

Hoover, Kevin D.

Working Paper

The Ontological Status of Shocks and Trends in Macroeconomics

CHOPE Working Paper, No. 2013-15

Provided in Cooperation with:

Center for the History of Political Economy at Duke University

Suggested Citation: Hoover, Kevin D. (2013) : The Ontological Status of Shocks and Trends in Macroeconomics, CHOPE Working Paper, No. 2013-15, Duke University, Center for the History of Political Economy (CHOPE), Durham, NC

This Version is available at:

<https://hdl.handle.net/10419/149706>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

THE ONTOLOGICAL STATUS OF SHOCKS AND TRENDS IN
MACROECONOMICS

BY KEVIN D. HOOVER

CHOPE Working Paper No. 2013-15

September 12, 2013



CENTER FOR THE
HISTORY OF POLITICAL ECONOMY
AT DUKE UNIVERSITY

The Ontological Status of Shocks and Trends in Macroeconomics*

Kevin D. Hoover

Department of Economics
and
Department of Philosophy
Duke University
Box 90097
Durham, NC 27708-0097

Tel. (919) 660-1876
E-mail kd.hoover@duke.edu

Revised, 12 September 2013

*Prepared for the conference on Ontology & Methodology at Virginia Tech 4-5 May 2013.

Abstract

of

Kevin D. Hoover's

The Ontological Status of Shocks and Trends in Macroeconomics

Modern empirical macroeconomic models, known as structural autoregressions (SVARs) are dynamic models that typically claim to represent a causal order among contemporaneously valued variables and to merely represent non-structural (reduced-form) co-occurrence between lagged variables and contemporaneous variables. The strategy is held to meet the minimal requirements for identifying the residual errors in particular equations in the model with independent, though otherwise not directly observable, exogenous causes (“shocks”) that ultimately account for change in the model. In nonstationary models, such shocks accumulate so that variables have discernible trends. Econometricians have conceived of variables that trend in sympathy with each other (so-called “cointegrated variables”) as sharing one or more of these unobserved trends as a common cause. It is possible for estimates of the values of both the otherwise unobservable individual shocks and the otherwise unobservable common trends to be backed-out of cointegrated systems of equations. The issue addressed in this paper is whether and in what circumstances these values can be regarded as observations of real entities rather than merely artifacts of the representation of variables in the model. The issue is related, on the one hand, to practical methodological problems in the use of SVARs for policy analysis – e.g., does it make sense to estimate of shocks or trends in one model and then use them as measures of variables in a conceptually distinct model? The issue is also related to debates in the philosophical analysis of causation – particularly, whether we are entitled, as assumed by the developers of Bayes-net approaches, to rely on the causal Markov condition (a generalization of Reichenbach’s common-cause condition) or whether cointegration generates a practical example of Nancy Cartwright’s “byproducts” objection to the causal Markov condition.

Keywords: shocks, trends, causation, latent variables, Bayes-nets, causal Markov condition, principle of the common cause, scientific realism, counterfactual policy analysis, structural vector-autoregression

JEL Codes: B4, C18, C32

The Ontological Status of Shocks and Trends in Macroeconomics

1. Shocks and Trends

In a preamble to a criticism of the use of the causal Markov condition in causal modeling, Nancy Cartwright tells us:

Metaphysics and methodology should go hand in hand. Metaphysics tells what something is: methodology, how to find out about it. Our methods must be justified by showing that they are indeed a good way to find out about the thing under study, given what it is. Conversely, if our metaphysical account does not tie in with our best methods for finding out, we should be suspicious of our metaphysics. [Cartwright 2007, p. 132]

While I do not share her views of the causal Markov condition, the idea that metaphysics and methodology must go hand in hand seems right. My aim here will be to explore the interaction of the metaphysics and methodology of macroeconometrics in a specific case.

Economics is a data-rich, and economic data display complex interrelationships both contemporaneously among the values of different variables and intertemporally between values of the same or different variables at different times. These relationships may be captured mathematically in a variety of representations that equally “save the phenomena” (van Fraassen 1980, ch. 3). Yet it is widely agreed that economics – especially since it frequently aspires to be a source of normative guidance or policy advice – cannot be indifferent among these phenomenologically equivalent representations but should privilege representations that support counterfactual analysis and provide a map of effective interventions in the economy.

Consequently, economists are committed implicitly to an ontology of causes. Indeed, some economists maintain that the distinction between *econometrics* and *statistics* applied to economic data is the distinction between methods suitable for identifying and measuring *causes* as opposed to methods suitable only for identifying and measuring *associations* (Heckman 2000, p. 45; Hoover 2004).

Econometric models that aim to capture causal relations in support of counterfactual analysis are referred to as *structural* and center on a decomposition into a deterministic structure, which represents the causal connections among variables and unobservable stochastic terms that are the entry point for probabilistic considerations. At least implicitly, economists are concerned with systems of such equations, conceived of in a way that maps fairly cleanly onto to so-called Bayes-net or graph-theoretic analyses of causation (Spirtes, Glymour, Scheines 2000; Pearl 2000).¹

Epistemological issues complicate economists’ understanding of causal ontology. The fundamental vision of many economists can be traced to Walras’s (1874[1954] formalization of general equilibrium – i.e., of the systemic interdependence of economic agents such that directly or indirectly every economic quantity depends on every other economic quantity. The general-equilibrium vision inspires the Walrasian fantasy of a complete agent-by-agent representation of the economy, not dissimilar to Laplace’s fantasy of the evolution of the physical universe as the unfolding of a set of initial conditions through a complete set of deterministic laws (Hoover 2001, p. 102).

There is a tension between this fantasy and the need to keep actual economic analysis manageable. A practicable structural equation may relate an effect to its causes and a residual.

For example,

$$(1) \quad z = \alpha x + \beta y + \varepsilon,$$

¹ Most discussions of Bayes nets focus on directed acyclical graphs, in which causation is strictly one-way. But this is not essential, and its characteristic of structural econometrics – going back to its foundations in Haavelmo’s (1943, 1944) seminal work – to be concerned with simultaneous (mutual) or cyclical causal relations.

where x , y , and z are variables, α and β are parameters, and ε is an unobservable residual; causes are written on the right-hand side of the equal sign; effects on the left.² The residual ε is frequently referred to as an *error* term. And, indeed, it may represent at least two types of error. It may, for instance, reflect measurement error. It may also reflect specification error. For example, the functional form of (1) is linear. But if this is only an approximation to a true nonlinear structure, then ε will capture the approximation error. Both of these artifacts of representation raise practical issues for econometrics, but I wish to set them to one side. The residual may also reflect genuine features of the world – also in at least two ways. First, variables that are omitted from the structural representation may, in effect, be impounded in the residual. Such an interpretation is invoked when it is said that the residual reflects causes that are not the primary focus of an investigation. Second, the residual may, in fact, simply reflect fundamental indeterminism – the manner in which the probabilistic character of the causal relationship is represented. In the first case, ε is a portmanteau term for omitted or latent variables that are real causes of z . In the second case, ε is part of the fundamental characterization of z itself.

Practically, there may be no operational distinction between the first, epistemological interpretation and the second, ontological interpretation. It would only be in the impossible situation in which we knew *all* of the distinct causes of z that we could know which of the two applied in a particular case. Economists refer to residuals under either interpretation as *shocks*. Shocks are not noncausal artifacts, but causally significant, unobserved or unobservable real features of the world. In practice, residuals could be a composite of shocks (omitted causes +

² “Causes on the right” is a common, but not always respected, convention in econometrics. It is sometimes helpful to represent the causal relationship more explicitly: Cartwright (2007, p. 13) suggests “ $\overset{c}{=}$ ” and Hoover (2001, p. 40) suggest “ \leftarrow ” as directed *causal* equalities.

fundamental indeterminism) and (measurement + specification) error. Our interest here, however, is only with shocks.

Equation (1) is a contemporaneous or static relationship in which we assume that variables are not time dependent. In contrast, an example of a time-dependent or dynamic relationship is provided by the stochastic difference equation:

$$(2) \quad z_t = \alpha x_t + \beta x_{t-1} + \gamma z_{t-1} + \varepsilon_t,$$

where γ is a new parameter and the subscripts indicate time periods $t = 0, 1, 2, \dots$. If there are no issues of approximation or specification error and (2) is causally apt, then the residual ε_t represents a shock. The cumulation over time of such shocks forms a (stochastic) *trend*.³

One version of scientific realism is characterized by the assertion that the theoretical posits of our best scientific explanations are real. Some critics of scientific realism ground their criticism in the unobservability of some of these entities (van Fraassen 1980). There is something question-begging in grounding the criticism in a stipulation that observation is restricted to cases involving the unaided senses. Not only is that a hard line to draw (do I cease to observe when I put on my glasses, though I observe precious little without them), it also defies the common usage of ordinary people and scientists who include in “observation” knowledge gained through acquaintance mediated through various instruments. Observations may be obtained through shorter or longer chains of instrumentation and with greater or lesser cooperation from other theories and models. At a rather ill-defined stage, indirect observation is perhaps better termed “inference.” Yet, neither observation nor indirect inference makes sense without something real to be observed or inferred. A fundamental question for econometric

³ There can also be deterministic trends in which the value of a variable depends noncausally and deterministically on time. Our concern here is only with stochastic trends.

analysis is whether shocks and trends are real and, if so, whether and how they can be observed or inferred.

2. Recovering Shocks

2.1 CAUSAL SUFFICIENCY AND THE IDENTIFICATION OF SHOCKS

Early econometrics typically treated shocks as a nuisance and focused on the problem of estimating the parameters of structural equations.⁴ To keep the analysis practically manageable variables were divided into *endogenous* variables to be “explained” by the econometric model and *exogenous* variables that were regarded as having been determined outside the model.

Counterfactual analysis was mainly thought of as assessing the paths of endogenous variables given alternative paths of exogenous variables – either on the assumption that the shocks would average out to zero or that their variability would be translated into the probability distribution of the endogenous variables.

Given the dominant Walrasian worldview, the distinction between endogenous and exogenous variables was a best pragmatic. In principle, since everything depended on everything else, all variables should be endogenous. Once exogenous variables are abandoned, counterfactual analysis must come down to either working out the consequences of alternative choices of parameters or tracing through the effects of shocks. Shocks can no longer be regarded as a nuisance but take pride of place as the fundamental drivers of the variables within a *regime* defined by a constant set of parameters. But shocks are epistemically opaque.

The problem can be easily understood by considering equation (1) as a single equation embedded in a *directed acyclical graph*, the mainstay of the Bayes-net approach. Think about z

⁴ Haavelmo (1944), Koopmans (1950), Hood and Koopmans (1953). See Duarte and Hoover (2013) for a history of the role of shocks in econometrics)

as sitting at the bottom of the causal order: x and y cause z , but z does not cause x or y .

Econometricians typically have observations on the variables and have to estimate the parameters. Typical estimation methods amount to probabilistic conditioning, so that

$$(3) \quad z = \hat{\alpha}x + \hat{\beta}y + \hat{\varepsilon},$$

where the “hat” over the parameters and the shock indicate that these are estimates formed by conditioning, and the shock is governed by a probability distribution.⁵ The key point is that the shock ε is not directly observable, but can be indirectly inferred – provided that (1) is causally apt – by backing an estimate of it out of (3):

$$(4) \quad \hat{\varepsilon} = z - \hat{\alpha}x + \hat{\beta}y.$$

The epistemic opaqueness arises from the fact that the data alone may not provide enough information to guarantee that (1) was a causally apt characterization of the structure in the first place. Suppose that z was determined by (5) instead of by (1), in which now x is still causally upstream but y is causally downstream from z :

$$(5) \quad z = \alpha x + \varepsilon.$$

An estimate of (5) based on the same set of data as used in (4) yields

$$(6) \quad \tilde{\varepsilon} = z - \tilde{\alpha}x,$$

where the tildes indicate estimates. The critical point is that in general $\hat{\alpha} \neq \tilde{\alpha}$ and $\hat{\varepsilon} \neq \tilde{\varepsilon}$. We require additional knowledge about which structure (1) or (5) is correct. Where is that to come from?

In practice, the most common modern macroeconomic models, known as *structural autoregressions* (or *SVARs*), involve systems of equations like (2) in which all variables are

⁵ The distribution is typically assumed to have mean of zero, which can be guaranteed by including a constant in the conditioning set, which is equivalent to an innocuous normalization of the measurements of the variables; and the distribution is often assumed to be normal or Gaussian, which can be justified in many circumstances by considerations that need not concern us here.

endogenously determined. The key issues at this point do not involve time dependence or dynamics, so it is easier to think of systems of equations like (1):

$$(7.A) \quad x = \alpha_{xy}y + \alpha_{xz}z + \varepsilon^x,$$

$$(7.B) \quad y = \alpha_{yx}x + \alpha_{yz}z + \varepsilon^y,$$

$$(7.C) \quad z = \alpha_{zx}x + \alpha_{zy}y + \varepsilon^z.$$

System (7) can be written more compactly in matrix notation as

$$(8) \quad \mathbf{A}\mathbf{X} = \begin{bmatrix} 1 & -\alpha_{xy} & -\alpha_{xz} \\ -\alpha_{yx} & 1 & -\alpha_{yz} \\ -\alpha_{zy} & -\alpha_{zy} & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} \varepsilon^x \\ \varepsilon^y \\ \varepsilon^z \end{bmatrix} = \mathbf{E}.$$

The first row of \mathbf{A} , which is the first matrix after the first equal sign, corresponds to the parameters of (7.A), the second row to those of (7.B), and the third to those of (7.C).

The epistemic opacity of system (7) or (8) is reflected in part in the fact that the variables do not provide enough independent information to estimate all of the parameters of the system. The classic approach of the Cowles Commission is to *identify* the parameters by restricting their number at least to be no greater than the available distinct sources of information provided by the data. This amounts to assuming *a priori* that some of the parameters (α_{ij} 's) equal zero and/or some of the correlations among the shocks vanish.⁶ The approach is rationalistic in that it is supposed to be grounded in theoretical considerations independent of the available empirical evidence.

Bayes-net approaches (e.g., the PC algorithm) provide an empirical approach to identification—essentially placing zeroes according to the evidence of various conditional independence relationships discoverable from the data themselves (Spirtes et al. 2000, ch. 5,

⁶ Other sorts of restrictions may also be identifying so long as they reduce the dimensionality of the parameter space to be no greater than the dimensionality of the data space.


Cooper 1999). The best case for the application of Bayes-net algorithms occurs when the data set is causally sufficient. Data are *causally sufficient* when every variable that causes more than one variable in the data set is included in the data set. Causal sufficiency guarantees that, if the system of equations excludes any variable from explicit consideration, it can be impounded in *one* of the shocks and that the shocks are not probabilistically dependent on each other. Causal sufficiency is a stringent condition. But even when causal sufficiency obtains, if the unobservable real underlying structure does not contain zero parameters ($\alpha_{ij} = 0$ for some variables i and j) or does not place the zero parameters in the right places, a search algorithm may be unable in principle to recover the structure completely from the data.

If the data are causally sufficient and if the causal structure of the system is known – either *a priori* or empirically – then the shocks can be easily be recovered from the data just as in equation (4). Figure 1 shows a causal graph that is in fact recoverable with Bayes-net algorithms for acyclical graphs. It corresponds to the estimated version of (8) with absent possible causal linkages indicated by zeroes in the estimated \mathbf{A} matrix:

$$(9) \quad \hat{\mathbf{A}}\mathbf{X} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -\hat{\alpha}_{zy} & -\hat{\alpha}_{zy} & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} \hat{\varepsilon}^x \\ \hat{\varepsilon}^y \\ \hat{\varepsilon}^z \end{bmatrix} = \hat{\mathbf{E}},$$

where the “hats” once again indicate estimated values.

Figure 1



2.2 MACROECONOMICS AND THE REALITY OF SHOCKS

The central problem is to understand how we should think about the estimated shocks: Are they real? Or are they simply artifacts of the representation and estimation procedure? These are not merely philosophers’ questions, but are in fact questions that economists themselves raise – often only implicitly, but sometimes explicitly. A good example of addressing the questions explicitly is found in Glenn Rudebusch’s “Do Measures of Monetary Policy in a VAR Make Sense?” (1998a).⁷ For Rudebusch (1998a, p. 907) monetary policy analysis – that is counterfactual analysis – would be easy in the context of a complete structural model. The SVAR model is one that focuses on the causal articulation of only the contemporaneous variables, as in system (7) or equation (8), but allows the dynamic structure (omitted from our examples) to reflect only the associations or correlations among variables rather than the complete causal articulation. Focusing on a single equation in the SVAR for the variable that represents the monetary policy instrument (usually the overnight interbank borrowing rate of interest known as the Federal funds rate), Rudebusch (p. 908) asks two questions: First, does the equation correctly represent the reaction of monetary policymakers – that is, does it contain the correct causal variables and estimate the correct parameters? Second, do the residuals plausibly represent the shocks to monetary policy? Rudebusch argues that SVARs are useful guides to monetary policy only if both questions can be answered, yes. He is, in fact, skeptical.

One of Rudebusch’s (p. 908) concerns is that parameter estimates in the analogue to (7) and (8) may not be stable over time. Though he does not analyze this issue further, there are two quite distinct possibilities. First, the system may be an accurate causal articulation, but through a policy action or a “natural experiment,” the underlying parameter has actually changed. An estimate of the system that failed to account for such a change would generate parameters that

⁷ Sims (1998) comments on the original paper, and Rudebusch (1998b) replies.

were not structural, but artifacts that combined the information of distinct regimes. The issues are, however, not the ontological ones of the correctness of the causal structure or the reality of the shocks, but are epistemological or methodological. On the epistemological side, the values assigned to the parameters and shocks are questionable. On the methodological side, we are pointed to the necessity of properly accounting for regime change.⁸

The second possibility is different. If the system fails to be causally sufficient, the actual estimates of parameters and residuals are really just complex conditional associations without straightforward causal interpretations. Worse still, such estimates need not be stable, as changes in variables omitted from the causal representation may induce changes in these conditional associations. The residuals cannot legitimately be regarded as shocks, because they lack the causal articulation that is the hallmark of a real variable. The observed instability of the estimates is, in this case, a signal of their unreality. These two possibilities are conceptually distinct. One of the central difficulties of econometrics is that they are hard to distinguish in practice.

The second source of instability, as it is grounded in the failure of causal sufficiency, is related to Rudebusch’s (p. 916) more general concern with the omission of structure. It is possible that a simpler system could be nested within a more complex system and yet both be causally sufficient. But such a nesting relationship is very special, and a general worry is that any overly simplified system would not be causally sufficient and that, as a result, the residuals could not be legitimately taken to represent shocks. It is puzzling and inconsistent that Rudebusch (p. 926) argues that the failure of different SVAR specifications to agree amongst themselves on the actual values for monetary-policy shocks “would seem to be a critical failing.”

⁸ In fact, Hoover (2001, chs. 8-10) investigates methods of using such changes of regime within a stable causal structure to extract information about the causal structure itself.

The notion that estimates are more reliable if they are robust to alternative specification is an appealing one to many economists and to some philosophers, but it is hardly defensible if the estimates are meant to characterize real causal parameters and shocks; for alternative specifications will rarely all satisfy causal sufficiency (see Cartwright 1991; Hoover 1995; Hoover and Perez 2000). Rudebusch (p. 928) comes round to endorsing this view indirectly when he observes that the failure of different SVARs to agree on the shocks “may reflect the fact that different authors have mined the data in different ways . . .”

What could be evidence for the success of an SVAR at recovering real shocks?

Rudebusch argues that consistency of SVAR evidence with evidence from a completely different source would add weight to the identification. The suggestion sounds similar to the robustness standard that we just rejected, but is really quite different. It is more closely related to the variety-of-evidence thesis (see Claveau 2013) and is exemplified by the philosopher’s science stock example of reality of molecules being underwritten by the multitude of different methods that produce consistent estimates of Avogadro’s number (Perrin 1913[1990]). Rudebusch’s deployment of the standard does not lead to the same happy outcome: he claims that methods of estimating shocks based on financial markets and the historical record of monetary policy actions by the Federal Reserve fail to conform to the shocks typically recovered from actual SVARs (pp. 920-927).

Related to using a *variety of evidence* as a desideratum of ontological independence, Rudebusch (p. 928) suggest another standard that is both evidence for, and a valuable property of, independent shocks – namely, *transferability*. In particular, he suggests that if we are successful at identifying ontologically independent shocks, then the estimates of these shocks themselves may be used as instrumental variables in estimating the coefficients of other causal

structures. This is an example of Hacking’s (1983, pp. 22-24) criterion “if you can spray them, then they are real,” used as evidence for the ontological independence of unobservable posits of theory. The key idea is that an estimated value may be theory-laden (or in Rudebusch’s case, model-laden), but its ontological independence is established by using it (the “spraying”) in an independent context.

Sims (1999) provides an example in which the ontological independence of the shocks is assumed. He poses a counterfactual question: Had the Federal Reserve managed monetary policy the same way that the modern Federal Reserve under Alan Greenspan’s chairmanship managed monetary policy, would the Great Depression have been avoided? His approach is to estimate an SVAR for both a period that includes the Great Depression as its last subperiod and for the Greenspan era (1987-end of the sample). In each, he identifies the equation for a short-term rate of interest as the monetary-policy rule. He substitutes the rule for the Greenspan era into the system for the Great Depression. Finally, using the shocks recovered from the original SVAR for the Great Depression, he simulates the paths for important economic variables to see how they would have behaved facing the same external conditions (i.e., the shocks) but confronting them with a latter-day monetary policy. The counterfactual experiment is problematic in that there are good reasons to question the legitimacy of the assumption that the rest of the SVAR for the Great Depression would remain stable in the face of the changed monetary-policy rule (see Hoover and Jorda 2001 and Hoover 2011). Nonetheless, the main point for our purpose is the implicit assumption that the estimated shocks for the Great Depression can be legitimately treated as observed variables to drive the counterfactual simulation.

3. The Ontology of Latent Variables

Sims’s use of shocks in counterfactual analysis presupposes that the SVAR from which they were estimated was a causally apt articulation of a causally sufficient set of variables. Causal sufficiency, however, is a stringent condition that may be hard to fulfill and that may not be readily verifiable. Little effort is, in fact, made to verify it in typical econometric applications. Causal sufficiency can easily fail if a common cause is omitted from the SVAR – either a variable that could be easily measured or one that is truly *latent* – i.e., one that is causally important but for which there is no independent method of measurement. Reichenbach’s (1956, p. 157) *principle of the common cause* states that if two variables are probabilistically dependent either one causes the other or there is a common cause of both. The omission of a common cause, therefore, may result in treating variables as directly causally connected that are, in fact, probabilistically independent conditional on their common cause.


Bayes-net causal search algorithms based on the causal Markov condition, a generalization of Reichenbach’s principle, may discover the traces of omitted variables in the patterns of probabilistic dependence and independence among variables (Sprites *et al.* 2000, ch. 10). What is important at this juncture is not the details, but the fact that this may present an opportunity for recovering a latent common cause. For example, suppose that as in Figure 2, C is a common cause of A and B , then Reichenbach’s principle supposes that conditioning on C would render A and B probabilistically independent. C is said to *screen off* A from B (and vice versa).

Cartwright (1993) objects that we have no right to assume that C will act as a screen. In her particular example, she claims treats A and B as the deterministic byproducts of an

indeterministic process.⁹ When *C* is *off* neither *A* nor *B* is produced; when *C* is *on*, *A* and *B* are both produced (*A* if, and only if, *B*), but they are produced only with a probability less than one. In this case, *C* does not screen off *A* and *B*. Since *C* being *on* provides only probabilistic knowledge that *A* or *B* is produced; while the additional knowledge that *A* is produced makes it certain that *B* is produced.

Cartwright’s assumption that *A* and *B* are deterministic byproducts is inessential to her point. The essence of the point is that she envisages the causal linkage itself to be indeterministic, and in such as case, the probability of *B* conditional on *C* and *A* will be greater than the probability of *B* conditional on *C* alone, even if the former probability is less than one. Cartwright challenges the ontological assumptions at the heart of the econometrician’s and the Bayes-net practitioner’s partitioning of the world into structure (i.e., the deterministic relationships among variables) and shocks (the bearer of probabilistic variation). They see the arrows in Figure 2 as deterministic linkages; Cartwright sees the arrows themselves as probabilistic.


Figure 2



⁹ See Hoover (2001, pp. 92-97) for a detailed discussion of Cartwright’s case.

Papineau (1991) analyzes such situations differently. He argues that there is an omitted variable (F) that indicates whether or not the causal linkage is operative (i.e., when the cause has “fired”). So, Figure 3 replaces Figure 2, and the arrows correspond to deterministic links. If F is itself stochastic (i.e., if it is driven by shocks not shown in the figure), then naturally C will not, just as Cartwright suggests, screen off A from B . But F will act as a screen. Cartwright replies simply that no F may be available to us: wanting it does not make it so.


Figure 3



Surely, Cartwright is correct that we may not have any independent method to observe the necessary F . But could it be a latent variable? We could easily calculate what F would have had to be in order to serve as a screen. In other words, we could use the relationship between A and B as a device to measure or indirectly observe F (cf. Boumans 2005, ch. 5, appendix). How should we regard the output of such a measurement? One possibility is that F is just an artifact: we can always calculate it *ex post*, but it has no ontological independence of its own. Another possibility is that we can show that the F calculated in one context does work in another. For example, suppose that we recover F from A and B and that we then find that our estimated F screens off the probabilistic dependence of two different variables X and Y . This might suggest

that, ignoring C , the data were related as in Figure 4. Let \hat{F}_{AB} be the estimate of F recovered as the screen of A from B and analogously for other variables. The ontological independence of F could receive substantial support if it turns out that $\hat{F}_{AB} \approx \hat{F}_{AD} \approx \hat{F}_{AE} \approx \hat{F}_{BD} \approx \hat{F}_{BE} \approx \hat{F}_{DE}$. This could be seen as an example of the evidential variety or transferability that were regarded as hallmarks of ontological independence in section 2.2. Such a finding would go a long way to undermine the salience of Cartwright’s skepticism of the implicit causal ontology shared by structural econometrics and Bayes-net approaches.

Figure 4



The logic of extracting a latent variable as an implicit common cause is essentially one of finding what a variable would have to be to turn a causally insufficient set of variables into a causally sufficient set. It can work only when, in addition, we are lucky enough that the causal structure (or at least the relevant part of it) is recoverable. Bayes-net methods, for instance, sometimes are able to recover only an equivalence class of causal structures, rather than a unique

structure.¹⁰ The approach requires a good deal of systemic knowledge and relies on the delicate statistical tests for conditional independence, which can mislead when the departures from causal sufficiency are too great. Recovering ontologically independent shocks from SVARs thus poses a significant methodological challenge for causal modeling. The matter may, as we will suggest presently, be more tractable if we turn from recovering shocks to recovering trends. Since trends are accumulations of shocks, we must first examine the dynamics that we had previously set to one side.

4. Dynamic Processes

4.1 CAUSAL AMBIGUITY IN DYNAMIC MODELS

Dynamic processes are often complex, but the main points can be made with some simple examples. A very simple univariate dynamic process is one in which the value of a variable (x) in one period (t) depends on its value the period before plus a shock (ε):

$$(10) \quad x_t = \rho x_{t-1} + \varepsilon_t,$$

where ρ is a parameter. If $|\rho| < 1$, then x may be represented equivalently as a weighted sum of all the past shocks where the weights decline geometrically:¹¹

$$(11) \quad x_t = \sum_{j=0}^{\infty} \rho^j \varepsilon_{t-j}.$$

¹⁰ Acyclic graphs that have the same skeletons (i.e., the same variables are connected by causal arrows without attention to their direction) and the same unshielded colliders are observationally equivalent (Pearl 2000, theorem 1.2.8, p. 19). An unshielded collider is a pattern in which two variables that are not connected directly are the source of causal arrows that point into the same third variable. Figure 1 is an example of an unshielded collider.

¹¹ Equation (11) follows from (10) by lagging (10) and substituting the expression into the right-hand side of (10) to eliminate for x_{t-1} . This introduces a term $\rho^2 x_{t-2}$, which is eliminated in the same way. After n such substitutions, we are left with an expression involving the shocks and $\rho^n x_{t-n}$. As $n \rightarrow \infty$, this term approaches zero, and the expression on the right-hand side of (11) remains.

In principle, (11) shows that any shock has a permanent effect on the value of x ; but the fact that the weights decline, mean that practically the “memory” of the shock is wiped out relatively quickly.

Roughly speaking a *stationary* time series is one for which the moments (e.g., the mean and variance) are finite and independent of time, and for which the covariance (or correlation) between its values at different times depends on how far apart those times are but not on the absolute time (e.g, the calendar date). (See Hendry 1995, p. 42 for a more formal definition). If the shock ε_t is drawn from an identically, independent normal distribution, then it fulfills the definition of a stationary time series. But in fact, so does x_t in (10) or (11): the expressions for the mean, variance, and covariances will be more complicated, but they fulfill the criteria of stationarity nonetheless.

Consider a dynamic process involving two variables:

$$(12.A) \quad x_t = \beta y_t + u_t,$$

$$(12.B) \quad u_t = \rho u_{t-1} + \varepsilon_t.$$

The value of the variables in (12.A) can be given an equivalent representation:¹²

$$(13) \quad x_t = \rho x_{t-1} + \beta y_t + \rho \beta y_{t-1} + \varepsilon_t.$$

If we define $\Pi_1 = \rho$, $\Pi_2 = \beta$, and $\Pi_3 = \rho\beta$, (13) may be rewritten as:

$$(14) \quad x_t = \Pi_1 x_{t-1} + \Pi_2 y_t + \Pi_3 y_{t-1} + \varepsilon_t,$$

subject to the restriction that $\Pi_1 \Pi_2 = -\Pi_3$.

Both the system (12) and equation (14) reproduce the same path of values for x . They are phenomenally equivalent. But they are not ontologically equivalent. In (12.A) u is a real

¹² To derive (13), lag (12.A) one period, multiply by both sides ρ , subtract the result from the original (12.A); then substitute from (12.B) and rearrange to get (13).

variable – perhaps omitted from explicit analysis or latent, but a genuine cause with its own dynamic history nonetheless. In (14) u has disappeared. Equation (14) is equivalent to (12.A) but only in the very special case that the restriction on the Π_j 's is fulfilled. If, for example, the values of the Π_j 's were drawn at random, then the restriction could be fulfilled only on a set of Lebesgue measure zero. That is, not however, grounds to reject (14) as a legitimate representation; for there may be good economic reasons for the relevant restrictions to be fulfilled.¹³ The key point is that we can always back u_t out of (12.A); but what is its ontological status? If system (12) is an apt articulation of the causal structure, then this is a way to measure an ontologically independent latent variable. If equation (14) is the apt articulation of causal structure, then u_t is simply an artifact of the representation.

The problem posed by u_t will serve as an analogy for the more complex problem of trends. But before we can get there, we need to consider some other aspects of dynamics.

4.2 NONSTATIONARY PROCESSES

Consider equation (10) with the modification that $\rho = 1$, so that

$$(15) \quad x_t = x_{t-1} + \varepsilon_t.$$

Then (11) would be replaced by:

$$(16) \quad x_t = x_{t-n} + \sum_{j=0}^{n-1} \varepsilon_{t-j},$$

where x_{t-n} is an initial condition n periods earlier, which does not vanish as $n \rightarrow \infty$, since there is no weight (the equivalent to ρ^j in (11)) that gets smaller as n gets bigger – i.e., $\rho^j = 1$ for all j . Unlike in (11) a shock in (16) transmits its full effect permanently to x – it has perfect “memory.”

¹³ See Hoover (2001, pp. 166-170) for an analogous case.

A variable such as x_t is an example of a *nonstationary* time series. Its mean and variance are both functions of time (Hendry 1995, section 3.8). Yet, if we subtract x_{t-1} from both sides of (15) to yield

$$(17) \quad \Delta x_t = x_t - x_{t-1} = \varepsilon_t,$$

we see, since ε_t is a stationary random shock, that the first-difference of this nonstationary time series is itself stationary.

A time series that becomes stationary after being differenced once is said to be *integrated of order 1* – often written I(1). The order of integration is the number of times a time series must be differenced to achieve stationarity. Equation (15) is a particularly simple case, as all that remains after differencing is a single random shock. But cases in which what remains is a more complex, yet still stationary expression, such as the right-hand side of equation (11) are equally cases of integration of order 1. A time series that is already stationary does not need to be differenced to achieve stationarity and so it is integrated of order zero – that is, I(0).

A time series such as (15) is called a *random walk*. If x is the distance between an aimlessly wandering drunk and his front door, then (15) says that, at any step t , he is as likely to walk closer as further away. The best predictor of where he will be on the next step is just wherever he happens to be standing right now.

Two I(1) time series may be related. Imagine that the aimless drunk is followed by a deeply concerned friend. The friend follows a path dictated in part by his desire to walk safely, in part by his desire to keep out of sight, and in part by the overriding consideration not to lose sight of the drunk. The two friends may get closer or further apart, but they will never get too far apart. The path of the friend will differ in detail from the path of the drunk, but over the longer

run, it is the path of the drunk that determines the main line of the path of the friend. If the distance between them increases too much, the friend accelerates his pace toward the drunk.

We can think of the path of the drunk as a *stochastic trend*: it is just the cumulation of his lurches and stumbles – the shocks to his position. The path of the friend can be thought of as stationary deviations from the trend. These must be stationary, because if they were nonstationary, the variance would not be finite and the friend would at some point, contrary to our initial assumption, drift arbitrarily far from the drunk. For the same reason, the distance between the drunk and the friend must be stationary. As a result the path of each friend is dominated by the single stochastic trend: the drunk generates the trend; the friend adjusts to it. And this is a general property. Imagine that the younger brother of the friend wonders just what is up and follows his brother (not the drunk), but keeps his distance.¹⁴ Then his path will also be dominated by the stochastic trend of the drunk. There are now three paths – one trend and two adjusting. In any system of such variables, the number of trends plus the number of adjusting relationships adds up to the total number of variables.

Let x_t be the position of the drunk and y_t the position of the friend. The time series are said to be *cointegrated* if both are individually $I(1)$ and a linear combination of them is stationary: $x_t - \beta y_t =$ a stationary or $I(0)$ process. And if their positions are cointegrated, they will not drift too far apart.

Cointegration shares some features with correlation. Both are species of probabilistic dependence. When two time series are correlated, though not perfectly correlated, they move roughly together. Similarly, when two series are cointegrated they move roughly together,

¹⁴ After having tweaked the commonplace example of cointegration as man following a drunk with the addition of a brother following the man, I discovered the structurally identical explication in Murray (1994), in which the drunk is followed by two dogs rather than two brothers.

although this may be evident only if we take a longer time horizon. Thus, if we track the drunk and his friend for five minutes, we may not see a pattern to the friend’s path; yet if we track them for an hour, we see that the friend never actually loses sight of the drunk.

The common statistics used to analyze correlation and many related notions require stationary data. Thus, Pearson’s correlation coefficient – a mainstay of statistical analysis – is an exact analogue to the parameter that governs probabilistic dependence in a normal distribution, which is a stationary distribution. There is no similar analogy with a nonstationary distribution, which is the main source of the phenomenon of “spurious correlation.”¹⁵

Correlation is an intransitive relationship: e.g., if A and B are probabilistically independent and both cause E (a configuration referred to in the Bayes-net literature as an “unshielded collider”), A is correlated with E and E is correlated with B , but A is *ex hypothesi* not correlated with B . (Figure 1 is an example of an unshielded collider.) In contrast, cointegration is transitive: if A is cointegrated with E and E with B , then A and B must also be cointegrated. If this were not so, then A and B could drift infinitely far apart; but E cannot drift infinitely far from either of them; thus if E tracks A it loses B and if E tracks B it loses A – either way, a contradiction.

With cointegration, there is no analogue to conditional correlation. For example, if A , B , and C are stationary variables of a closed universe with deterministic causal connections and each pair is correlated, then there is one of the variables such that conditioning on that variable eliminates the correlation between the other two. In a similar scenario in which the three

¹⁵ Sober (2001), for example, suggests that two trending series, the price of bread in Britain and sea levels in Venice are correlated because they are both trending upward. This is *spurious correlation* because the series are trending and, therefore, nonstationary; so that ordinary correlation measures that were designed with reference to stationary probability distributions (e.g., Pearson’s correlation coefficient) do not bear the same relationship to nonstationary distributions and do not in the nonstationary case indicate probabilistic dependence. See Hoover (2003) for a detailed discussion of Sober’s example and Hoover (2009) for a more general discussion.

variables are nonstationary and pairwise cointegrated, nothing similar happens. In fact, if two variables are cointegrated, the fact of their cointegration is robust to the introduction of further variables. A new variable may not be cointegrated with the original pair or it may be, in which case transitivity guarantees that the cointegration relationship is unaffected by the third variable.

It is worth noting that, although cointegration is similar in some respects and different in others from correlation, a variable may be both cointegrated with another and correlated with it. A variable can be decomposed into a stationary and a nonstationary component. It is certainly true that the nonstationary components of two I(1) time series may be correlated. It may also be true, though I have not seen this worked out, that there is a nonstationary analogue to correlation and conditional correlation. It is just that cointegration is not a candidate to be that analogue.

5. Recovering Trends

5.1 NEAR DECOMPOSABILITY AND A HIERARCHY OF CAUSAL STRUCTURES

The decomposition of a nonstationary system of equations into a stationary and nonstationary part is closely related to Herbert Simon’s (1996, ch. 8) analysis of hierarchies of complex systems based on the notion of near decomposability. A system is *decomposable* in Simon’s account when some parts are tightly related to one another so that they form a unit, and the unit interacts with other units only as a whole. Decomposability is rare, yet near decomposability is more common. A system for Simon is *nearly decomposable* when

(1) . . . the short-run behavior of the component subsystems is approximately independent of the short-run behavior of the other components; [and] (2) in the long run the behavior of any one of the components depends in only an aggregate way on the behavior of the other components. [Simon 1996, p. 198]¹⁶

It is worth quoting Simon’s concrete example of a nearly decomposable system in full:

¹⁶ See Simon and Ando (1961) and Simon and Iwasaki (1988) for a formal treatments.

Consider a building whose outside walls provide perfect thermal insulation from the environment. We shall take these walls as the boundary of our system. The building is divided into a large number of rooms, the walls between them being good, but not perfect, insulators. The walls between rooms are the boundaries of our major subsystems. Each room is divided by partitions into a number of cubicles, but the partitions are poor insulators. A thermometer hangs in each cubicle. Suppose that at the time of our first observation of the system there is a wide variation in temperature from cubicle to cubicle and from room to room – the various cubicles within the building are in a state of thermal disequilibrium. When we take new temperature readings several hours later, what shall we find? There will be very little variation in temperature among the cubicles of each room, but there may still be a large temperature variation *among* the rooms. When we take readings again several days later, we find an almost uniform temperature throughout the building; the temperature differences have virtually disappeared. [Simon 1996, p. 198]

To take an economic example, on a quarter-to-quarter or year-to-year basis, the level of employment in an economy is dominated by a complex interaction among relatively quickly adjusting factors, such as wage rates, interest rates, private and government expenditure. Yet, on a century-to-century basis, the level of employment is dominated by the slowly changing level of overall population, which probably has little to do with the detailed factors that govern employment in the short run, though it may well have to do with the also relatively slowly changing level of technology.

Similarly, the rates of return on various financial assets on a minute-to-minute or day-to-day basis are determined mainly by the rapid process of arbitrage in financial markets, while the aggregate level of all interest rates on a year-to-year basis is determined by the more slowly developing profitability of industry in general.

Simon provides no criterion for distinguishing fast and slow adjustment. My two economic examples, in fact, suggest that it will often be a matter of degree or a matter of how the analysis is contextualized. In the case of simple dynamic processes, the speed of adjustment may be captured by parameters governing time dependence, such as ρ in equations (10)-(13): $|\rho|$ closer to zero corresponds to slow adjustment; $|\rho|$ closer to 1 to fast adjustment. Yet, it is

important to note that there is a *sharp* distinction between $|\rho| < 1$ and $|\rho| = 1$. The probability distributions governing process in which $|\rho| = 0$ and $|\rho| < 1$ differ in detail, but belong to a single family of stationary distributions. The probability distributions governing $|\rho| = 1$ belong to an entirely different family of distributions governing Brownian motion (Hendry 1995, ch. 3, section 10). The transition between the two families represents a kind of critical phenomenon familiar in certain physical applications (e.g., see Batterman 2001). There is a sense, then, in which the long-run causal structure emerges from the short-run structure at the point of transition between stationary ($I(0)$, $|\rho| < 1$) processes and nonstationary ($I(1)$, $|\rho| = 1$) processes. The short-run and long-run causal structures can have distinct ontologies and may, in fact, pull in opposite directions.

5.2 COINTEGRATION AND STOCHASTIC TRENDS

To work with a concrete example, a typical *structural cointegrated vector autoregression* (*structural CVAR*) for three variables takes the form:

$$(18) \quad \begin{bmatrix} 1 & \alpha_{xy} & \alpha_{xz} \\ \alpha_{yx} & 1 & \alpha_{yz} \\ \alpha_{zy} & \alpha_{zy} & 1 \end{bmatrix} \begin{bmatrix} \Delta x_t \\ \Delta y_t \\ \Delta z_t \end{bmatrix} = \text{stationary dynamic component} + \begin{bmatrix} \pi_{xx} & \pi_{xy} & \pi_{xz} \\ \pi_{yx} & \pi_{yy} & \pi_{yz} \\ \pi_{zx} & \pi_{zy} & \pi_{zz} \end{bmatrix} \begin{bmatrix} x_{t-1} \\ y_{t-1} \\ z_{t-1} \end{bmatrix} + \begin{bmatrix} \varepsilon^x \\ \varepsilon^y \\ \varepsilon^z \end{bmatrix}$$

$\mathbf{A} \qquad \Delta \mathbf{X}_t \qquad \qquad \qquad \mathbf{\Pi} \qquad \mathbf{X}_{t-1} \qquad \mathbf{E}$

The left-hand side of (18) is essentially the same as the left-hand side of (8), except (trivially) the minus signs are omitted in the \mathbf{A} matrix and (significantly) the dependent variables are expressed as first differences. If the variables in \mathbf{X} are nonstationary ($I(1)$), then those in $\Delta \mathbf{X}_t$ are stationary ($I(0)$), so that all the issues of causal structure that were raised for equation (8) (or system (7)) remain pertinent. In particular, the need for causal sufficiency as a foundation for an apt

representation of the causal structure and as a prerequisite to backing out the short-run component of the shocks is unchanged.

Our interest, however, is on the right-hand side, particularly on the term $\Pi\mathbf{X}_{t-1}$, which represents the long-run relationships among the data. If the variables in \mathbf{X} are I(1) and, if they are all cointegrated, then they must share at least one common trend. Here, in fact, they must have either one or two common trends. If they have zero, then, contrary to our assumption, the variables would be stationary; if they have three, then, again, contrary to assumption, the variables would not be cointegrated. We pointed out previously that for a single dynamic process, such as equation (10), $|\rho| = 1$ implies a stochastic trend. The analogue for a system of equations is that there is a stochastic trend for each eigenvalue of Π equal to one. In the vernacular of the time-series econometrician, such a system is said to have a *unit root*. It is also a mathematical fact that the rank of a matrix such as Π equals the number of variables minus the number of unit roots. Since $\Pi\mathbf{X}_{t-1}$ describes how the variables adjust to each other (i.e., the nature of their cointegrating relationships), this mathematical fact confirms the point made informally with respect to the random-walking drunk, his friend, and his brother that the number of cointegrating relationships and the number of stochastic trends add up to the number of variables.

It is possible statistically to separate the stationary from the nonstationary components of (18) (see Juselius 2006, ch. 7, section 1). The nonstationary component can be written as

$$(19) \quad \begin{array}{c} \left[\begin{array}{c} \Delta\tilde{x}_t \\ \Delta\tilde{y}_t \\ \Delta\tilde{z}_t \end{array} \right] \\ \Delta\tilde{\mathbf{X}}_t \end{array} = \begin{array}{c} \left[\begin{array}{ccc} \pi_{xx} & \pi_{xy} & \pi_{xz} \\ \pi_{yx} & \pi_{yy} & \pi_{yz} \\ \pi_{zx} & \pi_{zy} & \pi_{zz} \end{array} \right] \\ \mathbf{\Pi} \end{array} \begin{array}{c} \left[\begin{array}{c} \tilde{x}_{t-1} \\ \tilde{y}_{t-1} \\ \tilde{z}_{t-1} \end{array} \right] \\ \tilde{\mathbf{X}}_{t-1} \end{array} + \begin{array}{c} \left[\begin{array}{c} \tilde{\varepsilon}^x \\ \tilde{\varepsilon}^y \\ \tilde{\varepsilon}^z \end{array} \right] \\ \tilde{\mathbf{E}} \end{array},$$

where the tilde over variables indicates that the influence of the stationary component on the nonstationary has been removed statistically. Equation (19) reflects Simon’s idea of near decomposability: “decomposable” because of the division between the stationary and the nonstationary, but only “near” because the *levels* of variables (long run) influence the *changes* in the same variables (short run). In our analogy to the drunk and his friend, equation (19) shows how the distance between the drunk and his friend influences the next step of the friend independently of any influences from the friend’s own recent movements or from the change in the position of the drunk. For example, if the gap between them widens, the friend will have to pick up the pace just because the gap *is* wide and, not particularly because it has *become* wider after either his or the drunk’s last step, even though the fact of the gap being wide is not unconnected with the whole history of their steps.

The trends are not themselves represented in (19). Let us assume that there is only one stochastic trend, then in a move that is analogous to the transformation of equations (13) or (14) into system (12) we can represent the long-run component of the variables as

$$(20) \quad \begin{bmatrix} \tilde{x}_t \\ \tilde{y}_t \\ \tilde{z}_t \end{bmatrix} = \begin{bmatrix} x_0 \\ y_0 \\ z_0 \end{bmatrix} + \begin{bmatrix} \phi_x \\ \phi_y \\ \phi_z \end{bmatrix} \sum_{j=0}^{\infty} \tau_{t-j},$$

where the ϕ_i are constant parameters and τ is the shock to the trend, and the variables subscripted 0 are initial conditions. Thus, the trend is

$$(21) \quad T_t = T_{t-n-1} + \sum_{j=0}^n \tau_{t-j},$$

where T_0 is an initial condition. By analogy with equations (15) and (16), the trend in (21) may also be expressed as

$$(22) \quad T_t = T_{t-1} + \tau_t.$$

Equation (22) is also analogous to the dynamic process in system (12) with T in the role of u and τ in the role of ε .

It turns out, once we have estimates of the parameters (π_{ij}) and shocks (ε^j) in (19), that the values of T and the τ are uniquely determined. We can, therefore, always back estimate T and the τ out of equations such as (19) that take account only of observable variables. Without going into the details of the calculation, the approach is to find the sequence of values for τ_t that render the variables cointegrated.

It is possible that one of the variables itself is, in fact, the trend. For example, if z were the trend ($z = T$), then in (19) $\pi_{zx} = \pi_{zx} = 0$ and $\pi_{zz} = 1$ (and in (20) $\phi_z = 1$). These are statistically testable hypotheses on the basis of an estimate of (19).¹⁷

It need not be the case, that the trend is observable. That case raises one of the main ontological issues: which representation, (19) or (20) and (22) together, is causally apt? One possibility is that (19) reflects the ontologically independent causal structure, so that the matrix \mathbf{A} having a unit root is just a contingent fact about the parameters of the system (the elements π_{ij} of the $\mathbf{\Pi}$ matrix). In that case, the trend T and its shocks τ are just representational artifacts. On the other hand, it may be that the trend is ontologically independent, so that it should be regarded as an unobserved latent variable. In that case, the perfectly tuned values of the π_{ij} are just the phenomenal consequence of omitting the stochastic trend from the representation. The question of which representation is causally apt is precisely analogous to the question of whether system (12) (the analogue to (20) and (22) together) or equation (13) (the analogue to (19)) is causally apt.

¹⁷ When the hypothesis is accepted, z is said to be *weakly exogenous* (see Juselius 2006, ch. 11, sections 1 and 2).

The evidence that would favor the ontological independence of the stochastic trend T is of the same general type as favored the ontological independence of the common cause F in section 3 above – namely, evidential variety and transferability. Start with two cointegrated variables. Adding a third variable to the system will not alter that fact: the third variable may be cointegrated with the original two or it may add a new stochastic trend. Depending on the exact details of the causal connections among the variables, the new variable may substantially alter the estimates of the parameters of the CVAR such as equation (18). A recent paper by Søren Johansen and Katarina Juselius (2012) shows, while the precise values of estimates of the common stochastic trend (e.g., T) and the trend shock (τ) will be different, that asymptotically, as the number of observations becomes large, the estimates from the narrower and wider data set will converge. While we cannot rule out for any one set of data that the parameters are perfectly tuned to produce cointegration, it would quite generally require a different tuning to produce the same result in a wider data set. Thus, the stability of the estimates of the trend and its shock across increasingly wider data sets is strong evidence that the trend is, in fact, an ontologically independent, though not directly observed, variable.

The ontology of nonstationary processes conveys an epistemological advantage not enjoyed by stationary processes. The relationship used to identify or measure latent trends is transitive cointegration rather than intransitive conditional probabilistic independence. An important consequence is that knowledge about trends can be built up in a *piecemeal* fashion – we do not have to begin with the notion that everything depends on everything else. Cointegration is a robust property that, once a minimally sufficient set of variables has been identified, does not depend on the details of the causal structure. This is not the case for conditional probabilistic independence, which depends fundamentally on the match between the

underlying causal structure and the representation of that structure in the probability model. The consequence is that we can identify and measure latent stationary shocks only when the shocks together with the explicitly observed variables form a causally sufficient set and when we are fortunate enough that the patterns of causal dependence and independence in that set imply a unique causal structure. Our epistemic projects need not be so ontologically blessed. This is no counsel of despair; rather it is a recognition of real methodological challenges.

The advantages of nonstationarity should not be overstated. Although cointegration provides a smoother path to the identification and estimation of stochastic trends, it does not provide any particular advantages with respect to the determination of long-run causal structure. Having transformed trends into observable variables through indirect measurement, we can add them to the variables of a structural CVAR such as (18), but we will not have represented the long-term causal order of those variables until we can select the right set of restrictions on the Π matrix. To put it another way, the structural CVAR is only partly structural. The nonstationary elements remain a purely phenomenal description of the data.

The very features that were helpful in identifying and measuring trends – that cointegration is transitive and that cointegrating relationships are fundamentally deterministic – are unhelpful at uncovering asymmetrical causal relationships. To infer causal structure, we will typically require a set of methodological resources similar to those required in the nonstationary case: causal sufficiency and measures of conditional probabilistic dependence and independence. Indeed, there will a disadvantage in that the statistical methods for analyzing conditional probability are better developed and more familiar for the stationary case.

Both the advantages and disadvantages of nonstationarity arise from the emergence of the long-run structure in keeping with Simon’s account of near decomposability. Simon treats that

emergence as a case of aggregation. In his rooms/cubicle example, in the long run any one of the thermometers from any one of the cubicles represents the state of the rooms equally well – the detailed information cubicle information is redundant and only the common temperature of the room matters. Choosing one thermometer to represent the long-run state thus is an economical aggregation. Similarly, since all of the cointegrating variables index the same common stochastic trend, they are redundant, and it is possible for some purposes to economize on information. But in a state of long-run equilibrium, the thermometers are also uninformative about the internal structure of the rooms and the building, as they provide no basis for discrimination. Similarly, the commonness of the common stochastic trend is a marker of uninformative. The trends can provide useful information about the long-term causal structure only if they are resituated to display their relationship to short-term changes in the variables.

6. On the Interface of Ontology and Methodology

The discussion so far has been full of details, so it would perhaps be useful to end by drawing the main threads together. The structural vector autoregression is the dominant tool of empirical macroeconomics. It is sometimes taken to be a purely phenomenal description of the data. But economics is a policy science and it is, therefore, in the words of the title of James Woodward’s (2003) book, devoted to “making things happen” – or at least to understanding how they happen. As a result economics needs models that can support counterfactual analysis.

In building such models, economists have relied on ontology that partitions economic processes into a deterministic structure and indeterministic shocks. Such a division does not in itself dictate a particular representation of the causal structure nor provide a particular

identification or basis for measuring the not-directly-observable shocks. The debate between Rudebusch and Sims (section 2.2) simply reminds us that the issue of how to get the right causal structure and to recover the shocks is not an abstract philosophical concern but one that animates practitioners and is even expressed ontologically: are the measured shocks the *real* shocks?

Within the economists’ own ontological presuppositions, answering that question is a matter of ascertaining that the modeled causal structure is the correct one; then the shocks backed out from the model and the actual data will be correct. Cartwright, however, challenges those ontological presuppositions – particularly that the causal connections among variables can be regarded as deterministic. Her challenge has an epistemological grounding: one should not presume – as the strategy for backing out the shocks from a model does presume – that a variable must exist to screen off the probabilistic dependence between two causally unconnected variables unless one can point to the variable.

Presumably Cartwright’s challenge is not meant to be a generic rejection of claims for the reality of unobserved entities. So, an effective answer would be one that gives us grounds for preferring the assumption of indirectly observed, ontologically independent entities. The criteria of variety of evidence (i.e., independent methods of indirect measurement that provide essentially the same values for the shocks) and transferability (i.e., a shock identified in one context appears to be causally efficacious in another) are frequently held to support claims of ontological independence. The methodological challenge is to support such claims in particular cases. As the Bayes-net analysts have taught us, such support may rely on difficult-to-establish claims of causal sufficiency. Difficult, however, is not impossible.

The fact that economic data are frequently nonstationary both poses its own methodological challenges and opens up the ontological possibility of a distinct hierarchy of

causal structures. The cointegrated vector autoregression, which is a preferred technique among time-series econometricians for dealing with nonstationary data, offers an instance of Simon’s near decomposability as it divides dynamic processes into distinct short-run and long-run causal structures. Stochastic trends, which are the nonstationary counterpart to stationary shocks, turn out to be easier to identify and measure because of the ontology of cointegrating relations. Nonetheless, the challenge of the discovery of causal structure remains similar in the nonstationary case to the stationary – a point not much appreciated by economists.

References

- Batterman, Robert W. (2001) *The Devil in the Details : Asymptotic Reasoning in Explanation, Reduction and Emergence*. Oxford: Oxford University Press.
- Boumans, Marcel. (2005) *How Economists Model the World into Numbers*. London: Routledge.
- Cartwright, Nancy. (1991) “Replicability, Reproducibility, and Robustness: Comments,” *History of Political Economy* 23(1), 143-155.
- Cartwright, Nancy. (1993) “Marks and Probabilities: Two Ways to Find Causal Structure,” in F. Stadler, editor. *Scientific Philosophy: Origins and Development*. Dordrecht: Kluwer, pp. 113-119.
- Cartwright, Nancy. (2007) *Hunting Causes and Using Them: Approaches in Philosophy and Economics*. Cambridge: Cambridge University Press.
- Claveau, François. (2013) “The Independence Condition in the Variety-of-Evidence Thesis,” *Philosophy of Science* 80(1), 94-118.
- Cooper, Gregory F. (1999). ‘An Overview of the Representation and Discovery of Causal Relationships using Bayesian Networks’, in Clark Glymour and Gregory F. Cooper, editors. *Computation, Causation, and Discovery*, American Association for Artificial Intelligence, Menlo Park, CA and MIT Press, Cambridge, MA, pp. 3-64.
- Duarte, Pedro Garcia and Kevin D. Hoover (2012) “Observing Shocks,” in Harro Maas and Mary Morgan, editors. *Histories of Observation in Economics*, supplement to *History of Political Economy* 44(5), 226-249.
- Haavelmo, Trygve. (1943) “The Statistical Implications of a System of Simultaneous Equations,” *Econometrica* 11(1), pp. 1-12.
- Haavelmo, Trygve. (1944) “The Probability Approach in Econometrics,” *Econometrica* 12(Supplement), iii-115.
- Hacking, Ian. (1983) *Representing and Intervening: Introductory Topics in the Philosophy of the Natural Sciences*. Cambridge: Cambridge University Press.

- Heckman, James J. (2000) “Causal Parameters and Policy Analysis in Economics: A Twentieth Century Perspective,” *Quarterly Journal of Economics* 115(1), 45-97.
- Hendry, David F. (1995) *Dynamic Econometrics*. Oxford: Oxford University Press.
- Hood, William C. and Tjalling Koopmans, editors. (1953) *Studies in Econometric Method*, Cowles Commission Monograph vol. 14. New York: Wiley, pp. 1-26.
- Hoover, Kevin D. (1995) “Comments on Cartwright and Woodward: Causation, Estimation, and Statistics,” in Daniel Little, editor. *On the Reliability of Economic Models: Essays in the Philosophy of Economics*. Dordrecht: Kluwer, 1995.
- Hoover, Kevin D. (2001) *Causality in Macroeconomics*. Cambridge: Cambridge University Press.
- Hoover, Kevin D. (2003) “Nonstationary Time Series, Cointegration, and the Principle of the Common Cause,” *British Journal for the Philosophy of Science* 54(4), 527-551.
- Hoover, Kevin D. (2004) “Lost Causes,” *Journal of the History of Economic Thought* 26(2), pp. 149-164.
- Hoover, Kevin D. (2009) “Probability and Structure in Econometric Models,” in *Logic, Methodology and Philosophy of Science: Proceedings of the Thirteenth International Congress*. London: College Publications, pp. 497-513.
- Hoover, Kevin D. (2011) “Counterfactuals and Causal Structure,” in Phyllis McKay Illari, Federica Russo, and Jon Williamson, editors. *Causality in the Sciences*. Oxford University Press, pp. 338-360.
- Hoover, Kevin D. and Òscar Jordà. (2001) “Measuring Systematic Monetary Policy,” *Federal Reserve Bank of St. Louis Review* 83(4), 113-137.
- Hoover, Kevin D. and Stephen J. Perez (2000) “Three Attitudes towards Data-mining,” *Journal of Economic Methodology* 7(2), 195-210.
- Johansen, Søren and Katarina Juselius. (2012) “An Asymptotic Invariance Property of the Common Trends Under Linear Transformation of the Data,” Department of Economics, University of Copenhagen, Discussion Paper No. 10-30.
- Koopmans, Tjalling C., editor. (1950) *Statistical Inference in Dynamic Economic Models*, Cowles Commission Monograph 10. New York: Wiley.
- Murray, Michael P. (1994) “A Drunk and Her Dog: An Illustration of Cointegration and Error Correction,” *The American Statistician* 48(1), pp. 37-39.
- Papineau, David. (1991) “Correlations and Causes,” *British Journal for the Philosophy of Science*, 42(3), 397-412.
- Pearl, Judea. 2000. *Causality: Models, Reasoning, and Inference*. Cambridge: Cambridge University Press.
- Perrin, Jean (1913 [1990]), *Atoms*. Woodbridge, CT: Ox Bow Press.
- Reichenbach, H. (1956) *The Direction of Time*, Berkeley: University of California Press.
- Rudebusch, Glenn D. (1998a) “Do Measures of Monetary Policy in a VAR Make Sense?” *International Economic Review* 39(4), pp. 907-931.

- Rudebusch, Glenn D. (1998b) “Do Measures of Monetary Policy in a VAR Make Sense? – A Reply to Christopher A. Sims,” *International Economic Review* 39(4), pp. 943-948.
- Simon, Herbert A. (1996) *The Sciences of the Artificial*, 3rd ed. Cambridge, MA: MIT Press.
- Simon, Herbert A. and Albert Ando. (1961) “Aggregation of Variables in Dynamic Systems,” *Econometrica* 29(2), 111-138.
- Simon, Herbert A. and Yumi Iwasaki. (1988) “Causal Ordering, Comparative Statics, and Near Decomposability,” *Journal of Econometrics* 39(1/2), 149-173.
- Sims, Christopher A. (1998) “Comment on Glenn Rudebusch’s ‘Do Measures of Monetary Policy in a VAR Make Sense?’” *International Economic Review* 39(4), pp. 933-941.
- Sims, Christopher A. (1999) “The Role of Interest Rate Policy in the Generation and Propagation of Business Cycles: What Has Changed Since the ‘30s?’” in Jeffrey C. Fuhrer and Scott Schuh, editors., *Beyond Shocks: What Causes Business Cycles*. Federal Reserve Bank of Boston Conference Series, No. 42. Boston: Federal Reserve Bank of Boston, pp. 121-60.
- Sober, Elliot. (2001) “Venetian Sea Levels, British Bread Prices, and the Principle of the Common Cause,” *British Journal for the Philosophy of Science*, 52(2), 331-46.
- Spirtes, Peter, Clark Glymour, and Richard Scheines. (2000) *Causation, Prediction, and Search*, 2nd edition. Cambridge, MA: MIT Press.
- van Fraassen, Bas. (1980) *The Scientific Image*. Oxford: Oxford University Press.
- Walras, Leon. (1874[1954]) *Elements of Pure Economics*, translated by William Jaffe. Homewood, IL: Richard Irwin, 1954.
- Woodward, James B. (2003) *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.