable exception to such agnosticism is the popular book by Angrist and Pischke (2009). Given the authors' explicit interest in causal issues, and intent to provide guidance on empirical practice, their work provides a good measure of the relevance of the causally-founded problems raised above for the rationale behind, and interpretation of, the use of covariates. Of particular relevance for our concerns is the authors' consideration of instances in which inclusion of covariates can generate problems instead of resolving them - an issue not addressed in many textbook accounts. The relevant part of their account - see Angrist and Pischke (2009, 59-68) - focuses on what they call the problem of 'bad control': "Bad controls are variables that are themselves outcome variables in the notional experiment at hand. That is, *bad controls might just as well be dependent variables too*" (my emphasis, Angrist and Pischke 2009, 64)

The scenario they envision is illustrated in Figure 3.iii. The concern is that if a causal intermediary shares an unobserved common cause with the dependent variable, conditioning on it yields a biased coefficient *even* where the variable of interest has been the subject of a randomised trial. We suggest this scenario is conceptually of second-order relative to the issues raised by contrasting figures 3.i and 3.ii. Interestingly, the source of Angrist and Pischke's concern is that a researcher might include an effect of the treatment to avoid omitted variables bias on the basis that $corr(C, X) \neq 0$. Under RP it is easy to refute such logic: if a researcher is interested in the total effect of an induced change in the variable of interest, they need only include C in a regression if it is believed that the correlation with X implies that C is a cause of X, or shares with it a common cause, *and* has an independent causal effect on Y. In short, one would have to believe that the randomized trial departed from the ideal in some way. Alternatively, if for some reason the researcher wanted to include C because they are specifically interested in the causal effect of X *excluding* the channel through C, then it must be acknowledged that while omitting C in the relevant regression does avoid confounding by another variable, it also means that the causal parameter of interest is not identified. Instead of this nuanced argument, Angrist and Pischke (2009)'s recommendation to researchers is simply to never condition on *any* variable temporally subsequent to treatment; this is a conclusion that is overly strong and not justified by the argument they present. Consequently, while they advocate "clear reasoning about causal channels" - primarily by identifying temporal order or making assumptions in this regard - their analysis fails to do this in a systematic fashion.

By comparison, it is significant that the analysis by Pearl (2009) does not suffer these weaknesses. First, that work clearly and explicitly addresses the issue of direct and indirect effects and their policy-relevance (see for instance Pearl 2009, 126-128). Second, it emphasises that a perfectly successful randomisation serves to sever the link between a variable and all its causes in the causal system that are not related to the experiment; indeed this assertion is fundamental to that work. Finally, and perhaps most importantly, it considers the full range of causal structures subsumed under economists' correlation conditions. This is as we would expect, given that Pearl assumes $\mathbf{RP}$. As a consequence, he comes to more nuanced conclusions:

- Conditioning on an effect (x_j) of x_i that is affected *by some other latent cause* results in a biased estimate of the *direct effect* of x_i (excluding that via x_j)

- "if we are careful never to adjust for *any consequence of treatment*...no bias will emerge in randomized trials"

(my emphases, Pearl 2009, 339-340)

The first point recognises that it is latent factors that cause the conditioning problem and is likely to occur if the interest is in the *direct* effect.[30] He does not make the error of claiming that *any* temporally subsequent variable would induce a bias, though even in that account this nuance can be lost - as illustrated by the second point which risks conflating issues relating to total effects and bias. The point there is that conditioning on a consequence of treatment could mean estimating a direct rather than a total effect if that consequence is simply a causal intermediary; it is conditioning on a consequence affected by a latent factor that leads to bias.

These issues, relating to correct interpretation of coefficients in the presence of relationships between explanatory variables, are made explicit in structural models in econometrics, and in what is known as 'path analysis', which at one point was popular in other social sciences such as sociology. The concern with those approaches, most particularly the former, has been that they typically require strong a priori assumptions about relationships between variables. However, the effort to move away from the constraints of the structural approach, while nevertheless addressing issues of causality, has perhaps meant the neglect of some important methodological issues. In the case of multicollinearity, the issue concerns

[30] Pearl's graphical representation of this issue - Pearl (Figure 11.5, 2009, 339) - is very similar to our Figure 3.iii above, except that it does not include a direct (unmediated) arrow from X to Y.

the correct interpretation of multiple regression coefficients given non-spurious collinearity among explanatory variables. Reichenbach's principle brings clarity to this problem, but its adoption would also imply a fundamental shift in how microeconometricians approach causal inference and that must also be true for any methods - such as those proposed by Pearl (2009) - that take RP as axiomatic.

Conclusion

Our primary concern in this paper has been to address claims by Reiss (2005, 2008) that Reichenbach's principle is necessary for econometricians' instrumental variable analysis, and by Pearl (2009) that causal graph methods premised on RP can serve to resolve a key intra-disciplinary conflict between structural econometricians and experimentalists. In section 2 we argued that the first claim is strictly false and while not addressing the second claim directly, in sections 2 and 3 we showed that RP has important implications beyond those aspects of econometrics explicitly related to the methodological dispute in question. Throughout, our argument has been that philosophical issues relating to causality may be important for practice in econometrics, but that it is necessary to appreciate the full rationale of what applied economists actually do. In particular: the distinction between structural and regression errors; the assumptions implicit in economists' structural models; the difference between various types of regression model misspecifications; and the distinction between macro- and micro-econometrics are all important.

In addition, we have suggested that Reichenbach's principle would have important implications for the interpretation by economists of their own discipline, since it would imply that all results premised on correlations or covariances are in fact causal statements of some sort. With this caveat in mind, we use the examples of instrumental variables analysis and economists' treatment of collinearity between explanatory variables to demonstrate that acceptance of Reichenbach's principle would provide coherent foundations for methods that may be flawed (at least in their interpretation). In the case of instrumental variables, RP provides a clear link between specification of the causal reasons for confounding and the causal role required for instruments to satisfy the statistical conditions for identification of causal parameters. In the converse case where the confounding factor is used as a control variable, RP allows us to demonstrate a problem with the symmetric interpretation of coefficients in multiple regressions where statistically significant correlation between covariates is present. While these arguments and examples are, on the one hand, somewhat more subtle than the counterexamples proposed by Reiss (2005, 2008) to the current logic of instrumental variables, we suggest they do support his claim that Reichenbach's principle is relevant for (micro)econometric methodology.

Whether microeconometricians will accepts these propositions is an entirely different matter. It may be that the discipline will continue to prefer a greater number of more specific, even arbitrary, assumptions to justify conclusions that could otherwise be reached by assuming Reichenbach's principle. Furthermore, the metaphysical status of the principle remains open. There are substantive reasons to question its generality, as noted by Arntzenius (1992, 2010), and its validity may well vary *within* the domains covered by the discipline of economics as a whole. Nevertheless, given the current state of microeconometric methodology as we have characterised it, it would appear that causal inference in this area *is* in need of either a principle akin to that proposed by Reichenbach, or an expansion of the ex ante assumptions economists typically make about causal structure.

References

Angrist, J., G. Imbens, and D. B. Rubin (1996). Identification of causal effects using instrumental variables. *Journal of the American Statistical Association 91*, 444–455.

Angrist, J. D. and A. B. Krueger (2001). Instrumental variables and the search for identification: from supply and demand to natural experiments. *Journal of Economic Perspectives 15*(4), 69–85.

Angrist, J. D. and J.-S. Pischke (2009). *Mostly harmless econometrics*. Princeton: Princeton University Press.

Arntzenius, F. (1992). The common cause principle. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association 2*, 227–237.

Arntzenius, F. (2010). Reichenbachś common cause principle. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2010 ed.).

Blanchard, O. J. (1987). Comment. *Journal of Business and Economic Statistics 5*(4), 449–451.

Boland, L. (2010). Cartwright on "economics". *Philosophy of the Social Sciences 40*(3), 530–538.

Cartwright, N. (1999). *The Dappled World: A study of the boundaries of science*. Cambridge: Cambridge University Press.

Cartwright, N. (2007). *Hunting causes and using them: Approaches in philosophy and economics*. Cambridge: Cambridge University Press.

Cartwright, N. (2010). What are randomised controlled trials good for? *Philosophical studies 147*, 59–70.

Deaton, A. (2008, October 9th). Instruments of development: Randomization in the tropics, and the search for the elusive keys to economic development. Keynes Lecture, British Academy.

Deaton, A. (2010). Instruments, randomization, and learning about development. *Journal of Economic Literature 48*(2), 424–455.

DiNardo, J. (2008). Natural experiments and quasi-natural experiments. In S. N. Durlauf and L. E. Blume (Eds.), *The New Palgrave Dictionary of Economics*. Basingstoke: Palgrave Macmillan.

Farrar, D. E. and R. R. Glauber (1967). Multicollinearity in regression analysis: The problem revisited. *Review of Economics and Statistics* (49), 92–107.

Fennell, D. (2007). Why functional form matters: revealing the structure in structural models in econometrics. *Philosophy of Science* (74), 1033–1045.

Glymour, C. (2004). The automation of discovery. *Daedalus* (Winter), 69–77.

Greene, W. H. (2003). *Econometric Analysis* (5th ed.). Pearson Education.

Hausman, D. M. and J. Woodward (1999). Independence, invariance and the causal Markov condition. *British Journal for the Philosophy of Science 50*(4), 521–583.

Heckman, J. J. (2000). Causal parameters and policy analysis in economics: A twentieth century retrospective. *Quarterly Journal of Economics 115*, 45–97.

Heckman, J. J. and E. J. Vytlacil (2007a). Chapter 70 econometric evaluation of social programs, part i: Causal models, structural models and econometric policy evaluation. Volume 6, Part B of *Handbook of Econometrics*, pp. 4779 – 4874. Elsevier.

Heckman, J. J. and E. J. Vytlacil (2007b). Chapter 71 econometric evaluation of social programs, part II: Using the marginal treatment effect to organize alternative econometric estimators to evaluate social programs, and to forecast their effects in new environments. Volume 6, Part B of *Handbook of Econometrics*, pp. 4875–5143. Elsevier.

Hendry, D. F. and H.-M. Krolzig (2005). The properties of automatic Gets modelling. *Economic Journal 115*, C32–C61.

Holland, P. W. (1986). Statistics and causal inference. *Journal of the American Statistical Association 81*(3), 945–960.

Hoover, K. (2001). *Causality in macroeconomics*. Cambridge: Cambridge University Press.

Hoover, K. (2003). Nonstationary time series, cointegration, and the principle of the common cause. *British Journal for the Philosophy of Science 54*, 527–551.

Hoover, K. (2009). Probability and structure in econometric models. In W. W. Clark Glymour and D. W. hl (Eds.), *Logic, methodology and philosophy of science: Proceedings of the Thirteenth International Congress*, pp. 497–513. College Publications.

Keane, M. (2005, 17-19 September). Structural vs. atheoretic approaches to econometrics. Keynote Address at the Duke Conference on Structural Models in Labor, Aging and Health.

Keane, M. P. (2010a). A structural perspective on the experimentalist school. *Journal of Economic Perspectives 24*(2), 47–58.

Keane, M. P. (2010b). Structural vs. atheoretic approaches to econometrics. *Journal of Econometrics 156*(1), 3–20.

Levitt, S. D. (1997). Using electoral cycles in police hiring to estimate the effect of police on crime. *American Economic Review 87*(4), 270–290.

Levitt, S. D. (2002). Using electoral cycles in police hiring to estimate the effect of police on crime: Reply. *American Economic Review 92*(4), 1244–1250.

Mansfield, E. R. and B. P. Helms (1982). Detecting multicollinearity. *The American Statistician 36*(3), 158–160.

Manski, C. (1991). Regression. *Journal of Economic Literature 29*(1), 34–50.

Manski, C. F. (2007). *Identification for prediction and decision*. Cambridge (MA): Harvard University Press.

McKim, V. R. and S. P. Turner (1997). *Causality in Crisis? Statistical methods in the search for causal knowledge in the social sciences*. Notre Dame (IN): University of Notre Dame Press.

Morgan, S. L. and C. Winship (2007). *Counterfactuals and causal inference*. Cambridge: Cambridge University Press.

Murray, M. P. (2006). Invalid instruments and coping with weak instruments. *Journal of Economic Perspectives 20*(4), 111–132.

Pearl, J. (2000/2009). *Causality: Models,reasoning and inference*. Cambridge: Cambridge University Press.

Reichenbach, H. (1956/1991). *The direction of time*. Berkeley: University of California Press.

Reiss, J. (2005). Causal instrumental variables and interventions. *Philosophy of Science 72*, 964–976.

Reiss, J. (2007). Time series, nonsense correlations and the principle of the common cause. In F. Russo and J. Williamson (Eds.), *Causality and probability in the sciences*, pp. 179–196. London: College Publications.

Reiss, J. (2008). *Error in Economics: The Methodology of Evidence-Based Economics*. London: Routledge.

Rosenzweig, M. R. and K. I. Wolpin (2000). Natural "natural experiments" in economics. *Journal of Economic Literature 38*(4), 827–874.

Sober, E. (2001). Venetian sea levels, British bread prices, and the principle of the common cause. *British Journal for the Philosophy of Science 52*, 331–346.

Spirtes, P., C. N. Glymour, and R. Scheines (1993/2000). *Causation, prediction and search*. Cambridge(MA): MIT Press.

Steel, D. (2003). Making time stand still: A response to Sober's counter-example to the principle of the common cause. *British Journal for the Philosophy of Science 54*(2), 309–317.

Woodward, J. (1988). Understanding regression. In *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, pp. 255–269. Chicago: University of Chicago Press.

Woodward, J. (1995). Causation and explanation in econometrics. In D. Little (Ed.), *On the Reliability of Economic Models: Essays in the Philosophy of Economics*, pp. 9–62. Dordrecht: Kluwer.

Woodward, J. (2003). *Making things happen*. Oxford: Oxford University Press.

Wooldridge, J. M. (2002a). *Econometric analysis of cross section and panel data*. Cambridge(MA): MIT Press.

Wooldridge, J. M. (2002b). *Introductory econometrics: A modern approach*. Cambridge(MA): South-Western College Publishing.

southern africa labour and development research unit

The Southern Africa Labour and Development Research Unit (SALDRU) conducts research directed at improving the well-being of South Africa's poor. It was established in 1975. Over the next two decades the unit's research played a central role in documenting the human costs of apartheid. Key projects from this period included the Farm Labour Conference (1976), the Economics of Health Care Conference (1978), and the Second Carnegie Enquiry into Poverty and Development in South Africa (1983-86). At the urging of the African National Congress, from 1992-1994 SALDRU and the World Bank coordinated the Project for Statistics on Living Standards and Development (PSLSD). This project provide baseline data for the implementation of post-apartheid socio-economic policies through South Africa's first non-racial national sample survey.

In the post-apartheid period, SALDRU has continued to gather data and conduct research directed at informing and assessing anti-poverty policy. In line with its historical contribution, SALDRU's researchers continue to conduct research detailing changing patterns of well-being in South Africa and assessing the impact of government policy on the poor. Current research work falls into the following research themes: post-apartheid poverty; employment and migration dynamics; family support structures in an era of rapid social change; public works and public infrastructure programmes, financial strategies of the poor; common property resources and the poor. Key survey projects include the Langeberg Integrated Family Survey (1999), the Khayelitsha/Mitchell's Plain Survey (2000), the ongoing Cape Area Panel Study (2001-) and the Financial Diaries Project.

www.saldru.uct.ac.za
Level 3, School of Economics Building, Middle Campus, University of Cape Town
Private Bag, Rondebosch 7701, Cape Town, South Africa
Tel: +27 (0)21 650 5696
Fax: +27 (0) 21 650 5797
Web: www.saldru.uct.ac.za