Recoverability*

Ryan Chahrour
Boston College

Kyle Jurado
Duke University

July 23, 2018

Abstract

When can structural shocks be recovered from observable data? We present a necessary and sufficient condition that gives the answer for any linear model. Invertibility, which requires that shocks be recoverable from current and past data only, is sufficient but not necessary. This means that semi-structural empirical methods like structural vector autoregression analysis can be applied to models with non-invertible shocks. To illustrate these results, we propose a set of structural restrictions capable of separately identifying fundamental shocks from non-fundamental "noise" shocks. In an application to postwar U.S. data, we find that productivity shocks account for less than 15% of the business-cycle fluctuations in aggregate consumption, output, investment, and employment.

JEL classification: D84, E32, C32
Keywords: Invertibility, structural vector autoregression, noise shocks

*Chahrour: Department of Economics, Boston College, 140 Commonwealth Avenue, Chestnut Hill, MA 02467 (e-mail: `ryan.chahrour@bc.edu`); Jurado: Department of Economics, Duke University, 419 Chapel Drive, Durham, NC 27708 (e-mail: `kyle.jurado@duke.edu`).

1 Introduction

Economists often explain economic outcomes in terms of structural "shocks," which represent exogenous changes in underlying fundamental processes. Typically, these shocks are not directly observed; instead, they are inferred from observable processes through the lens of an economic model. Therefore, an important question is whether the hypothesized shocks can indeed be recovered from the observable data.

We present a simple necessary and sufficient condition under which structural shocks are recoverable for any linear model. The model defines a particular linear transformation from shocks to observables, and our condition amounts to making sure that this transformation does not lose any information. This can be done by checking whether the matrix function summarizing the transformation is full column rank. If it is, then the observables contain at least as much information as the shocks, and knowledge of the model and the observables is enough to recover the shocks.

Our approach differs from existing literature because we do not focus on the question of whether shocks are recoverable from only current and past observables. This more stringent "invertibility" requirement is often violated in economic models.[1] For example, it may be violated if structural shocks are anticipated by economic agents.[2] However, in many cases it is still possible to recover shocks using future observables as well. Because there is no reason in principle to constrain ourselves to recover shocks only from current and past data, we focus on the question of whether shocks are recoverable from the data without any temporal constraints.

Non-invertibility is often viewed as a problem from the perspective of using semi-structural empirical methods in the spirit of Sims (1980). The reason seems to be that that the first step of these methods usually involves obtaining an invertible reduced-form representation of the data. But if the structural model of interest is not invertible, so the reasoning goes, then it is impossible to obtain the structural shocks as a linear combination of contemporaneous reduced-form shocks. As a result, it is common practice to verify first that a model is invertible (using tests such as the one in Fernández-Villaverde et al. (2007)) and, if this can't be done, then to resort to fully structural methods, which impose additional theoretical restrictions on the data

[1] For some examples, see Hansen and Sargent (1980, 1991), Lippi and Reichlin (1993, 1994), Futia (1981), and Quah (1990).

[2] As in Cochrane (1998), Leeper et al. (2013), Schmitt-Grohé and Uribe (2012), and Sims (2012).

generating process.[3]

We respond to these concerns by adopting a different perspective on semi-structural methods.[4] We view the reduced-form model simply as a statistical way of characterizing the information in the autocovariance function of the observable processes. While a VAR is undoubtedly the most popular choice, there are generally many different reduced-form representations consistent with the same autocovariance function. Which representation is "best" is essentially a statistical question.

Given a characterization of the autocovariance function, the structural step involves imposing a subset of the economic model's theoretical restrictions to obtain a "structural representation" with shocks that are the shocks of interest. If the structural representation happens to be non-invertible, that does not present any inherent difficulty. In particular, we argue that the identification of non-invertible shocks does not require stronger theoretical restrictions than usual. It just means that some restrictions need to be relaxed, and other need to be introduced in the place.

Understood as a necessary condition for employing semi-structural methods, invertibility has represented a major obstacle in using these methods to answer certain important economic questions. One question that we take up in this paper is: how much of the business cycle is driven by fluctuations in productivity? In order to provide a complete answer to this question, it is necessary to control for the hypothesis that economic agents might partially respond to future developments in productivity before they occur. Not taking this possibility into account could lead us to understate the importance of productivity shocks. This is an especially important concern in light of the fact that most dynamic equilibrium models predict that current actions depend heavily on agents' expectations about future fundamentals.

However, it turns out that it is not possible to adequately control for this hypothesis if we also insist that the underlying structural shocks must be invertible. Suppose that agents receive noise-ridden signals that contain information about future productivity beyond what can be seen in current and past productivity alone. To the extent that they are forward-looking and incorporate this additional information,

[3]This is the original remedy proposed by Hansen and Sargent (1991), and has been adopted by a large part of the literature on anticipated shocks. See the arguments in Schmitt-Grohé and Uribe (2012); Barsky and Sims (2012); and Blanchard et al. (2013).

[4]In fact, this may have been Sims (1980)'s original view (cf. his description on p.15). In his application, he uses an invertible VAR as the reduced-form model, but neither invertibility nor VARs are described as essential features of his proposed empirical strategy.

their current actions will partially depend on future productivity shocks. But because agents cannot tell on the basis of current and past data whether their signals have changed because of an actual change in future productivity or just unrelated noise, an econometrician with the same information or less will not be able to do so either.

Blanchard et al. (2013) conclude from this observation that semi-structural methods are generally inapplicable in situations of this type.[5] However, that conclusion can be avoided by replacing invertibility with recoverability as the appropriate necessary condition for using these methods. As an application of our results, we show that our recoverability condition is satisfied in an analytically convenient model of consumption determination with productivity and non-productivity noise shocks. We then present a new set of structural restrictions that are sufficient to identify the importance of productivity shocks regardless of whether agents receive advance information about future productivity. Using a Monte Carlo exercise, we show how VAR-based semi-structural analysis can be applied in this situation.

Finally, we apply the same procedure to a sample of postwar U.S. data. We do find evidence in support of the hypothesis that current changes in consumption depend on future changes in productivity. However, even after controlling for potential future dependence, we find that less than 15% of the business-cycle variation in consumption can be attributed to productivity shocks. We find even smaller shares when applying the same procedure to output, investment, and employment. This finding represents a challenge for theories of business cycles that rely primarily on beliefs about productivity. In order for these theories to be consistent with such low productivity shares, beliefs about productivity must fluctuate in ways that are mostly unrelated to productivity itself.

Some existing papers have demonstrated that semi-structural methods are not necessarily inapplicable when invertibility fails. Lippi and Reichlin (1994) propose restrictions to set-identify structural shocks, which allow for some forms of non-invertibility. Sims and Zha (2006) propose an iterative algorithm to check whether certain structural shocks are "approximately invertible," even if they are not invertible. Mertens and Ravn (2010) propose restrictions to identify the effects of partially anticipated changes in government spending across a class of models with non-invertible shocks. Dupor and Han (2011) develop a four-step procedure to partially

[5]For a more detailed discussion of the limitations of using VAR analysis models with anticipated shocks, see the review article by Beaudry and Portier (2014).

identify structural impulse response functions even if some forms of non-invertibility are present. Plagborg-Møller (2017) suggests that estimating a moving average model rather than an autoregression can help avoid concerns of non-invertibility. Forni et al. (2017a) propose a method to identify productivity shocks and noise shocks in some models in which agents receive imperfect signals about future productivity, and Forni et al. (2017b) use these same restrictions in an asset-pricing context.

Relative to these papers, our contribution is to introduce recoverability as a more appropriate alternative to invertibility as the essential condition for employing semi-structural methods, and to present a general analysis of the difference between the two. We also provide new insights on the importance of productivity shocks for generating business-cycle fluctuations by proposing a new set of structural restrictions which is capable of separately identifying fundamental shocks from non-fundamental noise shocks regardless of whether agents receive signals about future fundamentals. As it turns out, none of the restrictions that have been proposed in the existing literature are capable of doing this. Throughout the paper, we highlight other points of connection with the literature when they arise.

2 Recoverability Condition

This section presents our main theorem. We begin with some notation and definitions. Consider an arbitrary n_ξ dimensional wide-sense stationary process $\{\xi_t\}$, where the parameter t takes on all integer values.[6] We let $\mathcal{H}(\xi)$ denote the Hilbert space spanned by the variables $\xi_{k,t}$ for $k = 1, \ldots, n_\xi$ and $-\infty < t < \infty$, closed with respect to convergence in mean square. Similarly, we let $\mathcal{H}_t(\xi)$ denote the subspace spanned by these variables over all k but only up through date t. We can then define recoverability in terms of the relationship between $\{\xi_t\}$ and another n_η dimensional stationary process $\{\eta_t\}$ with which it is stationarily correlated.

Definition 1. $\{\eta_t\}$ is "recoverable" from $\{\xi_t\}$ if

$$\mathcal{H}(\eta) \subseteq \mathcal{H}(\xi).$$

This says that each of the variables $\eta_{k,t}$ is contained in the space $\mathcal{H}(\xi)$.[7] That is,

[6] Continuous-time versions of our results can be found in the Supplementary Material.

[7] Plagborg-Møller and Wolf (2018) adopt this same definition of recoverability in their recent work on variance decompositions in linear projection instrumental variables models.

each of these variables is perfectly revealed by the information contained in $\{\xi_t\}$. In the Gaussian case, this can be expressed in terms of mathematical expectations as

$$\eta_{k,t} = E[\eta_{k,t}|\mathcal{H}(\xi)].$$

Another way to think about recoverability is in terms of the econometric concept of identification. If a certain process is recoverable from another, this means that the former is identified from the latter.[8]

Recoverability is different from the familiar concept of invertibility, which has to do with whether one collection of random variables can be recovered only from the current and past history of another.

Definition 2. $\{\eta_t\}$ is "invertible" from $\{\xi_t\}$ if

$$\mathcal{H}_t(\eta) \subseteq \mathcal{H}_t(\xi) \quad \text{for all } -\infty < t < \infty.$$

Since $\mathcal{H}_t(\xi) \subset \mathcal{H}(\xi)$, it is easy to see that invertibility is sufficient but not necessary for recoverability.

An equivalent characterization of recoverability can be given in terms of an appropriate Hilbert space of complex vector functions. We write the spectral representation of $\{\xi_t\}$ as

$$\xi_t = \int_{-\pi}^{\pi} e^{i\lambda t}\Phi_\xi(d\lambda), \tag{1}$$

where $\Phi_\xi(d\lambda)$ is its associated random spectral measure. We say that a $1 \times n_\xi$ dimensional vector function $\psi(\lambda)$ belongs to the space $\mathcal{L}^2(F_\xi)$ if [9]

$$\int \psi(\lambda)F_\xi(d\lambda)\psi(\lambda)^* \equiv \int \sum_{k,l=1}^{n_\xi} \psi_k(\lambda)\overline{\psi_l(\lambda)} f^{(\mu)}_{\xi,kl}(\lambda)\mu(d\lambda) < \infty,$$

where $F_\xi(d\lambda)$ denotes the spectral measure of $\{\xi_t\}$, $\mu(d\lambda)$ denotes any non-negative measure with respect to which all the elements of $F_\xi(d\lambda)$ are absolutely continuous, $f^{(\mu)}_{\xi,kl}(\lambda) = F_{\xi,kl}(d\lambda)/\mu(d\lambda)$ for $k, l = 1, \ldots, n_\xi$, and the asterisk denotes complex conjugate transposition.[10] If we define the scalar product

$$(\psi_1, \psi_2) = \int \psi_1(\lambda)F_\xi(d\lambda)\psi_2(\lambda)^*,$$

[8] Of course, this type of "identification" needs to be understood in an appropriately general sense, since its object is a family of random variables, rather than a vector of parameters.

[9] Whenever the limits of integration are omitted, we will understand them to be $[-\pi, \pi]$.

[10] Recall that $F_{\xi,kl}(\Delta) \equiv E[\Phi_{\xi,k}(\Delta)\overline{\Phi_{\xi,l}(\Delta)}]$ for $k, l = 1, \ldots, n_\xi$ and any Borel set Δ.

and do not distinguish between two vector functions that satisfy $\|\psi_1 - \psi_2\| = 0$, then $\mathcal{L}^2(F_\xi)$ becomes a Hilbert space.[11] Using these definitions, the following lemma gives an alternative characterization of recoverability. Its proof is in the Appendix, together with all other proofs.

Lemma 1. *$\{\eta_t\}$ is recoverable from $\{\xi_t\}$ if and only if there exists an $n_\eta \times n_\xi$ matrix function $\psi(\lambda)$ with rows in $\mathcal{L}^2(F_\xi)$ such that*

$$\eta_t = \int e^{i\lambda t}\psi(\lambda)\Phi_\xi(d\lambda) \quad \text{for all } t. \tag{2}$$

We will say that a process $\{\eta_t\}$ can be obtained from $\{\xi_t\}$ by a "linear transformation" whenever it has a representation of the form in equation (2), and we will call $\psi(\lambda)$ the "spectral characteristic" associated with this transformation. Using this language, Lemma (1) says that $\{\eta_t\}$ is recoverable from $\{\xi_t\}$ if and only if it can be obtained from $\{\xi_t\}$ by a linear transformation.

In this paper, we are interested in determining the conditions under which a collection of structural economic shocks can be recovered from a collection of observable variables. Letting $\{y_t\}$ denote the n_y dimensional observable process and $\{\varepsilon_t\}$ the n_ε dimensional structural shock process, we make the following two assumptions. The first puts weak technical restrictions on the observables, the second defines the theoretical economic model as an arbitrary linear transformation from the structural shocks to the observables.

Assumption 1. *$\{y_t\}$ is stationary in the wide sense and linearly regular. That is, $E[y_{k,t}]$ is a constant and $E[|y_{k,t}|^2] < \infty$ for all $k = 1, \ldots, n_y$, the function $B_{kl}(t,s) = E[y_{k,t}\overline{y_{l,s}}]$ depends only on $t - s$ for all $k, l = 1, \ldots, n_y$, and $\cap_{t=-\infty}^{\infty}\mathcal{H}_t(y) = 0$.*

Assumption 2. *$\{y_t\}$ can be obtained from $\{\varepsilon_t\}$ by a linear transformation with spectral characteristic $\varphi(\lambda)$, where $\{\varepsilon_t\}$ is a zero-mean process with orthonormal values. That is,*

$$y_t = \int e^{i\lambda t}\varphi(\lambda)\Phi_\varepsilon(d\lambda) \quad \text{for all } t \tag{3}$$

with $E[\varepsilon_{k,t}] = 0$ and $E[|\varepsilon_{k,t}|^2] = 1$ for all $k = 1, \ldots, n_\varepsilon$, and $E[\varepsilon_{k,t}\overline{\varepsilon_{l,s}}] = 0$ for $t \neq s$ and all $k, l = 1, \ldots, n_\varepsilon$, as well as for $k \neq l$ and all t, s.

[11] The generalization of the Riesz-Fischer Theorem that is required to establish this fact is proven in Lemma 7.1, Ch. 1, of Rozanov (1967).

Example 1. A special case of the model in equation (3) is when the observables are related to the structural shocks by a linear state-space structure of the form

$$\begin{aligned} &\text{(observation)} \qquad y_t = Ax_t \\ &\text{(state)} \qquad x_t = Bx_{t-1} + C\varepsilon_t, \end{aligned} \tag{4}$$

where x_t is an n_x dimensional state vector. In this case, the spectral characteristic $\varphi(\lambda)$ in equation (3) takes the form

$$\varphi(\lambda) = A(I_{n_x} - Be^{-i\lambda})^{-1}C. \tag{5}$$

The solution to a wide class of linear (or linearized) dynamic equilibrium models can be written in this form.[12] ◇

By Lemma (1), the model in equation (3) says that the observables are recoverable with respect to the structural shocks. Naturally, a knowledge of the inputs of the system is enough to perfectly reveal the outputs. We would like to know when the reverse is true. That is, when can the shocks be recovered from the observables? The following theorem provides the answer.

Theorem 1 (Recoverability)**.** *Under Assumptions* (1) *and* (2)*, the structural shocks* $\{\varepsilon_t\}$ *are recoverable from the observables* $\{y_t\}$ *if and only if* $\varphi(\lambda)$ *is full column rank almost everywhere.*

The intuition for this result can be understood by analogy with the static case. If $\varphi(\lambda) = \varphi$ is a constant matrix, not depending on λ, then the model in equation (3) reduces to

$$y_t = \varphi\varepsilon_t.$$

In order for φ to have a left-inverse, ψ, such that $\psi\varphi = I_{n_\varepsilon}$, it is necessary and sufficient that the matrix φ have full column rank. In that case we can pre-multiply both sides of the previous equation by ψ to obtain the solution $\varepsilon_t = \psi y_t$. It turns out that this logic continues to apply in the dynamic case when $\varphi(\lambda)$ does depend on λ, with the added proviso that this matrix function may fail to be full column rank on a set of at most measure zero.

Before moving on, we make a few remarks regarding the theorem.

[12] Some authors include errors in the observation equation as well as the state equation. Those representations can be rewritten in the form of equation (4) by augmenting the state vector.

Remark 1. A corollary of the theorem is that a necessary condition for the structural shocks to be recoverable is that there be at least as many observable variables as shocks, $n_y \geq n_\varepsilon$. This is intuitive; it isn't possible to recover n_ε separate sources of random variation without observations of at least n_ε random processes.

Remark 2. While in many cases, it is possible to check this condition analytically, there is also a simple numerical procedure that can be used as well. The linear regularity of $\{y_t\}$ ensures that its spectral density, and therefore $\varphi(\lambda)$, has a constant rank for almost all λ. This means that we can draw a number λ_u randomly from the interval $[-\pi, \pi]$ and check whether $\varphi(\lambda_u)$ is full column rank.

Remark 3. From a mathematical perspective, this theorem represents a multivariate generalization of Theorem 10 from Kolmogorov (1941); what we call recoverable he calls "subordinate" (подчиненной). Relative to that theorem, we have incorporated two additional assumptions; namely, that the observable process is linearly regular and that the shock process has orthonormal values. Removing the first leaves our theorem unchanged (although Remark 2 then no longer applies), and removing the second requires us to stipulate that $\varphi(\lambda)$ should be full rank almost everywhere with respect to the spectral measure $F_\varepsilon(d\lambda)$ (as opposed to the Lebesgue measure). We present these more general results in the Supplemental Material.

For the purposes of comparison, we would also like to have a set of necessary and sufficient conditions for the invertibility of the structural shocks. It does not seem that any conditions of this type have been proven in the existing literature, at least not at the level of generality we consider here.[13] Since invertibility is stronger than recoverability, the condition in Theorem (1) must always be satisfied if we are to recover the shocks from current and past observables. Therefore, we can suppose that $\varphi(\lambda)$ is full column rank almost everywhere as we look for the additional restrictions that are needed.

The key step is to recall that, using Wold's decomposition theorem, it is possible

[13]There are places where sufficient conditions appear, however. The condition of Fernández-Villaverde et al. (2007) is one example. Necessary and sufficient conditions for certain types of ARMA models can be found in Brockwell and Davis (1991). Any "fundamental" process, in the sense of Rozanov (1967), is invertible, but the converse is not true. Therefore, conditions that determine whether a process is fundamental are not necessary for invertibility as defined here.

to represent $\{y_t\}$ by a linear transformation of the form

$$y_t = \int e^{i\lambda t}\gamma(\lambda)\Phi_w(d\lambda), \tag{6}$$

where $\Phi_w(d\lambda)$ is the random spectral measure of an r_y dimensional mean-zero process with orthonormal values, $\{w_t\}$, which has the property that the variables w_s, $s \leq t$, form an orthonormal basis in $\mathcal{H}_t(y)$ at each date, and r_y is the rank of $f_y(\lambda)$ for almost all λ.[14] This implies that $\mathcal{H}_t(w) = \mathcal{H}_t(y)$ for all t, so $\{w_t\}$ is both invertible and recoverable from $\{y_t\}$. Using the spectral characteristic from this representation, we can state the following result.

Theorem 2 (Invertibility)**.** *Under Assumptions* (1) *and* (2)*, the structural shocks* $\{\varepsilon_t\}$ *are invertible from the observables* $\{y_t\}$ *if and only if they are recoverable and*

$$\frac{1}{2\pi}\int e^{i\lambda s}\psi(\lambda)\gamma(\lambda)d\lambda = 0 \quad \textit{for all } s < 0,$$

where $\psi(\lambda)$ *is any* $n_\varepsilon \times n_y$ *matrix function satisfying* $\psi(\lambda)\varphi(\lambda) = I_{n_\varepsilon}$ *almost everywhere, and* $\gamma(\lambda)$ *comes from some version of Wold's decomposition of* $\{y_t\}$*.*

The following example uses three very simple models to illustrate the different possible combinations of shock recoverability and invertibility that can arise. In later sections, we will consider models with more explicit motivations. But for now, we observe that the third model illustrates how recoverability can fail to hold even when the number of shocks is equal to the number of observables, due to a form of "dynamic multicollinearity." This situation is not uncommon in the literature on belief-driven fluctuations (e.g. Barsky and Sims, 2012; Blanchard et al., 2013).

Example 2. In each of the following three models, we assume that the structural shocks make up a zero-mean process with orthonormal values.

(i) Recoverable and invertible: $y_t = \varepsilon_{t+1}, \Rightarrow \varphi(\lambda) = e^{i\lambda}$.

(ii) Recoverable but not invertible: $y_t = \varepsilon_{t-1}, \Rightarrow \varphi(\lambda) = e^{-i\lambda}$.

(iii) Neither recoverable nor invertible:

$$\begin{matrix} y_{1,t} = \varepsilon_{1,t} + \varepsilon_{2,t+1} \\ y_{2,t} = \varepsilon_{1,t-1} + \varepsilon_{2,t} \end{matrix}, \quad \Rightarrow \quad \varphi(\lambda) = \begin{bmatrix} 1 & e^{i\lambda} \\ e^{-i\lambda} & 1 \end{bmatrix}.$$

◇

[14] See Rozanov (1967), Ch. 2, Section 3.

We conclude this section with a remark about our use of wide-sense stationarity.

Remark 4. Definitions (1) and (2) apply to wide-sense stationary processes. However, they can be generalized to allow for deviations from stationarity. For example, consider a process $\{\xi_t\}$ that is stationary only after suitable differencing. That is,

$$\Delta^p \xi_t = \zeta_t \tag{7}$$

for some integer $p > 0$, where $\{\zeta_t\}$ is a stationary process. In this case we can define a new process

$$\tilde{\xi}_t(\theta) \equiv \int e^{i\lambda t} \frac{1}{(1-\theta e^{-i\lambda})^p} \Phi_\zeta(d\lambda), \tag{8}$$

which is stationary for each value of θ in $[0, 1)$. We refer to a process $\{\eta_t\}$ as recoverable (or invertible) from $\{\xi_t\}$ whenever $\{\tilde{\eta}_t(\theta)\}$ is recoverable (or invertible) from $\{\tilde{\xi}_t(\theta)\}$ for almost all $\theta \in [0, 1)$.

3 Shock-Specific Recoverability

Given a system of the form in (3), the condition in Theorem (1) determines whether the entire vector process $\{\varepsilon_t\}$ can be recovered. That is, it determines whether $\{\varepsilon_{k,t}\}$ can be recovered for *all* $k = 1, \ldots, n_\varepsilon$. However, in certain situations it may be possible to recover some but not all of these shocks. This section extends our the results in Section 2 to cover situations in which, within a single model, some shocks may be recoverable while others may not.

The following example provides a simple illustration.

Example 3. Suppose that $n_y = n_\varepsilon = 3$ and that the spectral characteristic $\varphi(\lambda)$ is given by

$$\varphi(\lambda) = \begin{bmatrix} 1 & e^{i\lambda} & 0 \\ e^{-i\lambda} & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

This matrix function has a rank of 2 for almost all λ, since the first two columns are linearly dependent for any value of λ in $[-\pi, \pi]$. Given observations of $\{y_t\}$, it is not possible to disentangle $\{\varepsilon_{1,t}\}$ from $\{\varepsilon_{2,t}\}$, so according to Theorem (1), $\{\varepsilon_t\}$ is not recoverable. But, evidently, it is possible to recover $\{\varepsilon_{3,t}\}$, since $\varepsilon_{3,t} = y_{3,t}$ for all t. $\diamond$

What we would like is a necessary and sufficient condition that is capable of determining which, if any, of the scalar processes $\{\varepsilon_{k,t}\}$, $k = 1, \ldots, n_\varepsilon$ are recoverable. To find such a condition, we take advantage of the fact that recoverability is equivalent to error-free prediction. First, we find the best linear forecast of the values of the shock process $\{\varepsilon_t\}$ on the basis of the observables. This entails finding the projections $\tilde{\varepsilon}_{k,t}$ of the unobserved variables $\{\varepsilon_{k,t}\}$ on the subspace $\mathcal{H}(y)$. Then, we determine a set of conditions on the model in equation (3) which ensures that the errors in these forecasts are zero. This can happen if and only if $\varepsilon_{k,t}$ is an element of $\mathcal{H}(y)$; that is, if and only if $\{\varepsilon_{k,t}\}$ is recoverable from $\{y_t\}$.

In the course of solving the linear prediction problem, it is helpful for us to define the "pseudoinverse" of a matrix function.

Definition 3. Let $\varphi(\lambda)$ be an arbitrary $m \times n$ matrix function. Then its "pseudoinverse," $\varphi(\lambda)^\dagger$, is the $n \times m$ matrix function $\varphi(\lambda)^\dagger$ which satisfies the equations

$$\varphi(\lambda)^\dagger \varphi(\lambda)\varphi(\lambda)^* = \varphi(\lambda)^* \quad \text{and} \quad \varphi(\lambda)^\dagger(\varphi(\lambda)^\dagger)^*\varphi(\lambda)^* = \varphi(\lambda)^\dagger$$

for almost all λ.

The existence and uniqueness of this function follows immediately from the existence and the uniqueness of the ordinary matrix pseudoinverse; see, for example, Penrose (1955). The only difference is that uniqueness must be understood to mean that any other matrix function satisfying these equations is equal to $\varphi(\lambda)^\dagger$ almost everywhere. Analogously to the ordinary matrix case, the function $\varphi(\lambda)^\dagger$ has the following properties, which are understood to hold for almost all λ,

(i) $\varphi(\lambda)^\dagger \varphi(\lambda) = I_n$ if and only if $\text{rank}(\varphi(\lambda)) = n$

(ii) $\varphi(\lambda)^\dagger = \varphi(\lambda)^{-1}$ if and only if $\text{rank}(\varphi(\lambda)) = n = m$.

With this notation, we can state the following lemma.

Lemma 2 (Optimal Smoothing)**.** *Under Assumptions* (1) *and* (2)*, the stationary process* $\{\tilde{\varepsilon}_t\}$ *consisting of the best linear estimates of* $\{\varepsilon_t\}$ *on the basis of the values* $y_{k,s}$*,* $k = 1, \ldots, n_y$*,* $-\infty < s < \infty$*, is obtained from* $\{y_t\}$ *by a linear transformation of the form*

$$\tilde{\varepsilon}_t = \int e^{i\lambda t}\varphi(\lambda)^\dagger \Phi_y(d\lambda).$$

As in the case of Theorem (1), the logic behind this result can be understood by analogy with the static case. If $\varphi(\lambda) = \varphi$, where φ does not have full column rank, then the least-squares estimate, $\tilde{\varepsilon}_t$, of ε_t based on y_t is given by

$$\tilde{\varepsilon}_t = \varphi^{\dagger} y_t.$$

Returning to the dynamic case, we know that $\{\varepsilon_{k,t}\}$ is recoverable if and only if $\tilde{\varepsilon}_{k,t} = \varepsilon_{k,t}$ for all t. From the expression for $\tilde{\varepsilon}_t$ in Lemma (2), it is easy to see that this is true if and only if the k-th row of the product $\varphi(\lambda)^{\dagger}\varphi(\lambda)$ equals the k-th row of the n_ε dimensional identity matrix. This is the content of the next theorem.

Theorem 3 (Shock-Specific Recoverability)**.** *Under Assumptions* (1) *and* (2)*, the process* $\{\varepsilon_{k,t}\}$ *is recoverable from the observables* $\{y_t\}$ *if and only if*

$$\delta_k(I_{n_\varepsilon} - \varphi(\lambda)^{\dagger}\varphi(\lambda)) = 0$$

almost everywhere, where δ_k *denotes a* $1 \times n_\varepsilon$ *constant vector with components* $\delta_{kk} = 1$ *and* $\delta_{kl} = 0$ *for* $k \neq l$*.*

Remark 5. It can be helpful to consider how this result generalizes Theorem (1). If the condition in Theorem (3) is satisfied for all $k = 1, \ldots, n_\varepsilon$, then

$$\varphi(\lambda)^{\dagger}\varphi(\lambda) = I_{n_\varepsilon}$$

for almost all λ, which is equivalent to the requirement that $\varphi(\lambda)$ have full column rank almost everywhere.

Remark 6. The condition in the theorem can be easily checked numerically. We can draw a number λ_u randomly from its domain and compute the matrix $\varphi(\lambda_u)^{\dagger}$ numerically. An efficient way to do this is to use the singular value decomposition of $\varphi(\lambda_u)$, as Matlab does when given the command $\varphi(\lambda_u) = \texttt{pinv}(\varphi(\lambda_u))$. Then we can check which rows of the matrix

$$I_{n_\varepsilon} - \varphi(\lambda_u)^{\dagger}\varphi(\lambda_u)$$

are zero vectors. Those rows correspond to shocks that are recoverable. It is also possible to show that, in Matlab, an equivalent procedure is to execute the command

$$N = \texttt{null}(\varphi(\lambda_u)),$$

and check which rows of N are zero vectors. If this command returns an empty matrix, then $\varphi(\lambda_u)$ is full column rank, in which case all the shocks are recoverable.

Just as it is possible for the structural shocks to be partially recoverable, it is also possible for them to be partially invertible. Theorem (2) can be generalized in a straightforward way to cover this latter possibility. Namely, we can find the best linear forecast of the values of the shock process $\{\varepsilon_t\}$ only on the basis of the information contained in current and past observables. This involves finding the projections, $\hat{\varepsilon}_{k,t}$ of the unobserved variables $\varepsilon_{k,t}$ on the subspace $\mathcal{H}_t(y)$.[15]

Let us denote by $[\varphi(\lambda)]_+$ the matrix function

$$[\varphi(\lambda)]_+ = \sum_{s=0}^{\infty} \varphi_s e^{-i\lambda s}$$

for any matrix function $\varphi(\lambda)$ whose elements are square integrable, where $\{\varphi_s\}$ are the Fourier coefficients of $\varphi(\lambda)$. Then the solution to the linear filtering problem is given by the following theorem.

Lemma 3 (Optimal Filtering)**.** *The stationary process* $\{\hat{\varepsilon}_t\}$ *consisting of the best linear estimates of* $\{\varepsilon_t\}$ *on the basis of the values* $y_{k,s}$*,* $k = 1, \ldots, n_y$*,* $-\infty < s \leq t$*, is obtained from* $\{y_t\}$ *by a linear transformation of the form*

$$\hat{\varepsilon}_t = \int e^{i\lambda t} [\varphi(\lambda)^\dagger \gamma(\lambda)]_+ \gamma(\lambda)^\dagger \Phi_y(d\lambda),$$

where $\gamma(\lambda)$ *comes from some version of Wold's decomposition of* $\{y_t\}$*.*

Using this lemma, we arrive at the following theorem.

Theorem 4 (Shock-Specific Invertibility)**.** *Under Assumptions* (1) *and* (2)*, the process* $\{\varepsilon_{k,t}\}$ *is invertible from the observables* $\{y_t\}$ *if and only if*

$$\delta_k (I_{n_\varepsilon} - [\varphi(\lambda)^\dagger \gamma(\lambda)]_+ \gamma(\lambda)^\dagger \varphi(\lambda)) = 0$$

almost everywhere, where $\gamma(\lambda)$ *comes from some version of Wold's decomposition of* $\{y_t\}$*, and* δ_k *denotes a* $1 \times n_\varepsilon$ *constant vector with components* $\delta_{kk} = 1$ *and* $\delta_{kl} = 0$ *for* $k \neq l$*.*

Remark 7. It can be helpful to see how this result generalizes Theorem (2). If $\{\varepsilon_{k,t}\}$ is recoverable, then Lemma (2) implies that

$$\varepsilon_{k,t} = \int e^{i\lambda t} \delta_k \varphi(\lambda)^\dagger \gamma(\lambda) \Phi_w(d\lambda).$$

[15] This same basic approach was proposed by Sims and Zha (2006) and later used by Forni et al. (2017), but it has not been implemented at the same level of generality we consider here.

Comparing this with the expression for $\hat{\varepsilon}_t$ in Lemma (3), we can see that $\varepsilon_{k,t} = \hat{\varepsilon}_{k,t}$ if and only if $\delta_k[\varphi(\lambda)^\dagger\gamma(\lambda)]_+ = \delta_k\varphi(\lambda)^\dagger\gamma(\lambda)$ for almost all λ. Therefore, an alternative necessary and sufficient condition for the invertibility of $\{\varepsilon_{k,t}\}$ is that it is recoverable and

$$\frac{1}{2\pi}\int e^{i\lambda s}\delta_k\varphi(\lambda)^\dagger\gamma(\lambda)d\lambda = 0 \quad \text{for all } s < 0. \tag{9}$$

Of course, if this is true for all $k = 1, \ldots, n_\varepsilon$, then we again obtain the same condition as in Theorem (2).

Remark 8. In some situations, it may be desirable to have a measure of the "degree" of recoverability or invertibility. One natural measure of this is the variance of the prediction error $\tilde{\zeta}_{k,t} \equiv \varepsilon_{k,t} - \tilde{\varepsilon}_{k,t}$, or $\hat{\zeta}_{k,t} \equiv \varepsilon_{k,t} - \hat{\varepsilon}_{k,t}$, respectively. Using Lemmas (2) and (3), we have that

$$\begin{aligned}
\text{var}[\tilde{\zeta}_{k,t}] &= 1 - \frac{1}{2\pi}\int \delta_k\varphi(\lambda)^\dagger\varphi(\lambda)\delta_k^* d\lambda \\
\text{var}[\hat{\zeta}_{k,t}] &= 1 - \frac{1}{2\pi}\int \delta_k[\varphi(\lambda)^\dagger\gamma(\lambda)]_+\gamma(\lambda)^\dagger\varphi(\lambda)\varphi(\lambda)^*(\gamma(\lambda)^\dagger)^*[\varphi(\lambda)^\dagger\gamma(\lambda)]_+^*\delta_k^* d\lambda
\end{aligned}$$

Both of these numbers lie between zero and one, and values closer to zero represent higher degrees of recoverability or invertibility.[16] In the spirit of Sims and Zha (2006), Sims (2012), or Forni et al. (2017), it may be the case that "approximate" recoverability (i.e. values of $\text{var}[\tilde{\zeta}_{k,t}]$ numbers sufficiently close to zero) is enough to reach approximately correct economic conclusions using semi-structural methods. However, we do not explore this possibility further here.

4 Semi-Structural Analysis

This section describes the relationship between recoverability and semi-structural empirical analysis. We argue that recoverability is essential for performing this type of analysis, and explain why invertibility is not necessary. We then address some potential misconceptions regarding the implications of non-invertibility for semi-structural analysis. The section concludes by presenting a simple permanent-income example in which the shocks are recoverable but not invertible, and demonstrates how a semi-structural analysis might nevertheless be carried out.

[16] See also Plagborg-Møller and Wolf (2018), who instead subtract each of these variances from one, so that the number represents the explained degree of variation in the shock of interest.

4.1 Importance of Recoverability

The condition in Theorem (1) (or more generally, the condition in Theorem (3)) determines whether the structural shocks of interest can be recovered from the observables, given a complete knowledge of the structural model. That is, whether they can be recovered using the full set of restrictions embedded in the spectral characteristic $\varphi(\lambda)$. If so, then Lemma (2) provides an expression for the value of the shock(s) at each date as a function of the observables.

In an influential paper, Sims (1980) asks whether it is possible to recover the shocks of interest using only an appropriate *subset* of the theoretical restrictions implied by the model. If it is, then one's empirical conclusions can be interpreted as being robust across a range of different structural models that only need to agree on the relevant subset of theoretical restrictions. The motivation for this strategy was to combine the advantages of unrestricted econometric models and fully-specified dynamic equilibrium models, while minimizing the limitations of each. It has found wide acceptance in the macroeconomic literature, and we refer to it as the "semi-structural" approach.

In more detail, the semi-structural approach comprises two steps. The first "reduced-form" step involves using purely statistical methods to obtain an empirically adequate characterization of the spectral density (equivalently, the autocovariance function) of the observable process. The goal of this step is to summarize the data. The second "structural" step involves imposing restrictions, derived from economic theory, which are not falsifiable but are nevertheless sufficient to recover the structural shocks of interest. The goal of this step is to identify the shocks using only a minimal amount of economic theory, and then to use the results to entertain and test economic hypotheses involving additional restrictions, which are possibly less plausible from an a priori perspective.

From this brief description, it is easy to see that recoverability is a necessary condition for adopting a semi-structural approach. If the shocks cannot be recovered even with the full set of structural restrictions, then there can be no hope of doing so with only a subset of those restrictions. However, it is also easy to see that invertibility is not a necessary condition. Non-invertibility of the structural shocks only implies that the set of theoretical restrictions used in the structural step of the analysis should not include the restriction that the shocks be recoverable using information from past and present observables only. But of course this does not preclude the possibility that the shocks can be recovered from some combination of information from past, present, *and future* observables.

Indeed, it seems that many existing arguments for the importance of invertibility are actually arguments for the importance of recoverability. An especially clear example can be found in Sims and Zha (2006), who appear to equate these two concepts.

> "We would like to determine the extent to which the econometrician can recover from observations on [six economic] variables the underlying structural shocks, particularly those to monetary policy. We approach the problem in two steps. First, we examine "invertibility." That is, we ask to what extent the econometrician, if he knew the true parameters of the model, could construct the structural shocks from observations on this list of six variables…It should be clear that the second stage is pointless if the first stage concludes that the structural shocks cannot be recovered even with the full knowledge of the model coefficients." (p.243)

The difficulty with this line of reasoning is that invertibility is only sufficient for determining whether an econometrician can construct the structural shocks from his list of observables. As such, its failure does not provide reason to conclude that a semi-structural analysis of the shocks is pointless.

Of course, the fact that the structural shocks are recoverable does not, in itself, provide any direction regarding which specific subset of theoretical restrictions should be imposed in the structural step. There may be several different subsets that could be chosen. The choice of restrictions is ultimately an economic one, and will depend on the details of the structural models and particular questions under consideration. However, a shift in attention away from invertibility and towards the more general concept of recoverability can help to encourage the search for new economically interesting structural restrictions, some of which may have been previously considered out of bounds. In the next section we will propose one such set of restrictions, which we think is capable of providing new insights on an old question.

Once an appropriate set of theoretical restrictions has been chosen and used to recover the shocks, it is straightforward to compute many other objects of economic interest, such as impulse responses, variance shares, or shock decompositions. For example, the response of $y_{k,t+s}$ to a unit impulse in the shock $\varepsilon_{l,t}$ is

$$\mathrm{IR}_{kl}(s) = \frac{1}{2\pi}\int e^{i\lambda s}\varphi_{kl}(\lambda)d\lambda,$$

and the share of the variance in the process $\{y_{k,t}\}$ due to the shock process $\{\varepsilon_{l,t}\}$ over the frequency range $\mathit{\Delta} = [\lambda_1, \lambda_2]$ is

$$\mathrm{VS}_{kl}(\mathit{\Delta}) \;=\int_{\mathit{\Delta}} |\varphi_{kl}(\lambda)|^2 d\lambda \left(\int_{\mathit{\Delta}} f_{y,kk}(\lambda)d\lambda\right)^{-1}.$$

Lastly, the fluctuations in $\{y_{k,t}\}$ due only to the shock process $\{\varepsilon_{l,t}\}$ are given by the projections $\tilde{y}_{k,t}(l)$ of $y_{k,t}$ onto the subspace of $\mathcal{H}(\varepsilon)$ spanned by $\varepsilon_{l,t}$ for all t,

$$\tilde{y}_{k,t}(l) = \int e^{i\lambda t}\varphi_{kl}(\lambda)\Phi_{\varepsilon,l}(d\lambda).$$

4.2 Potential Misconceptions

In this subsection, we consider two questions that might give practitioners pause when applying semi-structural methods in situations with recoverable but non-invertible shocks. First: if the shocks are not invertible, doesn't that make it inappropriate to estimate a VAR? Second: if the shocks aren't invertible, won't we need stronger structural restrictions to identify the shocks? The answer to both questions is in the negative, but understanding why can be helpful for avoiding potential misconceptions about the nature of semi-structural analysis. In what follows we raise and briefly develop these questions before providing our responses.

First: if the structural shocks are not invertible, doesn't that make it inappropriate to estimate a VAR? After all, VARs are, by definition, invertible representations. If the shocks in a given structural model are not invertible, then it follows that the model does not admit a VAR representation. But then doesn't this make conclusions drawn from estimated VARs suspect at best, and inconsistent at worst?

In response to this question, the most important thing to realize is that the appropriateness of using VARs is really a statistical issue, not an economic one. The purpose of the reduced-form model is just to characterize the spectral density of observables, and it is statistical criteria for evaluating goodness of fit that can and should be used to guide this part of the analysis. For example, even if the structural shocks are invertible, it might be the case that a VAR provides an inferior fit compared to another statistical model. Furthermore, nothing about the semi-structural approach ties it to estimating a VAR. Common practice, following the application in Sims (1980), has been to estimate a VAR, but of course nothing about his proposal requires this. One could consider models with moving average terms as well as autoregressive terms, or one could even take an entirely non-parametric approach. But insofar as a VAR does provide an adequate empirical fit, it can be used for semi-structural analysis, regardless of whether the structural shocks are invertible.

Our empirical application in Section (5) will provide an example that uses a reduced-form VAR to analyze non-invertible shocks. But here we can provide a simpler example to help illustrate this point.

Example 4. Let $n_y = n_\varepsilon = 1$. In the reduced-form step, suppose that we determine that a first order autoregression does a good gob characterizing the spectral density of observables. The resulting spectral density estimate is

$$f_y(\lambda) = \frac{1}{2\pi} \frac{\sigma_u^2}{(1 - be^{-i\lambda})(1 - be^{i\lambda})},$$

where the autoregressive parameter $|b| < 1$ and $\sigma_u > 0$. However, suppose further that economic theory tells us that $E[y_t \varepsilon_t] \geq 0$ and $\mathcal{H}^t(\varepsilon) = \mathcal{H}^{t+1}(y)$ for all t, where $\mathcal{H}^t(\varepsilon)$ is the closed subspace spanned by the variables ε_s for $s \geq t$. This means that the shock at each date is recoverable from information in future observables, which implies that the shocks are not invertible. The unique reduced-form representation of $\{y_t\}$ consistent with these restrictions is

$$y_t = by_{t+1} + \sigma_u \varepsilon_{t-1},$$

which means that $\varepsilon_t = \sigma_u^{-1}(y_{t+1} - by_{t+2})$ for all t. Hence, imposing the restrictions from theory allows us to recover the structural shocks even though the spectral density was obtained using an invertible reduced-form model.[17]

It is also interesting to consider what light these observations shed on studies that seek to determine whether the structural model admits a VAR representation, such as Fernández-Villaverde et al. (2007) or Ravenna (2007). On the one hand, as we have already argued, if the structural model does not admit a VAR representation, it still

[17]Note that it is not possible to obtain the structural shocks in this example from a linear transformation of Wold shocks with a spectral characteristic that is a Blaschke matrix when viewed as a function of $z = e^{-i\lambda}$. This is because Blaschke matrices are analytic in the unit disk. Therefore, this representation falls outside the scope of, for example, Lippi and Reichlin (1994).

may be perfectly appropriate to estimate a VAR in the reduced-form step. On the other hand, even if the structural model does admit a VAR representation, it still may be more appropriate to use a different reduced-form model. For example, perhaps the structural model's implied VAR representation has so many lags (although possibly a finite number), so that in fact a low order mixed autoregressive moving average model achieves a superior statistical fit. This means that the question of whether a structural model admits a VAR representation isn't necessary or sufficient for a VAR to be an appropriate reduced-form model. Granted, these conditions might be helpful for selecting appropriate theoretical restrictions in the structural step. But otherwise, these conditions might be best viewed as sufficient conditions for recoverability.

Second: if the structural shocks are not invertible, won't we need stronger structural restrictions to identify the shocks? Normally, the structural step only requires us to impose restrictions that amount to a static rotation of the reduced-form shocks. That is, we look for a unitary matrix U such that

$$\varepsilon_t = U w_t,$$

where $\{w_t\}$ is the orthonormal shock process from some version of Wold's representation of $\{y_t\}$. However, if the shocks are non-invertible, then we are required to find the *dynamic* rotation that maps reduced-form shocks to structural shocks. That is, we need to find a unitary matrix function $U(\lambda)$ such that

$$\varepsilon_t = \int e^{i\lambda t} U(\lambda) \Phi_w(d\lambda).$$

Since the set of all dynamic rotations is much larger than the set of all static rotations, doesn't this mean that the conditions required to achieve identification in this case are substantially more demanding?

In response to this question, the most important thing to realize is that just because structural restrictions are *different*, that does not make them *stronger*. In general, it is difficult to precisely characterize what it means for one set of restrictions to be stronger than another, because the strength of a set of restrictions undoubtedly depends on the range of different plausible economic models that those restrictions exclude. But a rough criterion that is sufficient for our purposes here is to say that one set of restrictions is stronger than another if the latter is a proper subset of the former. And in this sense, all semi-structural identification schemes require restrictions of the same strength.

The spectral density of the observables defines a set of observationally equivalent reduced-form representations of the form

$$y_t = \int e^{i\lambda t}\rho(\lambda)\Phi_u(d\lambda),$$

where $\{u_t\}$ is an r_y dimensional process with orthonormal values and r_y is the rank of $f_y(\lambda)$ for almost all λ. This can be seen by the fact that each of these representations is in one-to-one correspondence with the $n_y \times r_y$ matrix function $\rho(\lambda)$, with rows in $\mathcal{L}^2(F_u)$, which satisfies the equation

$$f_y(\lambda) = \frac{1}{2\pi}\rho(\lambda)\rho(\lambda)^*.$$

If the structural shocks are recoverable, then the structural representation is a member of this set. And the role of the structural restrictions is to pick out that representation. Therefore, all structural restrictions require us to rule out every observationally equivalent reduced-form representation but one.

From this perspective, it is possible to see how the language of "static" versus "dynamic" rotations can easily be misleading. True, the structural shocks in a non-invertible representation are dynamic rotations of the shocks from Wold's decomposition, but they are also static rotations of the shocks from other observationally equivalent reduced-form representations. Conversely, if the structural shocks are static rotations of Wold shocks, they are also dynamic rotations of other reduced-form shocks. But what is so special about Wold's decomposition in this context? While crucially important for certain tasks (such as forecasting), the shocks in Wold's decomposition don't occupy a place of priority when it comes to recovering the structural shocks.

In the end, invertibility must be viewed as an economic restriction, not a statistical one. And it would seem to be an especially bad idea to impose it in situations when we have good economic reasons for thinking it isn't satisfied. Various papers since Hansen and Sargent (1991) have illustrated that it is not uncommon for this to be the case. However, the correct interpretation of these papers is not that they represent situations in which semi-structural analysis is inapplicable or would require considerably more theoretical input. Rather, they illustrate situations when invertibility should be replaced by something else. The precise form that appropriate alternative restrictions should take will vary depending on the context. But as long as the structural shocks are recoverable, we know that the structural model does contain within itself restrictions that are sufficient to identify the shocks. The main task then becomes deciding which of these restrictions to relax.

4.3 Simple Economic Example

To tie together the discussion in this section, we now consider a simple but familiar economic example to illustrate how semi-structural methods can be applied even when the structural model is not invertible. This example is borrowed from Fernández-Villaverde et al. (2007), who use it to illustrate a situation when their invertibility condition fails to hold. We will show that the structural shocks are not invertible with respect to the observables, but nevertheless that the shocks are recoverable. We then show how to perform a semi-structural analysis of the underlying shocks by imposing an appropriate subset of the structural model's economic restrictions.

Example 5. An econometrician tries to recover labor income shocks $\{\varepsilon_t\}$ from observations of changes in surplus income, $s_t \equiv z_t - c_t$, where c_t is date-t consumption and z_t is date-t labor income, which satisfies

$$z_t = \varepsilon_t, \quad \varepsilon_t \overset{\text{iid}}{\sim} \mathcal{N}(0,1).$$

The optimal path for consumption is a random walk

$$c_t = c_{t-1} + \left(1 - \frac{1}{R}\right)\varepsilon_t,$$

where $R > 1$ is the constant gross real interest rate.[18] Combining the previous two equations with the definition of surplus income, it follows that

$$y_t \equiv s_t - s_{t-1} = \frac{1}{R}\varepsilon_t - \varepsilon_{t-1}. \tag{10}$$

Therefore, the observable changes in surplus income, $\{y_t\}$, follow a first-order moving average process.

According to this model, the labor income shocks are not invertible. Heuristically, it is not possible to "solve" for ε_t as a function of previous values of y_t. Iterating backward T periods,

$$\begin{aligned}\varepsilon_t &= Ry_t + R\varepsilon_{t-1}\\ &= Ry_t + R^2 y_{t-1} + R^3 y_{t-3} + \cdots + R^T \varepsilon_{t-T}.\end{aligned}$$

[18] See Sargent (1987), Chapter XII for a presentation of this model.

But since economic theory tells us that $R > 1$, this solution explodes as $T \to \infty$.

Nevertheless, these shocks are recoverable. In particular, we can identify the shock at date t based on the information in subsequent observables.

$$\begin{aligned}\varepsilon_t &= -y_{t+1} + \frac{1}{R}\varepsilon_{t+1} \\ &= -y_{t+1} + \frac{1}{R}y_{t+2} + \left(\frac{1}{R}\right)^2 y_{t+3} + \cdots + \left(\frac{1}{R}\right)^{T-1} \varepsilon_{t+T} \\ &= -\sum_{s=0}^{\infty}\left(\frac{1}{R}\right)^s y_{t+s+1},\end{aligned}$$

where the last equality follows from taking limits as $T \to \infty$.

To demonstrate these two results more formally, consider the spectral characteristic linking the shocks to observables in equation (10),

$$\varphi(\lambda) = \frac{1}{R} - e^{-i\lambda}.$$

It is easy to see that $\varphi(\lambda)$ is full rank (i.e. nonzero) for all λ except $\lambda_0 = -\ln(1/R)/i$. Therefore, by Theorem (1), the shocks are recoverable.

To establish non-invertibility, notice that

$$y_t = w_t - \frac{1}{R}w_{t-1}$$

is a version of Wold's decomposition of $\{y_t\}$, which corresponds to the spectral characteristic $\gamma(\lambda) = 1 - R^{-1}e^{-i\lambda}$. Since $\varphi(\lambda)$ is nonzero for almost all λ, we have $\psi(\lambda) = \varphi(\lambda)^{-1}$. This implies that the Fourier coefficient of $\psi(\lambda)\gamma(\lambda)$ for $s = -1$ does not vanish, since, for $R > 1$,

$$\frac{1}{2\pi}\int_{-\pi}^{\pi} e^{-i\lambda}\left(\frac{R - e^{-i\lambda}}{1 - Re^{-i\lambda}}\right) d\lambda = \frac{1}{R^2} - 1.$$

Hence, by Theorem (2), the income shocks are not invertible.

Even though the shocks are not invertible, they are recoverable, and we can adopt a semi-structural approach. In the reduced-form step, we can use a statistical model to characterize the spectral density of $\{y_t\}$. Given $f_y(\lambda)$, we can impose the following economic restrictions in the structural step:[19]

(i) $\mathcal{H}^t(\varepsilon) = \mathcal{H}^{t+1}(y)$ for all t, and

[19] As in Example (4), $\mathcal{H}^t(\varepsilon)$ denotes the subspace of $\mathcal{H}(\varepsilon)$ spanned by ε_s for all $s \geq t$.

(ii) $E[y_t \varepsilon_t] \geq 0$,

These structural restrictions uniquely determine the spectral factor $\hat{\varphi}(\lambda)$ such that

$$f_y(\lambda) = \frac{1}{2\pi} |\hat{\varphi}(\lambda)|^2,$$

where the Fourier coefficients of $\hat{\varphi}(\lambda)$ vanish for all $s > 1$. The solution to this problem is

$$\hat{\varphi}(\lambda) = -e^{-i\lambda} \hat{\gamma}(\lambda)^*,$$

where $\hat{\gamma}(\lambda)$ is the spectral characteristic from the version of Wold's decomposition of $\{y_t\}$ with $E[y_t w_t] \geq 0$. Since the Fourier coefficients of the function $\hat{\gamma}(\lambda)$ vanish for $s < 0$, it follows that the Fourier coefficients of $\hat{\gamma}(\lambda)^*$ vanish for $s > 0$. Multiplying by a factor $-e^{-i\lambda}$ ensures that the Fourier coefficients of $\hat{\varphi}(\lambda)$ vanish for $s > 1$, and that the restriction $E[y_t \varepsilon_t] \geq 0$ is satisfied. If the structural model is correctly specified, then $\hat{\varphi}(\lambda) = \varphi(\lambda) = R^{-1} - e^{-i\lambda}$. ◇

The permanent-income example just discussed is a situation in which invertibility fails to hold because agents inside the model have more information at each date than the information contained in preceding values of the econometrician's observables. Their date-t information set is given by the subspace $\mathcal{H}_t(\varepsilon)$, while the information contained by preceding observables is $\mathcal{H}_t(y)$. When $R > 1$, we have shown that $\mathcal{H}_t(\varepsilon)$ is not contained in $\mathcal{H}_t(y)$. If the set of observables is expanded so that the agents' information set at each date coincides with the information contained in the preceding values of the observables, then of course the structural shocks would be invertible from past observables (the agents know their current income shocks). However, there are situations in which, even if the observable variables are the same ones that the agents themselves observe, it would still not be possible to identify the structural shocks using the information in current and past observables. Models with imperfect information about fundamentals are one example, and we consider these in more detail in the following section.

5 Productivity and the Business Cycle

How much of the business cycle is driven by changes in productivity? Initial calculations by real business cycle theorists during the 1980s suggest that productivity is largely (or even entirely) responsible for aggregate fluctuations in economic activity. However, a number of empirical studies published early in the following decade suggest that these initial calculations suffer from various limitations, and that evidence from VARs indicates that productivity's role is quite small. One way to rehabilitate the original real business cycle view would be to adopt Cochrane (1994)'s hypothesis that current economic fluctuations may partly depend on *future* changes in productivity. Perhaps earlier studies estimate a small role for productivity shocks because they all assume that economic agents do not have any information about future productivity beyond the history of current and past productivity realizations? He presents some preliminary evidence consistent with this hypothesis, but leaves further evaluation for future research.

The challenge in assessing this "news" hypothesis when estimating the importance of productivity is that it requires us to entertain the possibility that the economy's underlying structural shocks may not be invertible. To see why, imagine that in addition to observing current and past productivity, agents also receive partially informative signals about future productivity. At the time they make their decisions, agents are unable to determine whether a high signal realization reflects high future productivity or just unrelated noise. Because the agents themselves cannot disentangle future productivity from noise based on the current and past history of observables, an econometrician with the same information or less will not be able to do so either. But if it is impossible to disentangle future productivity from noise, then it is also impossible to determine exactly how much of the business-cycle variation in economic activity is actually due to productivity.

This line of reasoning, originally due to Blanchard et al. (2013), has led a number of researchers to conclude that semi-structural methods are simply incompatible with the hypothesis that agents receive imperfect signals about future fundamental shocks.[20] The usual suggestion is that to make progress the econometrician should adopt a fully-structural empirical approach. However, this conclusion crucially rests on the premise that invertibility is a necessary condition for using semi-structural methods; a premise that so far we have argued is not true. Therefore, in this section we present a new set of structural restrictions that are capable of determining the importance of productivity shocks even when agents partially respond to these shocks in advance.

[20] Indeed, this is the main methodological conclusion drawn by Blanchard et al. (2013). See also the literature reviews by Beaudry and Portier (2014) and Lorenzoni (2011).

We verify the effectiveness of these restrictions in a Monte Carlo study and then apply exactly the same empirical procedure to a sample of U.S. data.

5.1 Illustrative Model

We begin by describing a simple model of consumption determination from Blanchard et al. (2013) according to which current consumption partially depends on future productivity shocks. This model helps to illustrate the point that in order to determine the importance of productivity shocks for explaining aggregate fluctuations it is generally inappropriate to assume that the structural shocks are invertible. However, as we will show, the fact that the structural shocks are recoverable suggests that a different set of identifying restrictions (implied by the model) would be appropriate for performing a semi-structural analysis of the shocks.

The model stipulates that at each date, aggregate consumption is equal to agents' long-run forecast of total factor productivity,

$$c_t = \lim_{j \to \infty} E_t[a_{t+j}]. \tag{11}$$

This forecast is made conditional on the current and past history of productivity and signals about future productivity, a_τ and s_τ for $\tau \leq t$. Productivity is a random walk,

$$a_t = a_{t-1} + \sigma_a \varepsilon_t^a, \tag{12}$$

$\sigma_a > 0$, and the signal about future productivity is given by

$$s_t = \left(\frac{1-\rho}{1+\rho}\right) \sum_{j=-\infty}^{\infty} \rho^{|j|} a_{t-j} + v_t. \tag{13}$$

The parameter $0 < \rho < 1$ controls how much information the signal contains about future productivity. When $\rho = 0$, $s_t = a_t + v_t$, so the signal contains no additional information beyond a_t itself. The process $\{v_t\}$ represents non-productivity noise, and is assumed to follow a law of motion of the form

$$v_t = 2\rho v_{t-1} - \rho^2 v_{t-2} + \sigma_v \varepsilon_t^v - (\beta + \bar{\beta})\sigma_v \varepsilon_{t-1}^v + \beta\bar{\beta}\sigma_v \varepsilon_{t-2}^v, \tag{14}$$

with $\sigma_v > 0$. The vector of productivity and noise shocks, $\varepsilon_t = (\varepsilon_t^a, \varepsilon_t^v)'$, is independent and identically distributed over time with zero mean and identity covariance matrix. There is also a nonlinear restriction on the parameters σ_a, σ_v, ρ, and β, which ensures that $\{a_t\}$ can be written alternatively as the sum of a permanent component with first-order autoregressive dynamics in first differences, and a transitory component with first-order autoregressive dynamics in levels.[21]

Letting $y_t = (\Delta a_t, \Delta c_t)'$, this model implies that the observable process $\{y_t\}$ can be obtained from the structural shocks $\{\varepsilon_t\}$ by a linear transformation with spectral characteristic

$$\varphi(\lambda) = \begin{bmatrix} \sigma_a & 0 \\ \\ \dfrac{\omega(1-e^{-i\lambda}) - \sigma_a^2(1-\rho)}{\sigma_a(\rho - e^{-i\lambda})} & \dfrac{\sqrt{(\sigma_a^2-\omega)(\rho\sigma_a^2+\omega)}(1-e^{-i\lambda})}{\sigma_a(1-\rho e^{-i\lambda})} \end{bmatrix}, \tag{15}$$

where $\omega < \sigma_a^2$ is the covariance between the optimal one-step-ahead forecast errors of productivity and consumption. It is easy to see that this matrix function is full column rank for almost all λ, and therefore, by Theorem (1), that the structural shocks are recoverable.

However, consistent with the intuition described above, these shocks are not invertible. To verify this, we first compute the version of Wold's decomposition of $\{y_t\}$ according to which $\mathcal{H}_t(y) = \mathcal{H}_t(w)$ for all t, $E[y_{k,t}w_{lt}] = 0$ for all $k < l$, and $E[y_{k,t}\varepsilon_{k,t}] \geq 0$ for all k. The corresponding spectral characteristic is

$$\gamma(\lambda) = \begin{bmatrix} \dfrac{\omega(1-\rho)+\rho\sigma_a^2(1-e^{-i\lambda})}{(1-\rho e^{-i\lambda})\sqrt{\rho\sigma_a^2+(1-\rho)\omega}} & \dfrac{e^{-i\lambda}(1-\rho)\sqrt{(\sigma_a^2-\omega)(\rho\sigma_a^2+\omega)}}{(1-\rho e^{-i\lambda})\sqrt{\rho\sigma_a^2+(1-\rho)\omega}} \\ \\ \dfrac{\omega}{\sqrt{\rho\sigma_a^2+(1-\rho)\omega}} & \dfrac{\sqrt{(\sigma_a^2-\omega)(\rho\sigma_a^2+\omega)}}{\sqrt{\rho\sigma_a^2+(1-\rho)\omega}} \end{bmatrix} \tag{16}$$

To apply Theorem (2), we need to analyze the properties of the matrix function $\alpha(\lambda) \equiv \psi(\lambda)\gamma(\lambda)$, where $\psi(\lambda)$ satisfies $\psi(\lambda)\varphi(\lambda) = I_2$ for almost all λ. Since $\varphi(\lambda)$ is square and has full rank for almost all λ, it follows that $\psi(\lambda) = \varphi(\lambda)^{-1}$. Therefore, we can combine equations (15) and (16) to obtain a closed-form expression for $\alpha(\lambda)$. For instance, the lower-left element of $\alpha(\lambda)$ reduces to

$$\alpha_{21}(\lambda) = \frac{(1-\rho)\sqrt{(\sigma_a^2-\omega)(\rho\sigma_a^2+\omega)}}{\sigma_a(\rho-e^{-i\lambda})\sqrt{\rho\sigma_a^2+(1-\rho)\omega}}.$$

[21] Blanchard et al. (2013) write the information structure in this alternative way, which is observationally equivalent to the one presented here. While the shocks in their formulation are not recoverable, they are also inappropriate for answering our economic question in this section. For more details on the link between their representation and the one presented above, see Chahrour and Jurado (2018).

However, the fact that

$$\frac{1}{2\pi}\int_{-\pi}^{\pi} e^{-i\lambda}\alpha_{21}(\lambda)d\lambda = -\frac{(1-\rho)\sqrt{(\sigma_a^2-\omega)(\rho\sigma_a^2+\omega)}}{\sqrt{\rho\sigma_a^2+(1-\rho)\omega}} \neq 0$$

implies that the Fourier coefficients of $\alpha(\lambda)$ do not vanish for all negative values of s, and therefore, by Theorem (1), that the structural shock process $\{\varepsilon_t\}$ is not invertible.

5.2 Structural Restrictions

While the shocks in the preceding model are not invertible, the fact that they are recoverable suggests that it may be possible to find an alternative set of structural restrictions that would be appropriate for use in a semi-structural analysis. One restriction, which arises naturally from the theory, is that productivity shocks are responsible for all of the variation in productivity. This, in turn, implies that non-productivity noise shocks must be orthogonal to productivity at all leads and lags, a restriction which is evidenced by the fact that the upper-right element of $\varphi(\lambda)$ is zero, as can be seen in equation (15).

To state our proposed restrictions more formally, we require that

(i) $\mathcal{H}_t(\varepsilon^a) = \mathcal{H}_t(\Delta a)$ for all t with $E[\Delta a_t \varepsilon_t^a] \geq 0$, and

(ii) $\mathcal{H}_t(\varepsilon^v) = \mathcal{H}_t(\zeta)$ for all t with $E[\zeta_t \varepsilon_t^v] \geq 0$,

where ζ_t is the residual from an orthogonal projection of Δc_t onto the space $\mathcal{H}(\Delta a)$. The first restriction says that the productivity shock comes from the Wold decomposition of $\{\Delta a_t\}$ with $E[\Delta a_t \varepsilon_t^a] \geq 0$. This implies that productivity shocks explain all of the variation in productivity. The second says that the noise shocks capture the fluctuations in current consumption growth that are orthogonal to productivity growth across all time periods.[22]

To demonstrate how these restrictions can be used as part of a semi-structural analysis, let us suppose that an econometrician has some characterization of the spectral density of the observable process $\{y_t\}$, obtained using purely statistical methods

[22]These restrictions might appear to coincide with the method proposed by Forni et al. (2017a,b); in fact, they do not. For example, their method does not correctly identify the shocks in the model from this section. However, our restrictions also work in their model.

(e.g. by first estimating a VAR). The structural step then involves obtaining a factorization of the spectral density of the form

$$f_y(\lambda) = \frac{1}{2\pi}\hat{\varphi}(\lambda)\hat{\varphi}(\lambda)^*, \tag{17}$$

where $\hat{\varphi}(\lambda)$ is uniquely determined by our theoretical restrictions. In this case, the restrictions require $\hat{\varphi}(\lambda)$ to have a lower-triangular form

$$\hat{\varphi}(\lambda) = \begin{bmatrix} \hat{\varphi}_{11}(\lambda) & 0 \\ \hat{\varphi}_{21}(\lambda) & \hat{\varphi}_{22}(\lambda) \end{bmatrix}, \tag{18}$$

where the Fourier coefficients of $\hat{\varphi}_{11}(\lambda)$ and $\hat{\varphi}_{22}(\lambda)$ vanish for all $s < 0$. Using the Fourier series expansion of $\hat{\varphi}(\lambda)$, these restrictions imply a structural moving average representation of the form

$$\begin{bmatrix} \Delta a_t \\ \Delta c_t \end{bmatrix} = \cdots + \underbrace{\begin{bmatrix} 0 & 0 \\ * & 0 \end{bmatrix}}_{\hat{\varphi}_{-1}} \begin{bmatrix} \varepsilon^a_{t+1} \\ \varepsilon^v_{t+1} \end{bmatrix} + \underbrace{\begin{bmatrix} * & 0 \\ * & * \end{bmatrix}}_{\hat{\varphi}_0} \begin{bmatrix} \varepsilon^a_t \\ \varepsilon^v_t \end{bmatrix} + \underbrace{\begin{bmatrix} * & 0 \\ * & * \end{bmatrix}}_{\hat{\varphi}_1} \begin{bmatrix} \varepsilon^a_{t-1} \\ \varepsilon^v_{t-1} \end{bmatrix} + \cdots,$$

where $\{\hat{\varphi}_s\}$ is the sequence of Fourier coefficients associated with $\hat{\varphi}(\lambda)$.[23]

Before presenting the solution to this spectral factorization problem, it is worth making two observations regarding the nature of these restrictions in relation to the original restrictions used by Sims (1980) in his empirical application. First, they also represent a form of "triangularization;" indeed, they amount to a dynamic generalization of the Cholesky factorization of a positive definite matrix. Second, relative to Sims' original restrictions, these trade one zero for another. They impose a zero-restriction on the upper-right entry of each coefficient matrix associated with a vector of past shocks, while at the same time relaxing a zero restriction on the lower-left entry of each coefficient matrix associated with a vector of future shocks. This makes especially transparent the fact that neither set of restrictions is any weaker than the other, in the sense of being a proper subset.

To solve our factorization problem, we begin by writing out equation (17) more explicitly, using equation (18), as

$$\begin{bmatrix} f_{\Delta a}(\lambda) & f_{\Delta a \Delta c}(\lambda) \\ f_{\Delta c \Delta a}(\lambda) & f_{\Delta c}(\lambda) \end{bmatrix} = \frac{1}{2\pi} \begin{bmatrix} |\hat{\varphi}_{11}(\lambda)|^2 & \hat{\varphi}_{11}(\lambda)\overline{\hat{\varphi}_{21}(\lambda)} \\ \hat{\varphi}_{11}(\lambda)\hat{\varphi}_{21}(\lambda) & |\hat{\varphi}_{22}(\lambda)|^2 + |\hat{\varphi}_{21}(\lambda)|^2 \end{bmatrix}. \tag{19}$$

[23] As in Example 4, in this case it is not possible to obtain the structural shocks in this representation from Wold shocks using a Blaschke matrix.

Restriction (i) says that $\hat{\varphi}_{11}(\lambda)$ can be computed from Wold's decomposition of $\{\Delta a_t\}$; this is unique and can be obtained in the usual way. The lower-left equation in (19) uniquely determines $\hat{\varphi}_{21}(\lambda)$ as a function of $f_{\Delta c\Delta a}(\lambda)$ and $\hat{\varphi}_{11}(\lambda)$, the first of which is known from data and the second of which has already been determined from the upper-left equation. The lower-right equation in (19) implies that

$$|\hat{\varphi}_{22}(\lambda)|^2 = 2\pi f_{\Delta c}(\lambda) - |\hat{\varphi}_{21}(\lambda)|^2$$

Together with restriction (ii), this means that $\hat{\varphi}_{22}(\lambda)$ is uniquely determined from Wold's decomposition of a process with spectral density $2\pi f_{\Delta c}(\lambda) - |\hat{\varphi}_{21}(\lambda)|^2$. Therefore, we have shown both that the factor $\hat{\varphi}(\lambda)$ is unique, and how to obtain it.

In the context of the model from the preceding subsection, it is straightforward to show that these restrictions correctly identify the structural representation. That is, the value of $\hat{\varphi}(\lambda)$ computed from the spectral density

$$f_y(\lambda) = \frac{1}{2\pi}\begin{bmatrix} \sigma_a^2 & \dfrac{\omega(1-e^{-i\lambda}) - \sigma_a^2(1-\rho)e^{-i\lambda}}{1-\rho e^{-i\lambda}} \\ \dfrac{\omega(1-e^{-i\lambda}) - \sigma_a^2(1-\rho)}{\rho - e^{-i\lambda}} & \sigma_a^2 \end{bmatrix}$$

according to the steps just described coincides with the value of $\varphi(\lambda)$ reported in equation (15).

5.3 Monte Carlo Study

To demonstrate how semi-structural methods can be applied in practice to models with partially informative signals about future fundamentals, we perform a Monte Carlo exercise using the illustrative model from this section. The exercise involves simulating data on consumption and productivity from the model, and placing ourselves in the position of an econometrician who has no knowledge of the true data generating process. He receives a finite sample of realizations, and is charged with estimating the importance of productivity shocks for consumption from that sample. To do so, he imposes only the structural restrictions (i) and (ii) from the previous subsection.

In practice, we simulate $N = 1000$ samples of $T = 276$ observations of consumption and productivity from the model. We use the following parameter values,

$$\rho = 0.8910, \quad \sigma_a = 0.6700, \quad \text{and} \quad \omega = 0.2258,$$

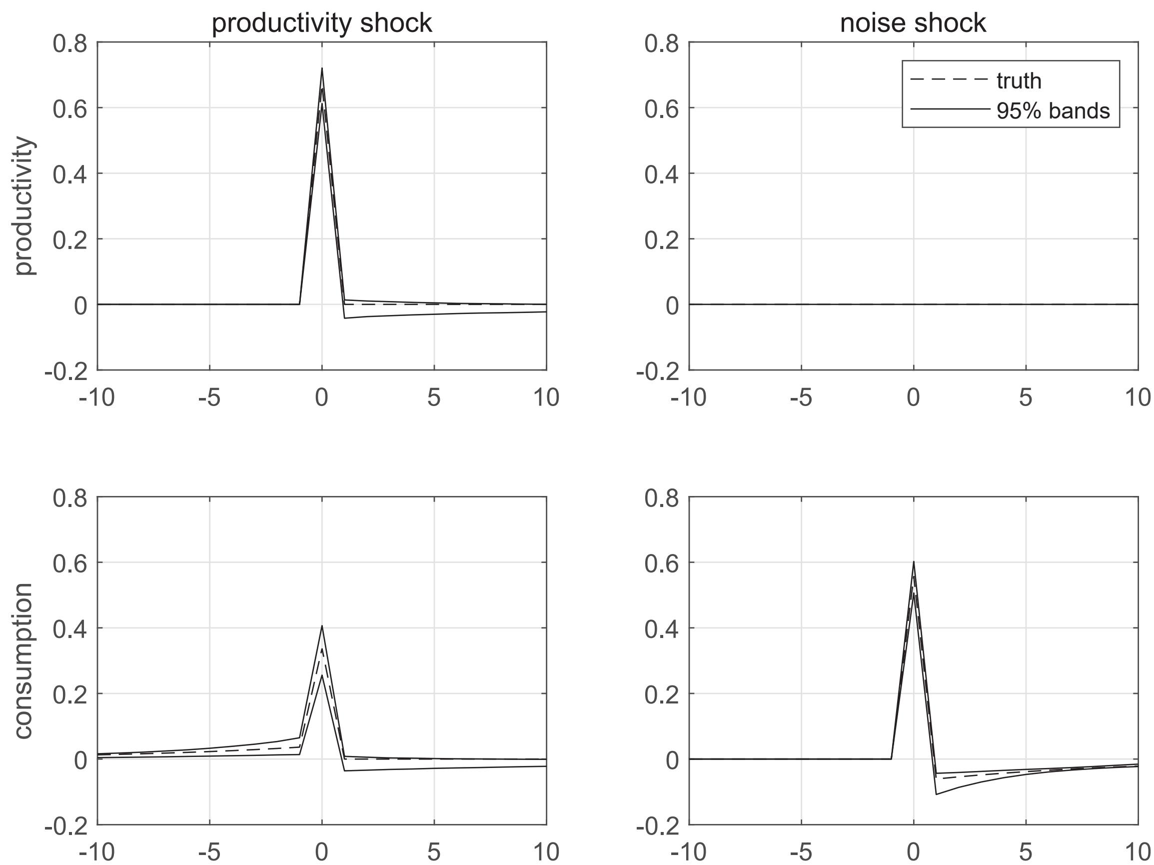

Figure 1: Impulse responses from simulated data. These figures display the responses of productivity growth and consumption growth to a one standard deviation impulse in the productivity shock and noise shock in period 0. The dashed lines are the true impulse response coefficients, and the solid lines are the 95% bands from the distribution of point estimates across the $N = 1000$ samples of length $T = 276$.

which correspond to the same parameter values chosen by Blanchard et al. (2013). The reduced-form model is an unrestricted VAR, which we fit to the data using the multivariate algorithm of Morf et al. (1978). The lag length is chosen to minimize the information criterion proposed in Hannan and Quinn (1979).

The dashed lines in Figure (1) display the true impulse response coefficients of productivity and consumption growth to productivity and noise shocks (i.e. the Fourier coefficients of the four functions in equation (15)). In the figure, period 0 represents the period in which the productivity shock is realized, affecting contemporaneous productivity growth. Negative values on the horizontal axis capture the effects of anticipation in the periods leading up to the period 0 shock, while positive values on the horizontal axis capture effects of the shock in the periods after it has taken place.

According to the model, productivity is a random walk, which can be seen in the upper-left panel by the fact that productivity growth only responds to contemporaneously to the productivity shock. Productivity is unrelated to noise shocks across all horizons, which is why all the coefficients in the upper-right panel are zero. The lower-left panel indicates that, while consumption growth does not respond to past productivity shocks, it does respond during the period of anticipation, through agents' reliance on their informative signal. According to the lower-right panel, consumption growth also increases contemporaneously with a noise shock, but decreases in response to noise shocks in the past. This is because as time passes, agents are able to learn that changes in their signal were due to noise and not an actual future changes in productivity. In response to noise shocks far enough in the past, consumption growth doesn't respond at all.

The solid lines in Figure (1) display 95% bands constructed from the point estimates across the N different simulated samples (i.e. the Fourier coefficients of the four functions in equation (18)). The places in the figure where the bands collapse to a single line reflect the zero restrictions imposed by our identification scheme: productivity doesn't respond to future productivity shocks or noise shocks, and consumption doesn't respond to future noise shocks. In all cases, the semi-structural VAR-based identification procedure does a good job of identifying these coefficients.

While informative about the effectiveness of the VAR-based procedure, the impulse response coefficients from Figure (1) are not strictly relevant to the question posed at the beginning of this section. To answer that question, we consider the share of the variance in consumption that is explained by productivity shocks over business cycle frequencies (6 to 32 quarters). This is reported in Figure (2). The vertical dashed line is the true productivity share (0.31), while the solid line is the distribution of point estimates across the N different samples. Again, the VAR-based procedure evidently delivers accurate estimates of the importance of productivity shocks. Based on the distribution of point estimates, it appears that the estimated importance of productivity shocks does exhibit some slight upward bias in samples of this size. However, a slight upward bias in this estimate would only strengthen the conclusions we reach in our empirical application.

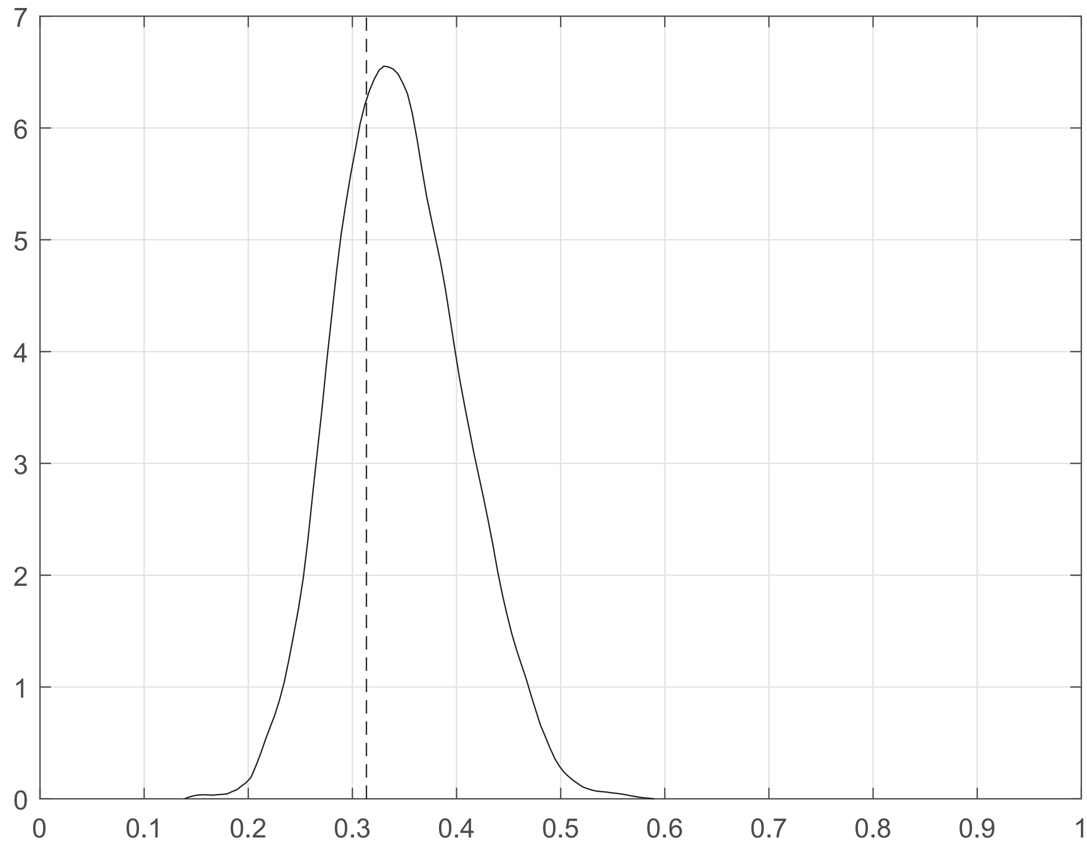

Figure 2: Variance share estimated from simulated data. This figure plots the share of the variance in consumption due to productivity shocks over business-cycle frequencies (6 to 32 quarters). The dashed line is true productivity share (0.31) and the solid line is the distribution of point estimates across the $N = 1000$ samples of length $T = 276$.

5.4 Application to U.S. Data

In this subsection, we apply the same semi-structural procedure used in our Monte Carlo exercise to actual U.S. consumption and productivity data. We measure consumption by the natural logarithm of real per-capita personal consumption expenditure (NIPA table 1.1.6, line 2, divided by BLS series LNU00000000Q) and productivity by the natural logarithm of utilization-adjusted total factor productivity (Basu et al., 2006). Our sample is 1948:Q1 to 2016:Q4, which gives $T = 276$ observations.

Before discussing the results, a cautionary remark is in order regarding the interpretation of noise shocks in actual data. In the illustrative model used in the Monte Carlo exercise, productivity is the only fundamental process, and agents have

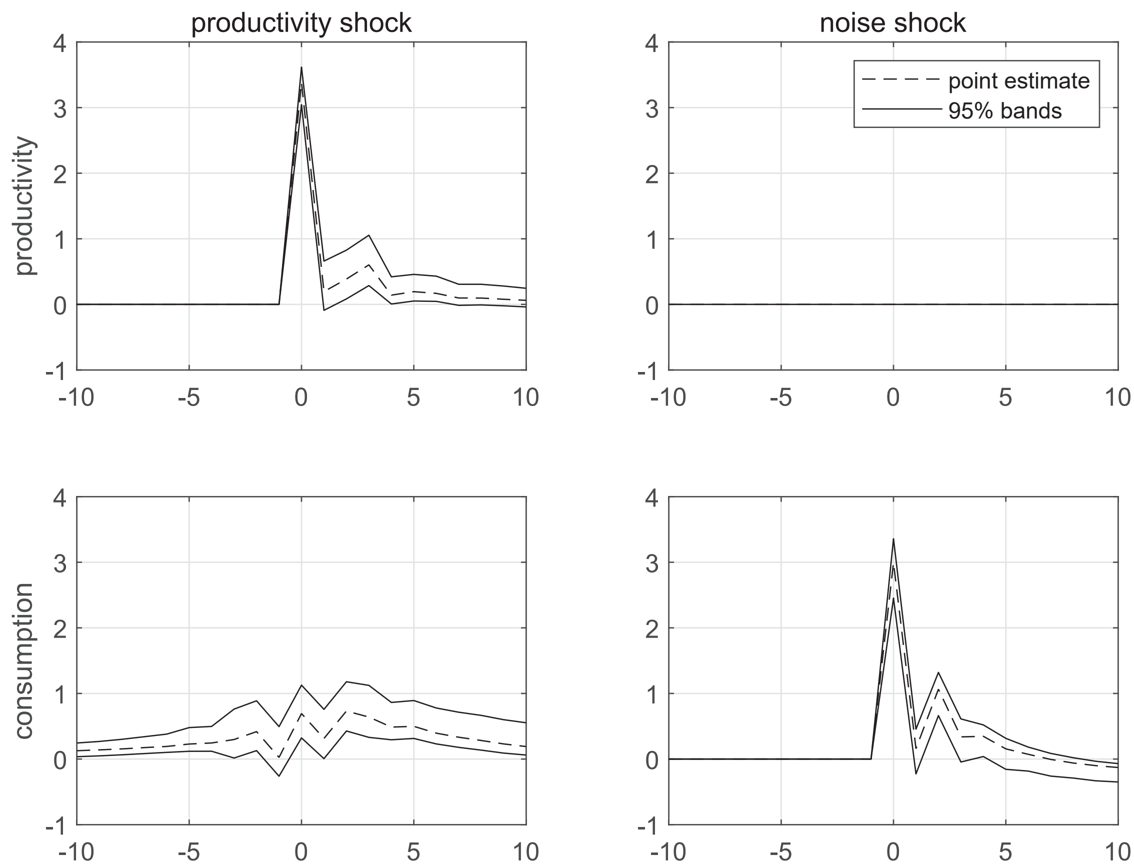

Figure 3: Impulse responses estimated from U.S. data. These figures display the responses of productivity growth and consumption growth to a one standard deviation impulse in the productivity shock and noise shock in period 0. The dashed lines are the point estimates (annualized percentage points), and the solid lines are the 95% bands from the distribution of $N = 1000$ bootstrap estimates. The sample period is 1948:Q1-2016:Q4.

rational expectations. As a result, the only reason that consumption can possibly move without some corresponding movement in current, past, or future productivity is because of rational errors induced by noisy signals. In the data, it is possible that consumption is driven by fundamentals other than productivity, by sunspots, or even by non-rational fluctuations in people's beliefs. Of course, this does not affect the interpretation of productivity shocks. However, it does mean that noise shocks should be interpreted broadly in this subsection as composite shocks that capture all *non-productivity* fluctuations in consumption.

Keeping that interpretation in mind, we turn to Figure (3), which displays the estimated impulse response coefficients of productivity and consumption growth to productivity and noise shocks. The figure show that the contemporaneous response of consumption to a one-standard-deviation productivity shock is 0.7 percent, consistent with the idea that positive productivity shocks are expansionary. Perhaps more interestingly, consumption growth also increases in the period of anticipation leading up to the shock. In particular, the fact that the response of current consumption to future productivity shocks is statistically different from zero is consistent with the idea that agents do observe information about future productivity above and beyond the history of aggregate productivity alone. However, the fact that responses are larger after period 0 also reveals that consumption growth is more sensitive to current and past productivity shocks than to future productivity shocks.

By contrast, the contemporaneous response of consumption growth to a noise shock is 3 percent, which is larger than the response of consumption growth to a productivity shock at any horizon. As in the model, consumption growth responds negatively to noise after several periods, consistent with the idea that over time agents can disentangle productivity and noise shocks. However, unlike the model, the consumption response is only negative for noise shocks that occurred more than six quarters in the past. To the extent that these shocks represent rational mistakes due to imperfect signals, this means that it takes longer for agents to recognize their mistakes in the data than in the model.

Turning to the importance of productivity shocks for consumption, Figure (4) plots the share of the variance in consumption explained by productivity shocks over business-cycle frequencies. The vertical dashed line is the point estimate, and the solid line is the distribution of point estimates across all bootstrap samples. We find that 14% of the business-cycle variation in consumption can be attributed to productivity shocks. This means that, even allowing for the possibility that consumption can partly depend on future productivity, productivity shocks are still not capable of explaining the bulk of the business-cycle variation in consumption.

An interesting observation regarding these estimates is that they are not sensitive to the fact that they are based on a bi-variate VAR. Often when using VARs, the choice of variables is a critical input into the analysis; results can sometimes change in economically significant ways by including one additional variable.[24] The reason for

[24]As in the case of the well-known "price puzzle" in semi-structural analyses of monetary shocks.

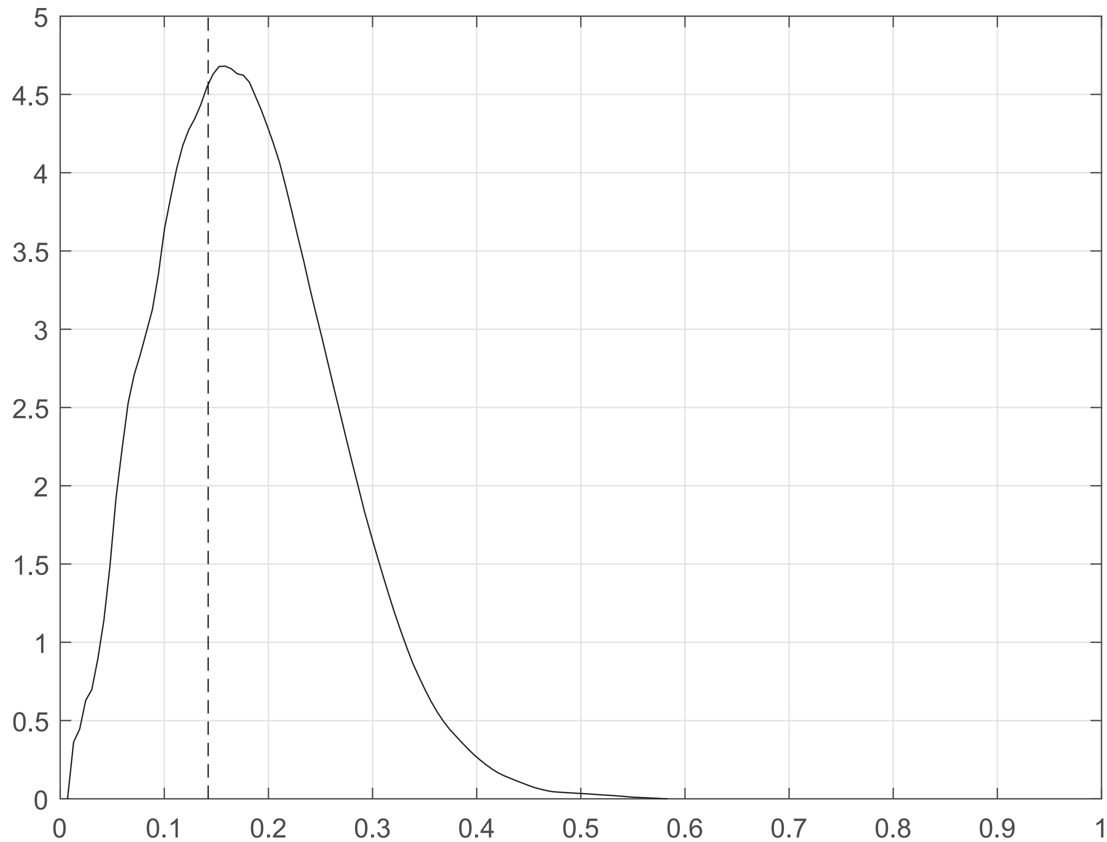

Figure 4: Variance share estimated from U.S. data. This figure plots the share of the variance in consumption due to productivity shocks over business-cycle frequencies (6 to 32 quarters). The dashed line is the point estimate (0.14) and the solid line is the distribution of estimates across the $N = 1000$ bootstrap samples.

this is that most existing VAR studies (unnecessarily?) require identification of the exact information set available to economic agents at each date. Understandably, this is a challenging task, and the concern that perhaps some important series observed by agents has been omitted always tends to loom in the background. This is the reason why one strand of the literature has sought to incorporate larger amounts of information into low-dimensional vectors by using data-compression methods such as principal component analysis.

However, for our purposes it is not necessary to assume that the current and past history of consumption and productivity fully summarizes agents' information set at each date — most likely it does not. Whatever *consumption-relevant* information agents do have about productivity is revealed by their observable consumption

Consumption	GDP	Investment	Employment
14	8	6	9
(4, 35)	(5, 24)	(4, 24)	(4, 24)

Table 1: Productivity shares for other macroeconomic aggregates. This table displays the share of the variance in each series due to productivity shocks over business-cycle frequencies (6 to 32 quarters). The first row contains the point estimates, and the second row contains 95% bootstrap confidence intervals.

choices. In other words, it is less important for us to identify exactly what signals agents observe than it is to observe how the information in those signals is ultimately incorporated into their actions. All that matters for our purposes is the joint bi-variate dynamics of consumption and productivity. This implies that all the estimates in Figures (3) and (4) are robust to the inclusion of any number of additional variables to the analysis.

While a bi-variate analysis is sufficient to compute the importance of productivity shocks for consumption, it is nevertheless natural to ask whether a similar result holds in the case of macroeconomic aggregates other than consumption. To answer this question, we repeat the preceding analysis after replacing consumption with each of the following variables: output, investment, and employment.[25] Table (1) displays the share of the variance in each series explained by productivity shocks over business-cycle frequencies. The table shows that productivity shocks can explain only a small portion of the variance of any of these variables. In fact, the 14% share explained for consumption is the largest of the group.[26]

None of this evidence is per se inconsistent with the hypothesis that business cycles are mainly (or even entirely, as in the illustrative model) driven by fluctuations in agents' *beliefs* about productivity. In order to evaluate such a hypothesis, it would

[25] We measure output by gross domestic product (NIPA table 1.1.6, line 1), investment by real gross private domestic investment (NIPA table 1.1.6, line 7), and employment by the natural logarithm of per-capita employment in the nonfarm business sector (BLS series CES0000000001). We take the natural logarithm of all variables, after expressing them in per-capita terms by dividing by the civilian noninstitutional population (divided by BLS series LNU00000000Q).

[26] Of course, our results depend on our chosen measure of productivity, which is utilization-adjusted. Using raw Solow residuals (with no adjustment), we estimate much larger shares.

be important to measure exactly what beliefs agents hold about productivity at each date. And in that case, as we have seen, the precise nature of the other variables included in the analysis would become relevant. However, even without undertaking such a project, the results from this section are helpful enough to press us into the following dilemma: either business-cycle fluctuations are not largely driven by fluctuations in beliefs about productivity, or else they are, but somehow agents' beliefs about productivity tend to fluctuate for reasons that are largely unrelated to actual changes in productivity.

6 Conclusion

At least since Hansen and Sargent (1991), economists have been keenly aware of the difficulties that non-invertible models pose for semi-structural methods of the type originally proposed by Sims (1980). Our purpose has been to argue that, at least from an econometric perspective, these difficulties don't necessarily represent difficulties at all. Nothing about the nature of the semi-structural approach requires either one's reduced-form model or one's structural model to be invertible.

Instead, we have argued that what is needed is the much weaker condition that the structural shocks be recoverable from observables. We have presented a simple necessary and sufficient condition that can be used to check for recoverability. We have also presented similar conditions for invertibility. Hopefully by clarifying the difference between invertibility and recoverability, and shifting attention to the later, our results will allow semi-structural methods to find greater applicability across a wider range of economic contexts.

References

Barsky, R. B. and E. R. Sims (2012). Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence. *American Economic Review 102*(4), 1343–1377.

Basu, S., J. G. Fernald, and M. S. Kimball (2006). Are Technology Improvements Contractionary? *The American Economic Review 96*(5), 1418–1448.

Beaudry, P. and F. Portier (2014). News-Driven Business Cycles: Insights and Challenges. *Journal of Economic Literature 52*(4), 993–1074.

Blanchard, O. J., J.-P. L'Huillier, and G. Lorenzoni (2013). News, Noise, and Fluctuations: An Empirical Exploration. *American Economic Review 103*(7), 3045–3070.

Brockwell, P. J. and R. A. Davis (1991). *Time Series: Theory and Methods.* New York: Springer.

Chahrour, R. and K. Jurado (2018). News or noise? the missing link. *American Economic Review 108*(7), 1702–1736.

Cochrane, J. H. (1994). Shocks. *Carnegie-Rochester Conference Series on Public Policy 41*(1), 295–364.

Cochrane, J. H. (1998). What do the VARs mean? Measuring the output effects of monetary policy. *Journal of Monetary Economics 41*(2), 277–300.

Dupor, B. and J. Han (2011). Handling Non-Invertibility: Theory and Applications. Working paper, Ohio State University.

Fernández-Villaverde, J., J. F. Rubio-Ramírez, T. J. Sargent, and M. W. Watson (2007). ABCs (and Ds) of Understanding VARs. *American Economic Review 97*(3), 1021–1026.

Forni, M., L. Gambetti, M. Lippi, and L. Sala (2017a). Noise Bubbles. *The Economic Journal 127*(604), 1940–1976.

Forni, M., L. Gambetti, M. Lippi, and L. Sala (2017b). Noisy News in Business Cycles. *American Economic Journal: Macroeconomics 9*(4), 122–52.

Forni, M., L. Gambetti, and L. Sala (2017). Reassessing structural VARs: beyond the ABCs (and Ds). Working paper, Bocconi University.

Futia, C. A. (1981). Rational Expectations in Stationary Linear Models. *Econometrica 49*(1), 171–192.

Hannan, E. J. and B. G. Quinn (1979). The Determination of the Order of an Autoregression. *Journal of the Royal Statistical Society. Series B (Methodological) 41*(2), 190–195.

Hansen, L. P. and T. J. Sargent (1980). Formulating and estimating dynamic linear rational expectations models. *Journal of Economic Dynamics and Control 2*(1), 7–46.

Hansen, L. P. and T. J. Sargent (1991). Two Difficulties in Interpreting Vector Autoregressions. In L. P. Hansen and T. J. Sargent (Eds.), *Rational Expectations Econometrics*, pp. 77–119. Boulder, CO: Westview Press.

Kolmogorov, A. N. (1941). Stationary Sequences in Hilbert Space. *Moscow University Mathematics Bulletin 2*(6), 1–40.

Leeper, E. M., T. B. Walker, and S.-C. S. Yang (2013). Fiscal Foresight and Information Flows. *Econometrica 81*(3), 1115–1145.

Lippi, M. and L. Reichlin (1993). The Dynamic Effects of Aggregate Demand and Supply Disturbances: Comment. *The American Economic Review 83*(3), 644–652.

Lippi, M. and L. Reichlin (1994). VAR analysis, nonfundamental representations, blaschke matrices. *Journal of Econometrics 63*(1), 307–325.

Lorenzoni, G. (2011). News and Aggregate Demand Shocks. *Annual Review of Economics 3*(1), 537–557.

Mertens, K. and M. O. Ravn (2010). Measuring the Impact of Fiscal Policy in the Face of Anticipation: A Structural VAR Approach. *The Economic Journal 120*(544), 393–413.

Morf, M., A. Vieira, D. T. L. Lee, and T. Kailath (1978). Recursive Multichannel Maximum Entropy Spectral Estimation. *IEEE Transactions on Geoscience Electronics 16*(2), 85–94.

Penrose, R. (1955). A generalized inverse for matrices. *Mathematical Proceedings of the Cambridge Philosophical Society 51*(3), 406413.

Plagborg-Møller, M. (2017). Bayesian Inference on Structural Impulse Response Functions. Working paper, Princeton University.

Plagborg-Møller, M. and C. K. Wolf (2018). Instrumental Variable Identification of Dynamic Variance Decompositions. Working paper, Princeton University.

Quah, D. (1990). Permanent and Transitory Movements in Labor Income: An Explanation for "Excess Smoothness" in Consumption. *Journal of Political Economy 98*(3), 449–475.

Ravenna, F. (2007). Vector autoregressions and reduced form representations of dsge models. *Journal of Monetary Economics 54*(7), 2048 – 2064.

Rosenberg, M. (1969). Mutual subordination of multivariate stationary processes over any locally compact abelian group. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete 12*(4), 333–343.

Rozanov, Y. A. (1967). *Stationary Random Processes.* San Francisco, CA: Holden-Day.

Sargent, T. J. (1987). *Macroeconomic Theory: Second Edition.* Bingley: Emerald Group.

Schmitt-Grohé, S. and M. Uribe (2012). What's News in Business Cycles. *Econometrica 80*(6), 2733–2764.

Sims, C. A. (1980). Macroeconomics and Reality. *Econometrica 48*(1), 1–48.

Sims, C. A. and T. Zha (2006). Does Monetary Policy Generate Recessions? *Macroeconomic Dynamics 10*(2), 231–272.

Sims, E. R. (2012). News, Non-Invertibility, and Structural VARs. In N. Balke, F. Canova, F. Milani, and M. A. Wyne (Eds.), *DSGE Models in Macroeconomics: Estimation, Evaluation, and New Developments*, pp. 81–135. Bingley, UK: Emerald Group Publishing.

Appendix

Proof of Lemma (1)**.** First, we observe that $\mathcal{H}(\xi)$ is isomorphic to $\mathcal{L}^2(F_\xi)$.[27] This can be seen by defining a correspondence between elements $h \in \mathcal{H}(\xi)$ of the form

$$h = \int \psi(\lambda)\Phi_\xi(d\lambda), \tag{20}$$

where

$$\int |\psi_k(\lambda)|^2 F_{\xi,kk}(d\lambda) < \infty, \quad k = 1, \ldots, n_\xi, \tag{21}$$

and the vector functions $\psi(\lambda) \in \mathcal{L}^2(F_\xi)$ which occur in the representation (20). This correspondence is linear, since $h_1 \leftrightarrow \psi_1$ and $h_2 \leftrightarrow \psi_2$ implies

$$\alpha_1 h_1 + \alpha_2 h_2 = \int (\alpha_1\psi_1(\lambda) + \alpha_2\psi_2(\lambda))\Phi_\xi(d\lambda) \leftrightarrow \alpha_1\psi_1 + \alpha_2\psi_2$$

for arbitrary scalars α_1, α_2. Moreover, it is isometric, since

$$(h_1, h_2) = \int \psi_1(\lambda)F_\xi(d\lambda)\psi_2(\lambda)^* = (\psi_1, \psi_2).$$

Because the closed linear manifold spanned by elements of the form (20) coincides with $\mathcal{H}(\xi)$, and the closed linear manifold spanned by elements $\psi(\lambda)$ of the form (21) coincides with $\mathcal{L}^2(F_\xi)$, it follows that correspondence we have defined can be extended by continuity to $\mathcal{H}(\xi)$ and $\mathcal{L}^2(F_\xi)$, preserving both its linearity and isometry.

Necessity: If $\mathcal{H}(\eta) \subseteq \mathcal{H}(\xi)$, then $\eta_{k,0} \in \mathcal{H}(\xi)$ for all $k = 1, \ldots, n_\eta$. Since $\mathcal{H}(\xi)$ is isomorphic to $\mathcal{L}^2(F_\xi)$, there exists a unique vector function $\psi(\lambda)$, whose rows are elements of $\mathcal{L}^2(F_\xi)$, such that

$$\eta_0 = \int \psi(\lambda)\Phi_\xi(d\lambda).$$

For every stationary process $\{\eta_t\}$, there exists a family of unitary operators U_t, $-\infty < t < \infty$, on $\mathcal{H}(\xi)$ such that

$$U_t\eta_{k,t} = \eta_{k,t+s}, \quad k = 1, \ldots, n_\eta$$

for any t, s. To the unitary operator U_t in $\mathcal{H}(\eta)$ corresponds the operator of multiplication by $e^{i\lambda t}$ in $\mathcal{L}^2(F_\eta)$; that is, for all $k = 1, \ldots, n_\eta$,

$$U_t\eta_{k,0} = U_t\left[\int \delta_k\psi(\lambda)\Phi_\xi(d\lambda)\right] = \eta_{k,t} = \int e^{i\lambda t}\delta_k\psi(\lambda)\Phi_\varepsilon(d\lambda),$$

[27]Recall that two Hilbert spaces are said to be "isomorphic" if it is possible to define a one-to-one correspondence between their elements which is linear and isometric.

where δ_k is a $1 \times n_\eta$ constant vector with components $\delta_{kk} = 1$ and $\delta_{kl} = 0$ for $k \neq l$. From this it follows that η_t has a representation of the form (2).

Sufficiency: Suppose there exists a function $\psi(\lambda)$ with rows in $\mathcal{L}^2(F_\xi)$ such that equation (2) holds. Then the function $e^{i\lambda t}\delta_k\psi(\lambda)$ is evidently also an element of $\mathcal{L}^2(F_\xi)$ for each $k = 1, \ldots, n_\eta$, since

$$\int e^{i\lambda t}\delta_k\psi(\lambda)F_\xi(d\lambda)\psi(\lambda)^*\delta_k^* e^{-i\lambda t} = \int \delta_k\psi(\lambda)F_\xi(d\lambda)\psi(\lambda)^*\delta_k^* < \infty.$$

Because $\mathcal{L}^2(F_\xi)$ is isomorphic to $\mathcal{H}(\xi)$, this means that $\eta_{k,t} \in \mathcal{H}(\xi)$ for $k = 1, \ldots, n_\eta$. Therefore, $\mathcal{H}(\eta) \subseteq \mathcal{H}(\xi)$. □

Proof of Theorem (1).[28]

Sufficiency: Equation (3) indicates that $\{y_t\}$ can be obtained from $\{\varepsilon_t\}$ by a linear transformation with spectral characteristic $\varphi(\lambda)$. This means that the random spectral measure of $\{y_t\}$ can be decomposed as[29]

$$\Phi_y(d\lambda) = \varphi(\lambda)\Phi_\varepsilon(d\lambda). \tag{22}$$

Because $\varphi(\lambda)$ has constant rank n_ε, there exists an $n_\varepsilon \times n_y$ matrix function $\psi(\lambda)$ such that

$$\psi(\lambda)\varphi(\lambda) = I_{n_\varepsilon}. \tag{23}$$

Combining equations (22) and (23), we get

$$\psi(\lambda)\Phi_y(d\lambda) = \Phi_\varepsilon(d\lambda).$$

Moreover, note that the rows of $\psi(\lambda)$ are elements of $\mathcal{L}^2(F_y)$ because for any $k = 1, \ldots, n_\varepsilon$, equations (22) and (23) imply that

$$\int \psi_k(\lambda)F_y(d\lambda)\psi_k(\lambda)^* = \frac{1}{2\pi}\int \psi_k(\lambda)\varphi(\lambda)\varphi(\lambda)^*\psi_k(\lambda)^* d\lambda = 1 < \infty.$$

Therefore $\{\varepsilon_t\}$ can be obtained from $\{y_t\}$ by a linear transformation with spectral characteristic $\psi(\lambda)$. By Lemma (1), it follows that the shocks are recoverable.

Necessity: To the contrary, suppose that the shocks are recoverable, so $\mathcal{H}(\varepsilon) \subseteq \mathcal{H}(y)$, but that $\varphi(\lambda)$ has rank different than n_ε on some set of positive measure.

[28]This proof is closely based on results from Rozanov (1967). Sufficiency is implied by Theorem 9.1, Ch.1, while necessity is implied by his subsequent discussion.

[29]More precisely, equation (22) means that $\Phi_y(\Delta) = \int_\Delta \psi(\lambda)\Phi_\varepsilon(d\lambda)$ for any Borel set Δ.

Because $\varphi(\lambda)$ has n_ε columns, its rank can never be greater than n_ε. Therefore, its rank on this set must be strictly less than this.

Now we find an element in $\mathcal{H}(\varepsilon)$ that is not in $\mathcal{H}(y)$, which is a contradiction. Because $\text{rank}(\varphi(\lambda)) < n_\varepsilon$ on some set of positive measure, there exists a $1 \times n_\varepsilon$ vector function $\psi(\lambda) \in \mathcal{L}^2(F_\varepsilon)$ such that $\|\psi(\lambda)\| \neq 0$ and

$$\varphi(\lambda)\psi(\lambda)^* = 0$$

for all λ. This would mean that the element

$$\eta = \int \psi(\lambda)\Phi_\varepsilon(d\lambda)$$

is orthogonal to $\mathcal{H}(y)$, because, for all $k = 1, \ldots, n_y$ and $-\infty < t < \infty$,

$$(y_{kt}, \eta) = \int e^{i\lambda t}\varphi_k(\lambda)\psi(\lambda)^* d\lambda = 0.$$

But this contradicts the hypothesis that $\mathcal{H}(\varepsilon) \subseteq \mathcal{H}(y)$. □

Proof of Theorem (2)**.** The fact that the variables w_s, $s \leq t$, form a basis in $\mathcal{H}_t(y)$ at each date means that a variable h is an element of $\mathcal{H}_t(y)$ if and only if it can be represented in the form of a series

$$h = \sum_{j=0}^{\infty} \alpha_j w_{t-j} \tag{24}$$

that converges in mean square. What we need to show is that each element of the vector ε_t has a representation of this form.

By the definition of $\psi(\lambda)$ and equation (6),

$$\varepsilon_t = \int e^{i\lambda t}\psi(\lambda)\Phi_y(d\lambda) = \int e^{i\lambda t}\psi(\lambda)\delta(\lambda)\Phi_w(d\lambda) \tag{25}$$

for all t. The rows of $\psi(\lambda)$ are elements of $\mathcal{L}^2(F_y)$, but they may not be square integrable with respect to the Lebesgue measure. On the other hand, the rows of $\alpha(\lambda) \equiv \psi(\lambda)\delta(\lambda)$ are square integrable, because $F_w(d\lambda) = \frac{1}{2\pi} I_{n_\varepsilon} d\lambda$. Therefore, $\alpha(\lambda)$ has a Fourier series expansion of the form

$$\alpha(\lambda) = \sum_{s=-\infty}^{\infty} \alpha_s e^{-i\lambda s}, \quad \text{where} \quad \alpha_s = \frac{1}{2\pi}\int e^{i\lambda s}\alpha(\lambda)d\lambda.$$

Combining this with equation (25), we can see that the elements of ε_t have a representation of the form (24) if and only if the Fourier coefficients $\{\alpha_s\}$ vanish for negative values of s, which is the condition stated in the theorem. □

Proof of Lemma (2)**.** By Lemma (1), the projections $\tilde{\varepsilon}_{k,t}$ form an n_ε dimensional stationary process $\{\tilde{\varepsilon}_t\}$ which is obtained from the process $\{y_t\}$ by a linear transformation,

$$\tilde{\varepsilon}_t = \int e^{i\lambda t}\psi(\lambda)\Phi_y(d\lambda),$$

where $\psi(\lambda)$ is some $n_\varepsilon \times n_y$ matrix function whose rows are elements of $\mathcal{L}^2(F_y)$. For the prediction errors $\varepsilon_{k,t} - \tilde{\varepsilon}_{k,t}$, $k = 1, \ldots, n_\varepsilon$, to be orthogonal to the space $\mathcal{H}(y)$, it must be that

$$E[(\varepsilon_t - \tilde{\varepsilon}_t)y_s^*] = \frac{1}{2\pi}\int e^{i\lambda(t-s)}\left[\varphi(\lambda)^* - \psi(\lambda)\varphi(\lambda)\varphi(\lambda)^*\right]d\lambda = 0$$

for any t and s. This is true if and only if

$$\varphi(\lambda)^* = \psi(\lambda)\varphi(\lambda)\varphi(\lambda)^* \tag{26}$$

for almost all λ. By definition, $\psi(\lambda) = \varphi(\lambda)^\dagger$ is a solution. Moreover, this solution is unique, in the sense that its rows are uniquely determined as elements of the space $\mathcal{L}^2(F_y)$. To see this, consider any other matrix function, $\psi(\lambda) \neq \varphi(\lambda)^\dagger$, whose rows are elements of $\mathcal{L}^2(F_y)$, which also satisfies (26). Then

$$\|\delta_k\varphi(\lambda)^\dagger - \delta_k\psi(\lambda)\|^2 = \int \delta_k(\varphi(\lambda)^\dagger - \psi(\lambda))\varphi(\lambda)\varphi(\lambda)^*(\varphi(\lambda)^\dagger - \psi(\lambda))^*\delta_k^* d\lambda = 0$$

for each $k = 1, \ldots, n_\varepsilon$, where δ_k denotes a $1 \times n_\varepsilon$ constant vector with components $\delta_{kk} = 1$ and $\delta_{kl} = 0$ for $k \neq l$. □

Proof of Theorem (3)**.** Using the optimal smoothing formula from Lemma (2),

$$\|\varepsilon_{k,t} - \tilde{\varepsilon}_{k,t}\|^2 = \frac{1}{2\pi}\int \delta_k(I_{n_\varepsilon} - \varphi(\lambda)^\dagger\varphi(\lambda))(I_{n_\varepsilon} - \varphi(\lambda)^\dagger\varphi(\lambda))^*\delta_k^* d\lambda,$$

which equals zero if and only if $\delta_k(I_{n_\varepsilon} - \varphi(\lambda)^\dagger\varphi(\lambda)) = 0$ almost everywhere. □

Proof of Lemma (3)**.** First we observe that the projections of $\varepsilon_{k,t}$ and $\tilde{\varepsilon}_{k,t}$ on $\mathcal{H}_t(y)$ coincide. Combining the representation of $\{\tilde{\varepsilon}_t\}$ from Lemma (2) with the Wold representation of $\{y_t\}$ in equation (6), we obtain

$$\tilde{\varepsilon}_t = \int e^{i\lambda t}\varphi(\lambda)^\dagger\gamma(\lambda)\Phi_w(d\lambda).$$

Using this representation of $\{\tilde{\varepsilon}_t\}$, we can see that the projections $\hat{\varepsilon}_{k,t}$ form a stationary process $\{\hat{\varepsilon}_t\}$ which is obtained from $\{w_t\}$ by a linear transformation of the form

$$\hat{\varepsilon}_t = \int e^{i\lambda t}[\varphi(\lambda)^\dagger \gamma(\lambda)]_+ \Phi_w(d\lambda).$$

Since $\gamma(\lambda)$ has full column rank for almost all λ, it follows that $\gamma(\lambda)^\dagger \gamma(\lambda) = I_{r_y}$, where r_y is the rank of $f_y(\lambda)$. Therefore

$$\Phi_w(d\lambda) = \gamma(\lambda)^\dagger \Phi_y(d\lambda).$$

Substituting this into the previous expression for $\Phi_w(d\lambda)$ gives the linear transformation reported in the lemma. Analogously to the proof of Lemma (2), the uniqueness of the projections $\hat{\varepsilon}_{k,t}$ implies that the spectral characteristic in this representation has rows which are all unique elements of $\mathcal{L}^2(F_y)$. □

Proof of Theorem (4)**.** Using the optimal filtering formula from Lemma (3),

$$\|\varepsilon_{k,t} - \hat{\varepsilon}_{k,t}\|^2 = \frac{1}{2\pi} \int \delta_k (I_{n_\varepsilon} - \alpha(\lambda)\gamma(\lambda)^\dagger \varphi(\lambda))(I_{n_\varepsilon} - \alpha(\lambda)\gamma(\lambda)^\dagger \varphi(\lambda))^* \delta_k^* d\lambda,$$

where $\alpha(\lambda) \equiv [\varphi(\lambda)^\dagger \gamma(\lambda)]_+$. This equals zero if and only if $\delta_k(I_{n_\varepsilon} - \alpha(\lambda)\gamma(\lambda)^\dagger \varphi(\lambda)) = 0$ for almost all λ. □

Supplemental Material

A Subordination Theory

As we have noted in Remark (3), our concept of recoverability is equivalent to the concept of "subordination," first introduced by Kolmogorov (1941). An important concern in the theory of subordination is to determine exactly when two processes are mutually subordinate to one another. Theorem 10 in Kolmogorov (1941) presents a first answer to this question in the univariate case. Theorems (1) and (3) in this paper extend this result to the multivariate case. However, both results invoke additional assumptions not present in Kolmogorov's original theorem. In this section, we prove a generalizations of our Theorems (1) and (3), which do not rely on these additional assumptions. In particular, we no longer require $\{y_t\}$ to be linearly regular, or $\{\varepsilon_t\}$ to have orthonormal values.

Before presenting the theorems, we introduce a preliminary result, which we will rely on in the proof. From Lemma (1), it follows that $\{\eta_t\}$ is recoverable from $\{\xi_t\}$ only if there exists an $n_\eta \times n_\xi$ matrix function $\psi(\lambda)$, with rows in $\mathcal{L}^2(F_\xi)$, such that the following relations hold:

$$F_\eta(d\lambda) = \psi(\lambda)F_\xi(d\lambda)\psi(\lambda)^* \tag{27}$$

$$F_{\eta_\xi}(d\lambda) = \psi(\lambda)F_\xi(d\lambda), \tag{28}$$

where $F_{\eta\xi}(d\lambda)$ denotes the joint spectral measure of $\{\eta_t\}$ and $\{\xi_t\}$; that is, $F_{\eta\xi,kl}(\mathit{\Delta}) \equiv E[\Phi_{\eta,k}(\mathit{\Delta})\overline{\Phi_{\xi,l}(\mathit{\Delta})}]$ for any Borel set $\mathit{\Delta}$. Theorem 8.1, Ch. 1, of Rozanov (1967) implies that these relations are also sufficient for $\{\eta_t\}$ to be recoverable from $\{\xi_t\}$.

Lemma 4. *$\{\eta_t\}$ is recoverable from $\{\xi_t\}$ if and only if there exits an $n_\eta \times n_\xi$ matrix function $\psi(\lambda)$, with rows in $\mathcal{L}^2(F_\xi)$, such that relations (27) and (28) hold.*

Proof. The necessity of the condition follows from Lemma (1).

Sufficiency: if relations (27) and (28) are satisfied, then by Lemma (1) we can define a process $\{\tilde{\eta}\}$, recoverable from $\{\xi_t\}$, for which

$$\tilde{\eta}_t = \int e^{i\lambda t}\psi(\lambda)\Phi_\xi(d\lambda).$$

With this definition, we have

$$F_{\tilde{\eta}}(d\lambda) = F_\eta(d\lambda) \quad \text{and} \quad F_{\tilde{\eta}\xi}(d\lambda) = F_{\eta\xi}(d\lambda).$$

Now we define an isometric operator T on the elements $\tilde{\eta}_{l,t}$ and $\xi_{k,t}$ in $\mathcal{H}(\xi)$ such that, for all $k = 1, \ldots, n_\xi$, $l = 1, \ldots, n_\eta$, and $-\infty < t < \infty$,

$$T\xi_{k,t} = \xi_{k,t} \quad \text{and} \quad T\tilde{\eta}_{l,t} = \eta_{l,t}.$$

This operator maps elements from $\mathcal{H}(\xi)$ into $\mathcal{H}(\xi, \eta)$, where the latter represents the Hilbert space generated by the variables $\xi_{k,t}$ and $\eta_{l,t}$. By continuity, this operator can be extended to an isometric mapping $\bar{T}$ from $\mathcal{H}(\xi)$ onto $\mathcal{H}(\xi, \eta)$. Therefore,

$$\bar{T}[\mathcal{H}(\xi)] = \mathcal{H}(\xi) \quad \text{and} \quad \bar{T}[\mathcal{H}(\xi)] = \mathcal{H}(\xi, \eta).$$

which imply that $\mathcal{H}(\xi) = \mathcal{H}(\xi, \eta)$. From this it follows that $\mathcal{H}(\eta) \subseteq \mathcal{H}(\xi)$. Furthermore, in order for $\bar{T}\xi_{k,t} = \xi_{k,t}$ for all k, $\bar{T}$ must be the identity operator, so $\tilde{\eta}_{l,t} = \eta_{l,t}$ for all l and t. □

Throughout the paper, we describe certain properties as holding "for almost all values of λ," which is always understood in reference to the Lebesgue measure. However, for the more general result in this section, we need to introduce the notion of a certain property holding for almost all values of λ with respect to an arbitrary spectral measure (which need not have absolutely continuous elements). Specifically, we say that a property is satisfied "almost everywhere with respect to $F_\xi(d\lambda)$" if it is satisfied for all λ, with the exception of some set Δ_0 for which $F_\xi(\Delta_0) = 0$. According to the definition of the integral used to define the space $\mathcal{L}^2(F_\xi)$, this means that

$$\int \psi(\lambda) F_\xi(d\lambda) \psi(\lambda)^* = 0$$

if and only if $\psi(\lambda) = 0$ almost everywhere with respect to $F_\xi(d\lambda)$.

Now we present the main result of this section.

Theorem 5. *Suppose that $\{y_t\}$ and $\{\varepsilon_t\}$ are jointly stationary, and that $\{y_t\}$ can be obtained from $\{\varepsilon_t\}$ by a linear transformation with spectral characteristic $\varphi(\lambda)$. Then $\{\varepsilon_t\}$ is recoverable from $\{y_t\}$ if and only if $\varphi(\lambda)$ has full column rank almost everywhere with respect to $F_\varepsilon(d\lambda)$.*

Proof. Necessity: By Lemma (1), $\{y_t\}$ is recoverable from $\{\varepsilon_t\}$. If $\{\varepsilon_t\}$ is also recoverable from $\{y_t\}$, then Lemma (4) implies that there exists a $n_\varepsilon \times n_y$ matrix function $\psi(\lambda)$ such that the following relations hold,

$$F_\varepsilon(d\lambda) = \psi(\lambda) F_y(d\lambda) \psi(\lambda)^* \qquad F_y(d\lambda) = \varphi(\lambda) F_\varepsilon(d\lambda) \varphi(\lambda)^*$$

$$F_{\varepsilon y}(d\lambda) = \psi(\lambda) F_y(d\lambda) \qquad F_{y\varepsilon}(d\lambda) = \varphi(\lambda) F_\varepsilon(d\lambda).$$

Combining the top two equations,

$$F_\varepsilon(d\lambda) = \psi(\lambda)\varphi(\lambda)F_\varepsilon(d\lambda)\varphi(\lambda)^*\psi(\lambda)^*.$$

Since both sides of this equation are finite, Lemma (1) implies that we can define a new process $\{\tilde{\varepsilon}_t\}$, recoverable with respect to $\{y_t\}$, by the linear transformation

$$\tilde{\varepsilon}_t = \int e^{i\lambda t}\psi(\lambda)\varphi(\lambda)\Phi_\varepsilon(d\lambda).$$

The spectral measures of $\{\tilde{\varepsilon}_t\}$ and $\{\varepsilon_t\}$ satisfy the following relations,

$$F_{\tilde{\varepsilon}}(d\lambda) = F_\varepsilon(d\lambda) \quad \text{and} \quad F_{\tilde{\varepsilon}y}(d\lambda) = F_{\varepsilon y}(d\lambda).$$

But then, as we saw in the proof of Lemma (4), it must be the case that the two processes $\{\tilde{\varepsilon}\}$ and $\{\varepsilon_t\}$ are the same. This implies that their difference, $\varepsilon_t - \tilde{\varepsilon}_t$, has a variance of zero,

$$\int (I_{n_\varepsilon} - \psi(\lambda)\varphi(\lambda))F_y(d\lambda)(I_{n_\varepsilon} - \psi(\lambda)\varphi(\lambda))^* = 0.$$

Therefore, $\psi(\lambda)\varphi(\lambda) = I_{n_\varepsilon}$ almost everywhere with respect to $F_\varepsilon(d\lambda)$, which is equivalent to the statement that $\varphi(\lambda)$ has full column rank almost everywhere with respect to $F_\varepsilon(d\lambda)$, since on any set of positive measure,

$$n_\varepsilon = \text{rank}(\psi(\lambda)\varphi(\lambda)) \leq \text{rank}(\varphi(\lambda)) \leq n_\varepsilon.$$

Sufficiency: In this case, $\{y_t\}$ is recoverable from $\{\varepsilon_t\}$ and $\varphi(\lambda)$ has full column rank almost everywhere with respect to $F_\varepsilon(d\lambda)$. As we just observed, the rank condition on $\varphi(\lambda)$ means that $\psi(\lambda)\varphi(\lambda) = I_{n_\varepsilon}$ almost everywhere with respect to $F_\varepsilon(d\lambda)$. Therefore, we can write

$$F_\varepsilon(d\lambda) = \psi(\lambda)\varphi(\lambda)F_\varepsilon(d\lambda)\varphi(\lambda)^*\psi(\lambda)^* = \psi(\lambda)F_y(\lambda)\psi(\lambda)^*$$
$$F_{\varepsilon y}(d\lambda) = \psi(\lambda)\varphi(\lambda)F_\varepsilon(d\lambda)\varphi(\lambda)^* = \psi(\lambda)F_y(d\lambda).$$

By Lemma (4), it follows that $\{\varepsilon_t\}$ is recoverable from $\{y_t\}$; and Lemma (2) indicates that the former can be obtained from the latter by a linear transformation with spectral characteristic $\psi(\lambda)$, which is uniquely defined in terms of the space $\mathcal{L}^2(F_y)$. □

Remark 9. Notice that if $F_\varepsilon(d\lambda) = I_{n_\varepsilon} d\lambda$, which occurs when $\{\varepsilon_t\}$ has orthonormal values as in Assumption (2), then the notions "almost everywhere with respect to $F_\varepsilon(d\lambda)$ and "almost everywhere with respect to the Lebesgue measure" coincide. This implies that Theorem (1) remains valid as stated regardless of whether or not $\{y_t\}$ is a linearly regular process.

Remark 10. Rosenberg (1969) also presents a generalization of Kolmogorov's Theorem 10, which applies to multivariate processes that are defined over any locally compact abelian group. The five conditions in his Theorem 3.4 can be shown to be equivalent to the one presented here in the case that the locally compact abelian group under consideration is the integers.

Here we also present a generalization of Theorem (3).

Theorem 6. *Under the assumptions in Theorem* (5), *the process* $\{\varepsilon_{k,t}\}$ *is recoverable from* $\{y_t\}$ *if and only if*

$$\delta_k(I_{n_\varepsilon} - \psi(\lambda)\varphi(\lambda)) = 0$$

almost everywhere with respect to $F_\varepsilon(d\lambda)$, *where*

$$\psi(\lambda) = f_\varepsilon^{(\mu)}(\lambda)\varphi(\lambda)^*[\varphi(\lambda)f_\varepsilon^{(\mu)}(\lambda)\varphi(\lambda)^*]^\dagger.$$

Proof. As in the proof of Lemma (2), let $\{\tilde{\varepsilon}_t\}$ denote the process whose elements are the projections of the elements of $\{\varepsilon_t\}$ onto $\mathcal{H}(y)$,

$$\tilde{\varepsilon}_t = \int e^{i\lambda t}\psi(\lambda)\Phi_y(d\lambda),$$

where $\psi(\lambda)$ is some $n_\varepsilon \times n_y$ matrix function whose rows are elements of $\mathcal{L}^2(F_y)$. The existence and uniqueness of such a function are guaranteed by Lemma (1). Then $\{\varepsilon_{k,t}\}$ is recoverable if and only if

$$\|\varepsilon_{k,t} - \tilde{\varepsilon}_{k,t}\|^2 = \int \delta_k(I_{n_\varepsilon} - \psi(\lambda)\varphi(\lambda))F_\varepsilon(d\lambda)(I_{n_\varepsilon} - \varphi(\lambda)^\dagger\varphi(\lambda))^*\delta_k^* = 0,$$

which will be satisfied if and only if $\delta_k(I_{n_\varepsilon} - \psi(\lambda)\varphi(\lambda)) = 0$ almost everywhere with respect to $F_\varepsilon(d\lambda)$.

Now, what remains is to find the solution for $\psi(\lambda)$ in terms of $\varphi(\lambda)$ and $F_\varepsilon(d\lambda)$. For all the prediction errors $\zeta_{k,t} \equiv \varepsilon_{k,t} - \tilde{\varepsilon}_{k,t}$ to be orthogonal to the space $\mathcal{H}(y)$, it must be that

$$F_{\zeta y}(d\lambda) = (I_{n_\varepsilon} - \psi(\lambda)\varphi(\lambda))F_\varepsilon(d\lambda)\varphi(\lambda)^* = 0.$$

By re-arranging this expression and using $F_\varepsilon(d\lambda) = f_\varepsilon^{(\mu)}(\lambda)\mu(d\lambda)$, we arrive at the equation

$$f_\varepsilon^{(\mu)}(\lambda)\varphi(\lambda)^* = \psi(\lambda)\varphi(\lambda)f_\varepsilon^{(\mu)}(\lambda)\varphi(\lambda)^*,$$

which holds almost everywhere with respect to $\mu(d\lambda)$. From this equation it is easy to see that the row space of the matrix $A(\lambda) = f_\varepsilon^{(\mu)}(\lambda)\varphi(\lambda)^*$ is contained in the row space of the matrix $B(\lambda) = \varphi(\lambda)f_\varepsilon^{(\mu)}(\lambda)\varphi(\lambda)^*$. Since $x^*B(\lambda)^\dagger B(\lambda) = x^*$ for any x in the row space of $B(\lambda)$, it follows that $\psi(\lambda) = A(\lambda)B(\lambda)^\dagger$ satisfies the preceding equation. Moreover, this solution is unique, in the sense that its rows are uniquely determined as elements of the space $\mathcal{L}^2(F_y)$. □

Remark 11. In the event that $\{\varepsilon_t\}$ has orthonormal values, we can take $\mu(d\lambda) = d\lambda$ so that $f_\varepsilon^{(\mu)}(\lambda) = I_{n_\varepsilon}$. Substituting these into the expression for $\psi(\lambda)$,

$$\psi(\lambda) = \varphi(\lambda)^*[\varphi(\lambda)\varphi(\lambda)^*]^\dagger = \varphi(\lambda)^\dagger,$$

so the condition from this theorem and the one from Theorem (3) coincide.

Remark 12. Nothing in the proofs of the theorems from this section require the parameter t only to take on integer values. Indeed, both results continue to hold as stated if t instead takes on all real values. We discuss the continuous time case in more detail in the following section.

B Continuous Time

Our discussion of recoverability and invertibility in the paper focuses on the case of discrete time, when the parameter t takes on all integer values. However, a similar analysis can be performed in the case of continuous time, when t takes on all real values. The main complication is that the idea of a structural shock process being an mutually uncorrelated random process $\{\varepsilon_t\}$ with a flat spectral density no longer applies in continuous time. Instead, we need to think of the structural shocks as mutually uncorrelated random measures. We can show that, with this change in the interpretation of structural shocks, Theorem (1) continues to hold exactly as stated, while Theorem (2) requires some slight changes.

First, we observe that Definitions (1) and (2) do not depend on time being discrete, so they remain the same. Lemma (1) also continues to hold in continuous time, provided that the limits of integration are set to $-\infty, \infty$ rather than $-\pi, \pi$.

This is because any continuous-time wide-sense stationary process $\{\xi_t\}$ such that the functions

$$B_{kl}(t) = E[\xi_{k,t+s}\overline{\xi_{l,s}}], \quad k, l = 1, \ldots, n_\xi$$

are continuous in the parameter t has a spectral representation that is similar to the one in equation (1), but with different limits of integration,

$$\xi_t = \int_{-\infty}^{\infty} e^{i\lambda t} \Phi_\xi(d\lambda).$$

Second, we extend Definitions (1) and (2) so they apply to an n_ζ dimensional random measure $\zeta(dt)$, defined on the Borel sets of the real line. We let $\mathcal{H}(\zeta)$ denote the Hilbert space spanned by the variables $\zeta_{k,t}(\Delta)$ for $k = 1, \ldots, n_\zeta$ and any Δ on the line $-\infty < t < \infty$. Similarly, we let $\mathcal{H}_t(\zeta)$ denote the subspace spanned by these variables over all k, but only for Δ lying in the half-line $(-\infty, t]$. We say that the random measure $\zeta(dt)$ is recoverable or invertible depending on whether

$$\mathcal{H}(\zeta) \subseteq \mathcal{H}(y) \quad \text{or} \quad \mathcal{H}_t(\zeta) \subseteq \mathcal{H}_t(y).$$

In continuous time, we consider an economic model of the form

$$y_t = \int_{-\infty}^{\infty} \varphi(\lambda) \Phi_\varepsilon(d\lambda). \tag{29}$$

This representation has the same form as its discrete-time analogue in (3), except that the limits of integration are set to $-\infty, \infty$ rather than $-\pi, \pi$.[30] But it is still the case that the random measures $\Phi_{\varepsilon,k}(d\lambda)$, $k = 1, \ldots, n_\varepsilon$ are mutually uncorrelated, i.e.,

$$E[\Phi_{\varepsilon,k}(\Delta_1)\overline{\Phi_{\varepsilon,l}(\Delta_2)}] = 0$$

if $k \neq l$, for any Borel sets Δ_1 and Δ_2 of the real line; and, moreover,

$$E[|\Phi_{\varepsilon,k}(d\lambda)|^2] = \frac{1}{2\pi} d\lambda$$

for all $k = 1, \ldots, n_\varepsilon$.

The difference is that in the continuous-time case, $\Phi_\varepsilon(d\lambda)$ cannot be the random spectral measure of an uncorrelated stationary process. This is because any such process would not have finite variance. Instead, we need to understand the structural

[30] In this subsection, when the limits of integration are omitted, they are understood to be $-\infty, \infty$.

shocks to be a collection of uncorrelated random measures $\varepsilon_k(dt)$, $k = 1, \dots, n_\varepsilon$, defined as the Fourier transforms of the measures $\Phi_{\varepsilon,k}(d\lambda)$, $k = 1, \dots, n_\varepsilon$. That is, we set

$$\varepsilon(\Delta) = \int \frac{e^{i\lambda t_2} - e^{i\lambda t_1}}{i\lambda} \Phi_\varepsilon(d\lambda),$$

for any semi-interval $\Delta = (t_1, t_2]$, and then take the extension of this measure to all Borel sets of the real line. With these changes, we obtain the following continuous-time "moving-average" representation of $\{y_t\}$

$$y_t = \int_{-\infty}^{\infty} \varphi_{t-s} \varepsilon(ds),$$

where $\{\varphi_s\}$ are the Fourier coefficients of the function $\varphi(\lambda)$.

In sum, we can replace Assumption (2) with the following assumption.

Assumption 3. $\{y_t\}$ *can be obtained from the* n_ε *dimensional mutually uncorrelated random measure* $\varepsilon(dt)$ *by a relation of the form*

$$y_t = \int e^{i\lambda t} \varphi(\lambda) \Phi_\varepsilon(d\lambda) \quad \text{for all } t, \tag{30}$$

where $E[|\varepsilon_k(dt)|^2] = dt$ *for all* k, *and* $E[\varepsilon_k(\Delta_1)\overline{\varepsilon_l(\Delta_2)}] = 0$ *for* $k \neq l$ *and any Borel sets* Δ_1 *and* Δ_2.

Example 6. As a special case of the model in equation (30), suppose that we have a state-space structure of the form

$$\begin{aligned} &\text{(observation)} \qquad y_t = Ax_t \\ &\text{(state)} \qquad dx_t = Bx_t dt + C d\varepsilon_t, \end{aligned}$$

where x_t is an n_x dimensional state vector, and $\{\varepsilon_t\}$ is an n_ε dimensional continuous-time process with orthogonal increments. The values $\varepsilon_{k,t}$ of this process are related to the mutually uncorrelated random measures $\varepsilon_k(dt)$ by the correspondence

$$\varepsilon_{k,t} = \varepsilon_k(\Delta), \quad \Delta = (-\infty, t],$$

for all t and $k = 1, \dots, n_\varepsilon$. This system corresponds to the the spectral characteristic

$$\varphi(\lambda) = A(i\lambda I_{n_x} - B)^{-1} C.$$

◇

It turns out that the recoverability or invertibility of the random measure $\varepsilon(dt)$ can be shown to coincide with the recoverability or invertibility of the stationary random process $\{\eta_t\}$, defined by

$$\eta_t = \int e^{i\lambda t} \Phi_\eta(d\lambda), \quad \Phi_\eta(d\lambda) = (1 + i\lambda)^{-1} \Phi_\varepsilon(d\lambda).$$

This is an n_ε dimensional process with the property that $\mathcal{H}(\eta) = \mathcal{H}(\varepsilon)$ and $\mathcal{H}_t(\eta) = \mathcal{H}_t(\varepsilon)$ for all t. This can be seen from the fact that, for any $\Delta = (t_1, t_2]$ with $t_2 \leq t$,

$$\begin{aligned} \varepsilon(\Delta) &= \int \frac{e^{i\lambda t_2} - e^{i\lambda t_1}}{i\lambda} (1 + i\lambda) \Phi_\eta(d\lambda) \\ &= \eta_{t_2} - \eta_{t_1} + \int_{t_1}^{t_2} \eta_s ds. \end{aligned}$$

Therefore our objective can be re-framed in terms of determining whether $\{\eta_t\}$ is recoverable or invertible, which facilitates matters.

We begin by showing that the recoverability condition in Theorem (1) continues to apply, exactly as stated, when time is continuous. The projections $\tilde{\eta}_{k,t}$ of the values $\eta_{k,t}$ onto the space $\mathcal{H}(y)$ form an n_ε dimensional stationary process $\{\tilde{\eta}_t\}$, which is obtained from $\{y_t\}$ by a linear transformation

$$\tilde{\eta}_t = \int e^{i\lambda t} \tilde{\psi}(\lambda) \Phi_y(d\lambda), \tag{31}$$

where, by Lemma (1), the rows of $\tilde{\psi}(\lambda)$ are elements of $\mathcal{L}^2(F_y)$. In terms of the process $\{\eta_t\}$, the economic model (30) can be written as

$$y_t = \int e^{i\lambda t} (1 + i\lambda) \varphi(\lambda) \Phi_\eta(d\lambda) \equiv \int e^{i\lambda t} \tilde{\varphi}(\lambda) \Phi_\eta(d\lambda).$$

This means that for any t and s,

$$E[(\eta_t - \tilde{\eta}_t) y_s^*] = \int e^{i\lambda(t-s)} \frac{1}{2\pi(1 + \lambda^2)} \left[\tilde{\varphi}(\lambda)^* - \tilde{\psi}(\lambda) \tilde{\varphi}(\lambda) \tilde{\varphi}(\lambda)^* \right] d\lambda,$$

which equals zero if and only if

$$\tilde{\varphi}(\lambda)^* = \tilde{\psi}(\lambda) \tilde{\varphi}(\lambda) \tilde{\varphi}(\lambda)^*$$

for almost all λ. Using the fact that $\tilde{\varphi}(\lambda) = (1 + i\lambda) \varphi(\lambda)$, it follows that

$$\tilde{\psi}(\lambda) = \tilde{\varphi}(\lambda)^\dagger = (1 + i\lambda)^{-1} \varphi(\lambda)^\dagger.$$

Therefore the (squared) distance between $\eta_{k,t}$ and the projection $\tilde{\eta}_{k,t}$ is

$$\|\eta_{k,t} - \tilde{\eta}_{k,t}\|^2 = \int \frac{1}{2\pi(1+\lambda^2)} \delta_k (I_{n_\varepsilon} - \varphi(\lambda)^\dagger \varphi(\lambda))(I_{n_\varepsilon} - \varphi(\lambda)^\dagger \varphi(\lambda))^* \delta_k^* d\lambda,$$

which equals zero if and only if

$$\delta_k (I_{n_\varepsilon} - \varphi(\lambda)^\dagger \varphi(\lambda)) = 0$$

for almost all λ. And this is the same condition stated in Theorem (3).

Theorem 7 (Recoverability: continuous time)**.** *Under Assumptions* (1) *and* (3)*, the measure $\varepsilon_k(dt)$ is recoverable from the observable process $\{y_t\}$ if and only if*

$$\delta_k (I_{n_\varepsilon} - \varphi(\lambda)^\dagger \varphi(\lambda)) = 0$$

almost everywhere, where δ_k denotes a $1 \times n_\varepsilon$ constant vector with components $\delta_{kk} = 1$ and $\delta_{kl} = 0$ for $k \neq l$.

Now we show how the invertibility condition in Theorem (4) carries over with some slight changes. The continuous-time version of Wold's decomposition theorem implies that it is possible to represent $\{y_t\}$ in the form of equation (6), where now $\Phi_w(d\lambda)$ is the random spectral measure associated with an r_y dimensional random measure with mutually uncorrelated values, $w(dt)$, which has the property that $\mathcal{H}_t(w) = \mathcal{H}_t(y)$ for all t.

As in the discrete-time case, we can find the projections $\hat{\eta}_{k,t}$ of the variables $\eta_{k,t}$ onto the subspace $\mathcal{H}_t(y)$ by projecting the variables $\tilde{\eta}_{k,t}$ onto this space. Combining equations (31) and (6),

$$\tilde{\eta}_t = \int e^{i\lambda t} \tilde{\varphi}(\lambda)^\dagger \gamma(\lambda) \Phi_w(d\lambda).$$

As in the discrete-time case, let us denote by $[\varphi(\lambda)]_+$ the matrix function

$$[\varphi(\lambda)]_+ = \int_0^\infty \varphi_s e^{-i\lambda s}$$

for any matrix function $\varphi(\lambda)$ whose elements are square integrable, where $\{\varphi_s\}$ are the Fourier coefficients of $\varphi(\lambda)$. Projecting on the subspace $\mathcal{H}_t(y)$, we obtain

$$\hat{\eta}_t = \int e^{i\lambda t} [\tilde{\varphi}(\lambda)^\dagger \gamma(\lambda)]_+ \Phi_w(d\lambda) = \int e^{i\lambda t} [\tilde{\varphi}(\lambda)^\dagger \gamma(\lambda)]_+ \gamma(\lambda)^\dagger \Phi_y(d\lambda).$$

Therefore,

$$\|\eta_{k,t} - \hat{\eta}_{k,t}\|^2 = \int \frac{1}{\pi(1+\lambda^2)} \delta_k (I_{n_\varepsilon} - \tilde{\alpha}(\lambda)\gamma(\lambda)^\dagger \tilde{\varphi}(\lambda))(I_{n_\varepsilon} - \tilde{\alpha}(\lambda)\gamma(\lambda)^\dagger \tilde{\varphi}(\lambda))^* \delta_k^* d\lambda,$$

where $\tilde{\alpha}(\lambda) \equiv [\tilde{\varphi}(\lambda)\gamma(\lambda)]_+$. This equals zero if and only if

$$\delta_k (I_{n_\varepsilon} - \tilde{\alpha}(\lambda)\gamma(\lambda)^\dagger \tilde{\varphi}(\lambda)) = 0$$

for almost all λ. We have therefore arrived at the following result.

Theorem 8 (Invertibility: continuous time)**.** *Under Assumptions* (1) *and* (3)*, the measure* $\varepsilon_k(dt)$ *is invertible from the observable process* $\{y_t\}$ *if and only if*

$$\delta_k (I_{n_\varepsilon} - [\tilde{\varphi}(\lambda)\gamma(\lambda)]_+ \gamma(\lambda)^\dagger \tilde{\varphi}(\lambda)) = 0$$

almost everywhere, where $\tilde{\varphi}(\lambda) = (1 + i\lambda)\varphi(\lambda)$*,* $\gamma(\lambda)$ *comes from some version of Wold's decomposition of* $\{y_t\}$*, and* δ_k *denotes a* $1 \times n_\varepsilon$ *constant vector with components* $\delta_{kk} = 1$ *and* $\delta_{kl} = 0$ *for* $k \neq l$*.*

Remark 13. As in the discrete-time case, it is possible to articulate an alternative necessary and sufficient condition for invertibility in terms of the Fourier coefficients of $\tilde{\varphi}(\lambda)^\dagger \gamma(\lambda)$. Namely, $\varepsilon_k(dt)$ is invertible if and only if it is recoverable and

$$\frac{1}{2\pi} \int e^{i\lambda s} \delta_k \tilde{\varphi}(\lambda)^\dagger \gamma(\lambda) d\lambda = 0$$

for all $s < 0$.