Synthese (2022) 200:185
https://doi.org/10.1007/s11229-022-03684-0

ORIGINAL RESEARCH

Econometric methods and Reichenbach's principle

Seán Mfundza Muller[1]

Received: 27 November 2020 / Accepted: 31 March 2022 / Published online: 21 April 2022

Abstract
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 (Philos Sci 72:964–976, 2005) has argued that the principle is necessary for instrumental variables methods in econometrics, and Pearl (In Causality: Models, reasoning and inference, Cambridge: Cambridge University Press, 2000/2009) builds a framework using it that he proposes as a means of resolving an important methodological dispute among econometricians. Through analysis of instrumental variables methods and the issue of multicollinearity, we aim to show 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 · Common cause principle · Instrumental variables · Econometrics · Causal inference

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, p. 157). Determined by frequency of occurrence, one might represent an improbable coincidence in probability terms as: $P(A \wedge B) > P(A)P(B)$.[1] 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,

[1] 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.

T.C.: All Things Reichenbach Lead Guest Editor: Dr. Erik Curiel and Flavia Padovan.

✉ Seán Mfundza Muller
seanm@uj.ac.za

1 Johannesburg Institute for Advanced Study (JIAS), University of Johannesburg, 1 Tolip Street, Westdene, Johannesburg 2092, South Africa

p. 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’ 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 Reichenbach’s Principle, 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, Reichenbach’s Principle 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 the suggestion that correlation may arise from a common cause. The difference is in emphasis: Reichenbach’s Principle suggests causal inference may proceed from correlations, while the popular mantra emphasises caution in doing this. This critical distinction is reflected in method. For instance, one might contrast the use by Glymour & Scheines (2000) of Reichenbach’s Principle 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.[2]

The paper examines the implications of Reichenbach’s Principle for econometrics through two specific contributions. The first critically examines the prior claim by Reiss (2005; 2008) that Reichenbach’s Principle is necessary for, and possibly implicit in, economists’ instrumental variables method. Unpicking the errors in that account enables a more nuanced assessment of the significance of Reichenbach’s Principle in that context. At base, Reiss’s analysis is premised on a misunderstanding of economists’ stated rationale for instrumental variable solutions to the identification problem. Nevertheless, we argue that the Principle does raise concerns about economists’ choice and justification of instruments when these are properly understood. The second contribution considers the related issue of ‘multicollinearity’: statistical correlation between explanatory variables. In particular, it demonstrates how acceptance of Reichenbach’s Principle has important implications for the interpretation of regression coefficients and thereby, if correct, renders economists’ traditional approach to multicollinearity untenable. 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 substantive understanding of the implications of Reichenbach’s Principle for econometrics than has been the case to date.

The modern literature on microeconometrics–the development and application of statistical methods for empirical analysis of microeconomic issues–is primarily con-

[2] The Spirtes et al. (2000) approach has also been employed in the analysis of time-series data.

cerned with empirical identification of plausibly unconfounded effects of one variable of interest on another, sometimes referred to as 'the identification problem'. A confounding factor ('confounder') is one that is associated with the variables of interest in such a way that it may lead to incorrect inferences about the presence, or extent, of a causal relationship between those variables if not adequately accounted for.[3] One approach is to 'control' for variation in possible confounding factors in empirical analysis when estimating causal relationships, by including possible confounders (or proxies thereof) in the analysis. 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, pp. 3–4). One solution, increasingly presented using the counterfactual-based 'Rubin Causal Model'–see Angrist et al. (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 & 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.[4] A variable representing the source of such variation 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 & Pischke (2009) frame reluctance to do this as characteristic of a statistician rather than an econometrician. This appetite for causal claims has, however, not been accompanied by engagement with issues identified by philosophers such as manipulation (Woodward 2003) and interventionist (Pearl 2009) accounts of causality. One notable lacuna for philosophers of science is the wholesale omission of Reichenbach's Principle. The relevance of the Principle for econometrics has received some philosophical attention in relation to a counterexample proposed by Sober (2001); see for instance Hoover (2001; 2003; 2009), Steel (2003), and Reiss (2007). As Sober'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.[5] In contrast to the relative neglect of philosophical isues by economists, work by Spirtes et al. (2000) and Pearl (2009) on the use of causal graphs for algorithmic identification of causal relationships is 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).

[3] As Pearl (2009, pp. 194–195) notes, there are some nuanced distinctions to be made between confounders and confounding.

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

[5] 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').

Our focus here will be on microeconometric examples and the use of purportedly quasi-random variation to identify causal effects. To start with, we consider Reiss's (2005; 2008) argument that Reichenbach's Principle 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. Having demonstrated the error in Reiss's approach, we will argue that Reichenbach's Principle does nevertheless provide important methodological insights into instrumental variable methods.

Our second contribution is to show how adopting Reichenbach's Principle has implications for empirical practice in economcis beyond those considered by Reiss, or indeed the proponents of causal graph approaches. In particular, we consider economists' approach to correlations among explanatory variables in regressions, known as 'multicollinearity'. As we explain below, multicollinearity is the obverse of the instrumental variables case, and therefore complements our first contribution.

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.[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, which is intended to correspond to the underlying relationship. 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 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'. Empirical work in the social sciences, including eco-

[6] See Woodward (1988) for what one might call a 'traditional' overview of the formalities of regression methods directed at philosophers. Our discussion relies more on the presentations by Manski (1991) and Wooldridge (2002a). A valuable additional reference is Morgan & Winship's (2007) book, which provides an overview of regression and graph-based methods within a detailed discussion of causal inference based on counterfactuals.

[7] Linearity in the parameters allows $E(y|x)$ to be a non-linear function of the explanatory variables. E.g. We could have $E(y|\mathbf{x}) = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_2^2$. The linearity assumption is convenient though not essential.

nomics, remains dominated by use of such regressions and for ease of exposition we will focus on these.

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 besides the researcher's interest. However, as noted, modern econometric analysis is interested in more than mere description of associations between variables. Consequently, textbook analyses typically begin with a *structural* equation of the dependent variable of interest (y) that may be explicitly or implicitly causal.[8] 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 by assumption; 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.

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, p. 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, following the logic of the standard textbook approach, we can write the structural equation:

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

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': it can be represented in terms of population statistics. [9]

In the usual case researchers do not have data on the entire population but rather a sample, which may be randomly drawn, of N observations. When economists refer to 'running a regression' they are typically referring to the final process of estimating parameters of an equation using sample analogues of the population statistics mentioned above. Whether that estimates causal parameters depends on how accurately the estimated equation represents the true structure. In this framework, the identification problem concerns obtaining an empirical estimate of the parameter of interest

[8] 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).

[9] 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:

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

This is obtained by rewriting (3) as $y = \mathbf{x}\boldsymbol{\beta} + u$, premultiplying both sides by the transposed vector $\mathbf{x}^T$, applying the expectations operator to both sides and solving for $\boldsymbol{\beta}$ (noting that $E[\mathbf{x}^T\mathbf{u}] = 0$).

that corresponds to the 'true' value of that parameter in the data generating process-represented by (1).

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 ($\hat{\beta}$) relative to the true parameter in (1): $E(\hat{\beta}) \neq \beta$. In practice economists do not know what the true parameters are, but textbook empirical methods are typically based on thought experiments of this kind.[10] In what follows we will use one important form of endogeneity known as 'omitted variables bias' to illustrate the instrumental variables method. The reasons for this are two-fold. First, concern with omitted variables bias is arguably the primary reason underlying economists' use of instrumental variables methods. Second, it is the logic of that approach that Reiss (2005; 2008) seeks to critique using Reichenbach's Principle.

As noted by Woodward (1988, p. 259), a key decision in specifying an estimable regression equation is determining which variables should be included. The traditional approach among applied economists has viewed mistaken inclusion of a variable not in the true structural equation as less problematic than excluding a relevant variable. As discussed later in the paper, that position is questionable once proper consideration is given to the range of potential causal relationships among explanatory variables. Exclusion of a relevant variable 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 (4) 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 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 $\mathbf{x}$.

$$E(y|\mathbf{x}, q) = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \cdots \beta_k x_k + \gamma q \tag{4}$$

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.[11] It is this method that is closest to to the philosophical assumption of modularity and associated definitions of causation through manipulation advocated by authors such as Hausman & Woodward (1999) and Holland (1986), and criticised by others such as Cartwright (2002) and Hoover (2011). For simplicity, however, our analysis will focus on the so-called 'instrumental variables' (henceforth, IV) solution to the omitted variable problem using *observational* data.

10 See Qin (2018) for a recent critique of that approach.

11 There is some disagreement about how this severance of relationships should be represented; see Pearl (2009, pp. 376–377).

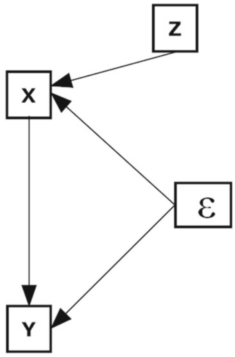

Fig. 1 The standard instrumental variables scenario

The *theoretical* IV solution stated in econometric textbooks 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 explaining 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.[12]

To illustrate, consider the problem of attempting to estimate the effect of additional years of schooling (X) on earnings (Y) in the labour market. The estimated parameter from a naive regression of earnings on schooling will likely be biased, upwards, by a range of omitted variables (ϵ) that affect both earnings and years of schooling. These include factors like family wealth, parental education, individual ability, and geographical location. One approach to addressing that would be to include as many of these variables, or proxies thereof, as possible as explanatory variables in the regression. Such a strategy faces not only data limitations but also the problem of whether researchers can be assumed to know all the relevant factors ex ante. The instrumental variables approach proposes a seemingly much simpler solution: find *one* variable that satisfies the above three conditions (IV1-IV3). One influential example utilised the month of birth of individuals in the United States (Angrist & Krueger 1991): the combination of legal stipulations that determined start age and the minimum age for leaving schooling meant that month of birth (Z) determined the amount of schooling completed by the time a student was presented with the option of leaving school.[13] The argument is that such variation in schooling is independent of possible confounding

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

[13] See Angrist & Pischke (2009, pp. 117–120) for a summary of this example in the context of an explanation of the instrumental variables approach.

factors, since month of birth is independent of them. However, the latter claim cannot be empirically tested and is therefore, in effect, an assumption premised on arguments that are typically causal in nature.

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.[14] And to estimate it one could substitute the sample analogues as before. A key point is that while *IV1* and *IV3* can be tested empirically, *IV2* cannot because it concerns the unobservable term u.[15] 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, p. 114).[16]

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.[17] 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. Following on work of this sort by other authors (Cartwright 1999; Hoover 2001; Reiss 2008), in this section we consider the implications of Reichenbach's Principle for the use of instrumental variables in econometrics. In particular, we critically examine the claim by Reiss (2005) that instrumental variables analysis requires Reichenbach's Principle, and in doing so we also seek to clarify a few potential misunderstandings regarding econometric methods.

Reiss's (2005) basic argument–also in Reiss (2008, pp. 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

14 The relevant expression is:

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

Where $\mathbf{z}$ is the $\mathbf{x}$ vector *including* z and *excluding* x_j.

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

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

17 Similar approaches are also utilised in macroeconometrics, however the nature of much macroeconomic data-in the form of time-series at the country level-is such that there are additional complications which are outside the scope of the present paper.

Fig. 2 Reiss's (2005) counterexample

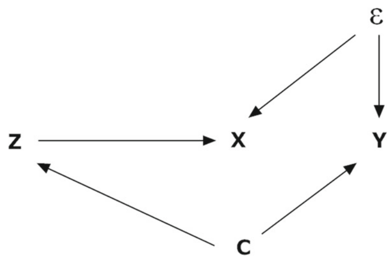

based on instrumental variables and divides these into 'stages' of analysis. Below we explain the flaws in that proposal, but first we state these stages for reference:

Stage 1 Assume Reichenbach's Principle (**RP**), causal transitivity (**T**) and 'functional correctness' (**FC**);[18]

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

Stage 3 Assume that the instrument (Z) is a 'causal instrumental variable': Z is a cause of X ('CIV1'); Z either is not a cause of Y or is only a cause of Y through X ('CIV2'); Z and Y either do not have common causes, or any common cause satisfies *CIV2* ('CIV3') (Reiss 2007, p. 973)

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

The primary problem with Reiss's analysis is the neglect of additional implicit, or definitional, assumptions in addition to the explicitly stated IV1 and IV2.[19] Two assumptions implicit in (4) are: That q and each x_i, $i = 1 \dots k$, have independent explanatory power for y in the sense that they explain variation in y independently of other factors; and, that *no* other observed or unobserved variable, including any associated with any of the *specified* explanatory variables, will have independent explanatory power for y.[20] Neglect of these assumptions means that Reiss's counterexample is structurally inequivalent to the situation econometricians concerned with omitted variable bias are seeking to address.

This oversight manifests in the counterexample in Fig. 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

18 Causal transitivity means that, "For any three variables A, B, and C, if A causes B and B causes C, then A causes C" (Reiss 2005, p. 969). And Reiss 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". In essence, this can be understood as meaning that an equation correctly represents the relationship between the magnitudes of variables in the equation, but those relations need not be causal.

19 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.

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

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, Fig. 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. Contra Reiss: it is not a failure of the method per se.[21]

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.[22] This oversight is related to misunderstanding the properties of error terms in (true) structural and estimated regression equations.[23] Indeed, as Reiss notes, if "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." (Reiss 2005, p. 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, p. 259).[24] It is, correspondingly, important to remain alive to the subtle distinctions between (micro)economists' methods of causal discovery and algorithmic approaches to 'hunting causes' (Spirtes et al., 2000; Pearl, 2009). 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

[21] Hoover (2007) draws attention to an analogous problem in Reiss's reasoning about 'collider' variables.

[22] 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, p. 517) and Pearl (Figure 7.8 (d), 2009, p. 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.

[23] There are two additional problems with Reiss's argument. First, his proposed stage 2 assumption is in fact the *standard* way of interpreting the error term in a *hypothesised* structural equation. Econometric textbooks often make this interpretation explicit–see for instance Greene (2003, p. 8)-and it is recognised by Woodward (1988, p. 261). Pearl 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, p. 162-163). Second, Reiss fails to note that one can describe the IV logic without reference to causality per se. Contra to Reiss's criticism that "textbooks contain 'recipes' for econometric inference that give the impression that econometrics can proceed without causal background assumptions" (Reiss 2005, p. 966), 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.

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

assumptions–like Reichenbach's Principle–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 without substantive assumptions about causal structure. As relates to common cause assumptions like Reichenbach's Principle, the analysis of Arntzenius (1992) would appear to indirectly support Cartwright's stance by virtue of his elaboration of counterexamples. Except that the counterexamples to Reichenbach's Principle are quite domain-specific, and his conclusion supports the view that there exist systems of interest under which it is a valid assumption. Furthermore, Hoover (2007) provides a detailed critique of purported counterexamples presented by Reiss (2007).

The difference in emphasis of the two approaches is directly connected to the necessity of Reichenbach's Principle. What Reiss has done is construct a system where the two basic correlative relationships in instrumental variable analysis are satisfied. The axioms **RP**, **T** and **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 have tended to be uninterested in algorithmic approaches, as can been seen in the very limited use of such methods in empirical work and 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 & Krolzig (2005).[25] While Spirtes et al. (2000) and Pearl (2009) place the burden of causal structure on Reichenbach's Principle 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 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. Recent work by Keane (2010) against the possibility of 'atheoretical' econometrics and Heckman & Vytlacil (2007) contributes to a long-standing literature (Liu, 1960; Sims, 1980) that has challenged such propositions; 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. We concur with Reiss's conclusion, but from a different premise. The basis of Reiss's argument is that the theoretical basis for the IV method already contains within it causal assumptions, As already argued above, this is wrong and it bears elaborating upon before moving

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

to agree with the proposition that causal content be explicitly incorporated into the IV approach.

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:

$$x_j = \gamma_0 + \boldsymbol{\gamma}\mathbf{x}_{\neg\mathbf{j}} + \theta z + \varphi \tag{5}$$

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 (5). 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 (5). To the contrary, as Wooldridge emphasises: "there is nothing necessarily structural about [the] equation" (Wooldridge 2002a, p. 84).

However, 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.[26] Absent *any* commitment by economists to the causal character of relationships between the endogenous, confounding and instrumental variables, we might ask what difference assuming Reichenbach's Principle would make. Reichenbach's Principle 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. That seems to be the case for most illustrative examples in economics, such as individual ability being a common cause of education and earnings, or neighbourhood wealth being a common cause of the number of police officers and the crime rate. Then causal transitivity rules-out the instrument being caused by x, since that would imply correlation between z and the omitted factor, violating IV2.

Therefore a useful point that emerges from combining Reichenbach's Principle with a correct understanding of econometric method, is that if economists fail to address the (potentially) causal origins of endogeneity, they cannot convincingly make a causal case for instrument validity. And this is consistent with earlier arguments in the causal graph literature (Spirtes et al., 2000). For instance, if economists accept **RP** and **T** then, as a methodological point, where the endogenous explanatory variable shares the omitted variable as a common cause with y valid instruments *must*:

[26] 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.

1. Be causes of the endogenous variable, or;[27]
2. Share a common cause with the endogenous variable.

As an example, consider Murray's (2006) 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. Thus Levitt begins with a causal story as to why a naive regression of crime on police numbers would likely produce biased estimates, then he presents further causal stories to support the claim that his proposed instruments are valid and will eliminate the endogeneity bias.

That most instruments and the 'stories' told to support them in the literature follow this logic–from endogeneity resulting from an omitted common cause to causal instruments–supports the assertion that the causal nature of IVs ought to be made explicit, and that this in turn would be a worthwhile addition to the standard textbook account. That suggestion is likely to be resisted by economists for a number of reasons, two of which can be illustrated by actual comments from applied microeconometricians at a seminar presentation of the basic argument. First, that the suggestion is akin to proposing a line in every textbook saying, "Don't be stupid". In other words: it is simply *obvious* that a valid instrument cannot be caused by the endogenous variable. The second comment was that no papers in the extant literature come to mind that use any instruments other than the two sorts described, *ergo* the insight is valueless since it will not change empirical practice.

The appropriate response to the second point is that it *supports* the view 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 first claim that the above implications of Reichenbach's Principle are 'obvious', we may challenge the econometrician who is skeptical about the merits of making causal assumptions explicit to explain why a variable satisfying the formal requirements, but caused by the endogenous variable, cannot be a valid instrument. 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. That being so, there is a strong case for making such assumptions explicit and Reichenbach's Principle would be an obvious basis on which to do so. A question that follows naturally from that is: what might such assumptions imply for other areas of the discipline if they were to be made explicit? It is to that question we now turn.

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

3 Multicollinearity and the interpretation of regression coefficients

To consider the possible import of Reichenbach's Principle, 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 & 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 & Pischke, 2009, p. 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.[28] A similar attitude is evident in Blanchard's (1987) statement that: "Multicollinearity is God's will, not a problem with [ordinary least squares] or statistical techniques in general" (Blanchard 1987, p. 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.[29] 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 & Glauber (1967) and Mansfield & 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 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 esti-

[28] 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.

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

mated 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.[30] This seems like a reassuring result for econometricians.

The result is somewhat misleading, however, since under Reichenbach's Principle the parameters estimated may be *conceptually* different to those from scenarios where covariates are uncorrelated. *If* correlation implies some form of causal relationship between variables, then the inclusion of additional explanatory variables correlated with the variable of interest necessitates a change in interpretation of the estimated parameter on that variable. But how the interpretation changes will depend on the precise nature and/or direction of the causal relationship.

To be specific, consider the causal systems represented in Fig. 3. Figs. 3.i and 3.ii illustrate two possible causal systems under **RP** if we have significant correlation between covariates.[31] 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 **RP**. The total effect of X on Y in Fig. 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 Fig. 3.ii. Comparing the two 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 partialled-out. 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 Fig. 3.i the reduction occurs because a confounding factor is correctly controlled for, Fig. 3.ii shows that such a reduction could occur due to partialling-out a portion of X's total effect.

The problem, then, is that recently-dominant empirical methods in economics do not demand causal assumptions regarding omitted variables or correlated explanatory variables. That 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 (and the associated risk of omitted variable bias), 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

[30] 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.

[31] 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.

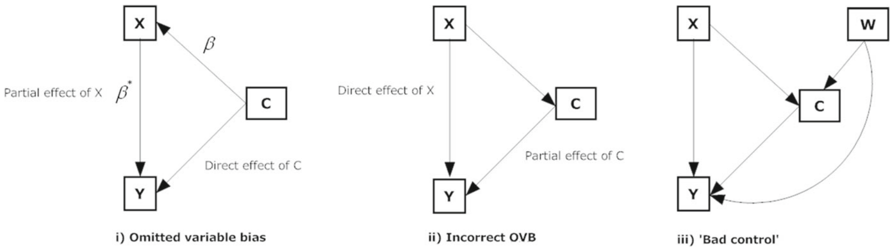

Fig. 3 Implications of multicollinearity under Reichenbach's Principle

to as 'the causal effect of X on Y', depending on what assumptions are made about causal intermediaries.[32]

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* **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. In other words, economists appeal to a causal notion that they have not coherently articulated, but articulation thereof would throw their preferred approach into question. As with the reluctance to commit to a causal representation of instruments, economists have typically taken the view that it is unnecessary to consider the implications of a causal aspect to collinearity.[33] 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 and surveyed in more detail by Hoover (2004).

One notable exception to such agnosticism is the popular book by Angrist & 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 (Angrist & Pischke, 2009, pp. 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 & Pischke, 2009, p. 64)

[32] Implicit in this graph is that, following Pearl (2009), all variables affecting *more than one* other variable in the system are shown.

[33] A partial exception is reflected in one historical practice of conducting and reporting results from multiple regressions. The practice was to report regression results in such a way as to show how the incremental inclusion of every additional explanatory variable ('control') altered the estimated parameter on the explanatory variable of primary interest.

The scenario they envision is illustrated in Fig. 3.iii. The concern is that if a causal intermediary shares an unobserved common cause (W) 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 Figs. 3.i and 3.ii, but it is nevertheless important. What is of particular interest for our purposes is that 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 $corr(C, X) \neq 0$. Yet by employing Reichenbach's Principle 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 & Pischke's (2009, p. 68) primary recommendation is that researchers not 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), which assumes **RP**, 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, pp. 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. As a consequence, Pearl 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, pp. 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.[34] 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

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

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; whereas 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. One response is represented by the graphical causal search approaches developed by Spirtes et al. (2000), which promises to deliver the causal structures considered in path analysis with weaker a priori assumptions. In microeconomics, the dominant response has been to move away from structural considerations entirely and utilise instead 'design-based' methods (Angrist & Pischke, 2010) that exploit properties of the data (including the use of experiments) to get at causal effects with minimal ex ante assumptions. However, these latter efforts to move away from the constraints of the structural approach have been associated with the neglect of some important methodological issues. In the case of multicollinearity, the issue concerns 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 Reichenbach's Principle as axiomatic.

4 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 Reichenbach's Principle 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 Reichenbach's Principle 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 otherwise be flawed (at least in their interpretation). In the case of instrumental variables, Reichenbach's Principle 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, Reichenbach's Principle 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 demonstrate 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.

Acknowledgements I am grateful to two anonymous reviewers for detailed comments and suggestions that led to improvements in this manuscript, and to Martin Wittenberg for helpful comments on the very first draft of this work. Participants at a SALDRU seminar in the School of Economics at the University of Cape Town and other anonymous reviewers provided comments that informed the first working paper version (Muller 2012). The usual disclaimers apply.

References

Angrist, J. D., & Krueger, A. B. (1991). Does compulsory schooling attendance affect schooling and earnings? *Quarterly Journal of Economics, 106*, 976–1014.

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

Angrist, J. D., & Krueger, A. B. (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., & Pischke, J.-S. (2009). *Mostly harmless econometrics*. Princeton: Princeton University Press.

Angrist, J. D., & Pischke, J.-S. (2010). The credibility revolution in empirical economics: How better research design is taking the con out of econometrics. *Journal of Economic Perspectives, 24*(2), 3–30.

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). Reichenbach's 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.

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

Cartwright, N. (2002). Against Modularity, the Causal Markov Condition, and Any Link Between the Two: Comments on Hausman and Woodward. *The British Journal for the Philosophy of Science, 53*(3), 411–453.

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

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

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

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

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

Heckman, J. J., & Vytlacil, E. J. (2007). 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., & Krolzig, H.-M. (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. (2004). Lost causes. *Journal of the History of Economic Thought, 26*(2), 149–164.

Hoover, K. (2007). Probability and structure in econometric models. Paper presented at the 13th International Congress of Logic, Methodology, and the Philosophy of Science at Tsinghua University, Beijing, China, 8–16 August 2007.

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.

Hoover, K. (2011). Counterfactuals and causal structure. In P. Illari, F. Russo, and J. Williamson (Eds.), *Causality in the Sciences*, pp. 338–360. Oxford University Press.

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

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.

Liu, T.-C. (1960). Underidentification, structural estimation, and forecasting. *Econometrica, 28*(4), 855.

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

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

McKim, V. R., & Turner, S. P. (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., & Winship, C. (2007). *Counterfactuals and causal inference*. Cambridge: Cambridge University Press.

Muller, S. M. (2012). Econometric methods and Reichenbach's principle. SALDRU Working Paper 85.

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.

Qin, D. (2018). Let's take the bias out of econometrics. *Journal of Economic Methodology, 26*(2), 81–98.

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 & 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., & Wolpin, K. I. (2000). Natural "natural experiments" in economics. *Journal of Economic Literature, 38*(4), 827–874.
Sims, C. A. (1980). Macroeconomics and Reality. *Econometrica, 48*(1), 1.
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., Glymour, C. N., Scheines, R. (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. (2002). *Econometric analysis of cross section and panel data*. Cambridge, MA: MIT Press.
Wooldridge, J. M. (2002). *Introductory econometrics: A modern approach*. Cambridge: South-Western College Publishing.