Recoverability*

Ryan Chahrour — Boston College

Kyle Jurado — Duke University

May 6, 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 even to models with non-invertible shocks. We illustrate these results in the context of a simple model of consumption determination with productivity shocks and non-productivity noise shocks. In an application to postwar U.S. data, we find that non-productivity shocks account for a large majority of fluctuations in aggregate consumption over business cycle frequencies.

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 almost everywhere. 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 perfectly infer 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 usually 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, then it is impossible that the reduced-form shocks be equal to the underlying structural shocks. As a result, it is common practice first to verify 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

[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).

additional theoretical restrictions on the data 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. Given this 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 structural shocks of interest. If the structural representation happens to be non-invertible, so be it. Just because it may be desirable to estimate an invertible model in the reduced-form step, that should not in any way tie our hands when we get to the structural step. There are generally many different representations consistent with the same autocovariance function, and it is the role of economic theory to help us pick out an economically interesting one.

From this perspective, it is also easy to see that the reduced-form model doesn't have to be invertible either. The econometrician could easily estimate a non-invertible or even non-parametric model in the reduced-form step. All that is required is to obtain a characterization of the autocovariance function of the observable processes. Naturally, some reduced-form models will do a better job than others in specific contexts. Our purpose in this paper is not to advocate for any particular one. Instead, it is to determine when it is possible to recover structural shocks of interest given a satisfactory reduced-form representation of the autocovariance structure of the data.

One strand of the macroeconomic literature in which semi-structural methods have been eschewed involves models with purely belief-driven fluctuations. In particular, Blanchard et al. (2013) argue that structural vector autoregression (VAR) analysis cannot be applied to models with non-fundamental noise shocks because they are inherently non-invertible. In a determinate rational expectations model, if economic agents could tell on the basis of current and past data that a shock was pure noise, they would not respond to it. Therefore it is impossible to recover noise shocks from current and past data.[5]

[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 is the original perspective taken by Sims (1980); see his description on p.15. In his application, he uses an invertible vector autoregression as the reduced-form model, but neither invertibility nor vector autoregressions are necessary features of his proposed empirical strategy.

[5]For a more extended discussion of the limitations of using structural VAR analysis in this

While it is true that noise shocks are not invertible, they are often recoverable. As an application of our results, we show that our recoverability condition is satisfied in an analytically convenient model of consumption determination with noise shocks taken from Blanchard et al. (2013). We then perform a Monte Carlo exercise to show how structural VAR analysis can be applied in this situation. Finally, we apply the same procedure to a sample of postwar U.S. data on consumption and productivity. We find that less than 15% of the business-cycle variation in consumption can be attributed to productivity shocks, with all remaining fluctuations attributed to non-productivity noise. This finding represents a challenge for theories of consumption determination that rely primarily on beliefs about productivity. It implies that in any such theory, beliefs about productivity must be fluctuating in ways that are mostly unrelated to productivity itself.

A few papers have suggested that semi-structural methods are not necessarily inapplicable when invertibility fails. Lippi and Reichlin (1994) examine a particular subset of non-invertible representations ("basic" ones) given an invertible reduced-form model. Sims and Zha (2006) propose an iterative algorithm to check whether certain structural shocks are "approximately invertible," even if they are not invertible. Dupor and Han (2011) develop a four-step procedure to partially identify structural impulse responses whether or not non-invertibility is present. Plagborg-Møller (2017) suggests that estimating a moving average reduced-form model rather than an autoregression can help avoid concerns of non-invertibility. In a paper closely related to our empirical application, Forni et al. (2017a) write down a particular model with noise shocks and show that it is possible to identify those shocks by finding appropriate dynamic rotations of reduced-form VAR residuals. Forni et al. (2017b) also perform a similar analysis in an asset-pricing context.

These papers tend to give the impression that non-invertibility is a problem, but one that can be circumvented in some special cases or with additional effort. For example, one might need to use a different reduced-form model, or rely on additional theoretical restrictions to select one non-invertible representation among many. By contrast, we argue that non-invertibility is *never* a problem, at least not when it comes to using semi-structural methods. What really matters is recoverability. Neither a particular reduced-form model, nor additional theoretical restrictions are required to use semi-structural methods.

literature, see the review article by Beaudry and Portier (2014).

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, 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

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

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.

It is important to observe that both recoverability and invertibility are *population* concepts. As such, they are closely related to the econometric concept of identification. If a certain process is recoverable from another, this can be taken to

[6]Analogues to all our results in the case of continuous time, when t takes on all real values, 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.

mean that the former is identified from the latter. Likewise, if a process is invertible form another, this means that the value of the former at each date is identified from the current and past history of the latter. Of course, this use of "identification" needs to be understood in an appropriately general sense. Its object here is not a finite-dimensional vector of numbers, as is typically the case; rather, it refers to an infinite-dimensional family of random variables.[8]

An alternative but 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]

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

In this expression, $F_\xi(d\lambda)$ denotes the spectral measure of $\{\xi_t\}$ and the asterisk denotes complex conjugate transposition.[10] If we define the scalar product

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

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. Using these definitions, the following lemma gives an alternative charaterization 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*

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

[8]This sense of identification can be made more precise by establishing a connection with the literature on the identification of nonparametric functions, but because it would take us too far afield, we do not pursue such a connection in this paper.

[9]From now on, whenever the limits of integration are omitted, we will understand them to be $-\pi, \pi$ unless otherwise indicated.

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

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 a 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.

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

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

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.[11] $\diamond$

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 same is true in the opposite direction. 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 for almost all* λ*.*

The logic behind 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 couple of remarks regarding the theorem.

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 simply check whether $\varphi(\lambda_u)$ is full column rank.

[11] 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.

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.[12] 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 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.[13] 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*

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

where $\psi(\lambda)$ *is any* $n_\varepsilon \times n_y$ *matrix function satisfying* $\psi(\lambda)\varphi(\lambda) = I_{n_\varepsilon}$ *for almost all* λ*, 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 economic motivations.

[12]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 autoregressive and moving average 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 sufficient but not necessary for invertibility as we have defined it.

[13]See Rozanov (1967), Ch. 2.

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:

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

◇

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

Remark 3. 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 can say that a process $\{\eta_t\}$ is 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 Partial Recoverability

This section extends our results from the previous section to cover situations in which, within a single model, some shocks may be recoverable while others may not. Even though the results in this section are technically more general, they require more of an investment in terms of notation and machinery. On a first reading, it is possible to skip this section and jump directly to Section (4).

Given a system of the form (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. We refer to these situations as ones of "partial recoverability." The following example provides a simple illustration.

Example 3. In a model of the form (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 $\{\varepsilon_{3,t}\}$, because

$$\varepsilon_{3,t} = y_{3,t}$$

for all t. ◇

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

(ii) $\varphi(\lambda)^\dagger = \varphi(\lambda)^{-1}$ if and only if $\varphi(\lambda)$ is nondegenerate

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 fully general 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 can be 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*

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

for almost all λ*, 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 4. 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 5. 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, so also it is 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 $\mathcal{H}_t(y)$.[14]

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.

[14]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.

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*

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

for almost all λ*, 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 6. It can be helpful to see how this result generalizes Theorem (2). If $\{\varepsilon_{k,t}\}$ is recoverable, then Lemma (2) implies that

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

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).

4 Semi-Structural Analysis

This section describes the relationship between recoverability and semi-structural empirical analysis. We argue that recoverability is the key condition for performing this type of analysis, and explain why invertibility is not necessary. We begin with a brief review, which helps to frame the discussion.

4.1 Importance of Recoverability

The semi-structural approach, which goes back to the seminar paper of Sims (1980), represents a hybrid between purely statistical and fully structural approaches to analyzing economic time series. Purely statistical models can provide good empirical fit, but are not amenable to economic interpretation. By contrast, fully structural models are amenable to economic interpretation, but require imposing what some have seen as “incredible” (to borrow Sims’ term) theoretical restrictions. The idea behind the semi-structural approach, which has by now found wide acceptance in the macroeconomic literature, is to combine the empirical flexibility of purely statistical models with a “credible” subset of the restrictions implied by a fully structural model (or a class of such models).

This combination is achieved in two steps: a reduced-form step and a structural step. The reduced-form step involves using 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 spectral density defines a *set* of observationally equivalent reduced-form representations of the observable process. The structural step then involves imposing economic restrictions to select one reduced-form representation from this set, which has the property that, under the null hypothesis that the structural model is correctly specified, it *is* the structural representation. There exists a reduced-form representation with this property only if the structural shocks are recoverable from the observables.

To explain these steps in more detail, we introduce the following definition of a reduced-form representation.

Definition 4. In a “reduced-form representation,” $\{y_t\}$ is related to a zero-mean process with orthonormal values $\{u_t\}$ by a linear transformation with spectral characteristic $\rho(\lambda)$, which has full column rank for almost all λ. That is,

$$y_t = \int e^{i\lambda t} \rho(\lambda) \Phi_u(d\lambda) \quad \text{for all } t \tag{10}$$

where $\rho(\lambda)$ is an $n_y \times r_y$ matrix function, r_y is the rank of $f_y(\lambda)$ for almost all λ, $E[u_{k,t}] = 0$ and $E[|u_{k,t}|^2] = 1$ for all $k = 1, \dots, n_u$, and $E[u_{k,t}\overline{u_{l,s}}] = 0$ for $t \neq s$ and all $k, l = 1, \dots, n_\varepsilon$, as well as for $k \neq l$ and all t, s.

Importantly, $\rho(\lambda)$ is full column rank for almost all λ, which, by Theorem (1) means that the reduced-form shocks $\{u_t\}$ are always recoverable. This is part and parcel of what it means for the shocks to be "reduced-form."

The reduced-form step involves using statistical methods to characterize the spectral density $f_y(\lambda)$ from a time series of observables. It is possible to do this whenever the observables are linearly regualr and wide-sense stationary, even if the structural model in equation (3) is not correctly specified. In other words, whenever Assumption (1) is satisfied, even if Assumption (2) is not. The spectral density defines a set of observationally equivalent reduced-form representations of the type in Definition (4). Each reduced-form representation corresponds to an $n_y \times r_y$ matrix function, with rows in $\mathcal{L}^2(F_u)$, which satisfies the equation

$$f_y(\lambda) = \frac{1}{2\pi}\rho(\lambda)\rho(\lambda)^* \tag{11}$$

for almost all λ. We let $\mathcal{R}(f_y)$ denote the set of all such functions.

Common practice in this first step, following the application in Sims (1980), is to use an autoregressive model to characterize the spectral density of observables. But of course, nothing requires the use of such a statistical model. One could consider models with moving average terms as well as autoregressive terms; one could consider non-invertible versions of these models as well as invertible ones; one could even take an entirely non-parametric approach. The objective is only to characterize the spectral density, and statistical criteria for evaluating goodness of fit can be used to guide this part of the analysis.

In the structural step, an estimate of the structural representation is chosen from among the set of reduced-form representations defined by equation (11). This is accomplished by imposing theoretical restrictions with the following two properties:

(P1) they identify exactly one reduced-form representation in this set, regardless of whether Assumption (2) is satisfied, and

(P2) the reduced-form representation they identify is the structural representation if Assumption (2) is satisfied.

These properties help to clarify the sense in which the restrictions imposed in the structural step are a "credible subset" of the restrictions implied by the fully structural model. The first says that they are credible because they respect all the properties of the data, regardless of whether the structural model is correctly specified. This is because they identify a reduced-form representation, which is always consistent with the spectral density of observables by definition. The second says they are a subset of the restrictions implied by the structural representation, because the full set of restrictions may not be consistent with the spectral density of observables if Assumption (2) is not satisfied. This can happen when the structural model is both "over-identified" and incorrectly specified.

When is it possible to impose restrictions of this type? Only if the structural shocks are recoverable from the observables. To see why, suppose that they are not recoverable. In that case, it is not possible to identify the structural representation from the set of reduced-form representations even if the structural representation is correctly specified, and the full set of restrictions are imposed. But then there can be no hope of identifying the structural representation by imposing only a subset of those restrictions. We summarize this point in the following theorem.

Theorem 5. *Suppose that Assumption* (1) *is satisfied. Then there exists a matrix function* $\hat{\varphi}(\lambda) \in \mathcal{R}(f_y)$ *with the property that* $\hat{\varphi}(\lambda) = \varphi(\lambda)$ *if Assumption* (2) *is also satisfied only if the structural shocks* $\{\varepsilon_t\}$ *are recoverable from the observables* $\{y_t\}$*.*

Of course, the condition that the structural shocks are recoverable doesn't provide any direction regarding which specific subset of theoretical restrictions should be imposed in the structural step. In general, there might be several different subsets that could be chosen. This choice of restrictions is ultimately an economic one, and will depend on the details of the structural model(s) under consideration. A number of different choices have been made in the existing literature, one of which is described in the following example.

Example 4. Perhaps the most familiar restrictions are the "triangularization" restrictions from the original application in Sims (1980):

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

(ii) $E[y_{k,t}\overline{\varepsilon_{l,t}}] = 0$ for all $k < l$ and $E[y_{k,t}\overline{\varepsilon_{k,t}}] \geq 0$ for all k.

The first restriction says that the information sets generated by past structural shocks and past observables always coincide, which implies that the shocks are both recoverable and invertible.[15] The second imposes a recursive causal structure on the shocks, by allowing only the first shock to affect the first observable contemporaneously, only the first and second shock to affect the second observable contemporaneously, and so on.

In the case that the spectral density of observables has full rank for almost all λ, these restrictions are sufficient to uniquely select one representation from among the set of all empirically consistent reduced-form representations. That representation corresponds to the version of Wold's decomposition of $\{y_t\}$ with spectral characteristic $\hat{\varphi}(\lambda)$ such that $\hat{\varphi}_0 = \frac{1}{2\pi}\int \hat{\varphi}(\lambda)d\lambda$ is a lower-triangular matrix with real and positive diagonal elements. In other words, for any version of Wold's decomposition of $\{y_t\}$ with spectral characteristic $\gamma(\lambda)$, $\hat{\varphi}_0$ is uniquely determined by the lower-triangular Cholesky factorization of the matrix $\gamma_0\gamma_0^*$, given by

$$\hat{\varphi}_0\hat{\varphi}_0^* = \gamma_0\gamma_0^*,$$

where $\gamma_0 = \frac{1}{2\pi}\int \gamma(\lambda)d\lambda$, and $\hat{\varphi}(\lambda)$ is uniquely determined by the equation

$$\hat{\varphi}(\lambda) = \gamma(\lambda)(\hat{\varphi}_0^{-1}\gamma_0)^*.$$

The spectral characteristic $\hat{\varphi}(\lambda)$ is well-defined regardless of whether the (implicit) structural model is correctly specified. But under the null hypothesis that it is, it will be true that $\hat{\varphi}(\lambda) = \varphi(\lambda)$. $\diamond$

Once an estimate of the structural representation has been obtained, it is straightforward to compute many objects of economic interest, such as the shocks themselves, shock decompositions, impulse responses, and variance shares. For example, under the null hypothesis that the structural model is correctly specified and the shocks are recoverable, the values of the shock process $\{\varepsilon_{k,t}\}$ are given by

$$\varepsilon_{k,t} = \int e^{-i\lambda t}\delta_k\varphi(\lambda)^\dagger\Phi_y(d\lambda). \tag{12}$$

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,

$$\tilde{y}_{k,t}(l) = \int e^{i\lambda t}\varphi_{kl}(\lambda)\Phi_{\varepsilon,l}(d\lambda). \tag{13}$$

[15] Following Rozanov (1967), Ch. 2, this is sometimes described as an assumption that the structural shocks are "fundamental." However, this statistical sense of the word "fundamental" is not to be confused with its economic sense.

The response of $y_{k,t+s}$ to a unit impulse in the shock $\varepsilon_{l,t}$ is

$$\mathrm{IR}_{kl}(s) = \frac{1}{2\pi} \int e^{i\lambda s} \varphi_{kl}(\lambda) d\lambda, \tag{14}$$

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

$$\mathrm{VS}_{kl}(\mathit{\Delta}) \;= \int_{\mathit{\Delta}} |\varphi_{kl}(\lambda)|^2 d\lambda \left(\int_{\mathit{\Delta}} f_{y,kk}(\lambda) d\lambda \right)^{-1}. \tag{15}$$

Finally, it may be that the structural shocks are only partially recoverable (or only some of the shocks are of interest). In this case, it is still possible to perform a semi-structural analysis of these shocks in exactly the same manner. However, the theoretical restrictions imposed in the structural step only need to have the following properties:

(P1′) they identify at least one reduced-form representation in this set, regardless of whether Assumption (2) is satisfied, and

(P2′) the reduced-form representations they identify all agree with the parts of the structural representation involving the shock(s) of interest if Assumption (2) is satisfied.

The following theorem presents the natural generalization of Theorem (5) to the case of partial recoverability.

Theorem 6. *Suppose that Assumption* (1) *is satisfied. Then there exists a matrix function* $\hat{\varphi}(\lambda) \in \mathcal{R}(f_y)$ *with the property that* $\hat{\varphi}(\lambda)\delta_k^* = \varphi(\lambda)\delta_k^*$ *if Assumption* (2) *is also satisfied only if the structural shock process* $\{\varepsilon_{k,t}\}$ *is recoverable from the observables* $\{y_t\}$.

4.2 Why Invertibility is Not Necessary

In the previous subsection, we showed that recoverability is necessary for performing a Sims (1980)-style semi-structural analysis of structural shocks. In this subsection, we explain why invertibility is not necessary. The basic reason is that neither properties (P1) and (P2), nor their weaker versions (P1′) and (P2′), require that the structural shocks be invertible.

Perhaps the easiest way to see this is to consider imposing the same structural restrictions as in Example (4), but in the opposite time direction. These restrictions are clearly consistent with Sims (1980)'s proposed empirical strategy, but they are applicable across a class of structural models with non-invertible structural shocks. The following example describes how a semi-structural analysis can be performed with restrictions of this type. We provide an example with more explicit economic motivation in the following subsection.

Example 5. For any wide-sense stationary random process $\{\xi_t\}$, let $\mathcal{H}^t(\xi)$ denote the closed subspace of $\mathcal{H}(\xi)$ spanned by the variables $\xi_{k,s}$ for all k and $s \geq t$. Instead of the restrictions from Example (4), consider the following restrictions:

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

(ii) $E[y_{k,t}\overline{\varepsilon_{l,t}}] = 0$ for all $k < l$ and $E[y_{k,t}\overline{\varepsilon_{k,t}}] \geq 0$ for all k.

The first restriction says that the information sets generated by future structural shocks and future observables always coincide, which implies that the shocks are recoverable but not necessarily invertible. The second imposes a recursive causal structure on the shocks, by allowing only the first shock to affect the first observable contemporaneously, only the first and second shock to affect the second observable contemporaneously, and so on.

In the case that the spectral density of observables has full rank for almost all λ, these restrictions are sufficient to uniquely select one representation from among the set of all empirically consistent reduced-form representations. That representation can be obtained using Wold's decomposition of $\{y_t\}$. From any version Wold's decomposition of $\{y_t\}$ with spectral characteristic $\gamma(\lambda)$, $\hat{\varphi}_0$ is uniquely determined by the lower-triangular Cholesky factorization of the matrix $\gamma_0^*\gamma_0$, given by

$$\hat{\varphi}_0\hat{\varphi}_0^* = \gamma_0^*\gamma_0,$$

where $\gamma_0 = \frac{1}{2\pi}\int \gamma(\lambda)d\lambda$, and $\hat{\varphi}(\lambda)$ is uniquely determined by the equation

$$\hat{\varphi}(\lambda) = \gamma(\lambda)^*(\hat{\varphi}_0^{-1}\gamma_0^*)^*.$$

The spectral characteristic $\hat{\varphi}(\lambda)$ is well-defined regardless of whether the (implicit) structural model is correctly specified. But under the null hypothesis that it is, it will be true that $\hat{\varphi}(\lambda) = \varphi(\lambda)$. ◇

But if invertibility is not necessary, then why has so much of the literature seen invertibility as an essential part of Sims (1980)'s proposal? One possibility is that the literature has effectively adopted a "do as he does, not as he says" interpretation of Sims's proposal. In his application, he used an autoregression as the reduced-form model, and imposed the structural restrictions described in Example (4), the first of which implies that the structural shocks are invertible. By conflating the general proposal with the specific application, the subsequent literature has understood Sims's empirical strategy to be applicable only when (i) the reduced-form model is an autoregression, and (ii) the structural shocks satisfy $\mathcal{H}_t(\varepsilon) = \mathcal{H}_t(y)$ for all t. We will now analyze these two restrictions in more detail.

Even if the reduced-form model is an autoregression, that does not require the structural shocks to be invertible. As we explained in the previous subsection, the only purpose of the reduced-form model is to provide a characterization of the spectral density. The specific structural factorization that is chosen depends on the properties of the structural model, not the reduced-form model. Moreover, it is difficult to see how it could be advantageous to insist, a priori, that the statistical model always take the form of an autoregression. Why not entertain any number of possible statistical models, and only adopt an autoregression if there is good empirical evidence that such a model is empirically adequate? That way, the fact that the reduced-form model is an autoregression becomes a result, not an assumption.

The following example illustrates a simple case in which the reduced-form model is an autoregression and the structural shocks are not invertible, but it is straightforward to perform a semi-structural analysis of the shocks.

Example 6. Let $n_y = n_\varepsilon = 1$. In the reduced-form step, we use a first-order autoregressive model to characterize the spectral density of observables, consistent with the constraint that only autoregressive models may be estimated. Suppose that the reduced-form step delivers

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

where $|b| < 1$ and $\sigma_u = (E[|u_t|^2])^{1/2} > 0$. However, suppose that economic theory tells us that (i.e. our fully structural model implies that) $E[y_t \overline{\varepsilon_t}] \geq 0$ and $\mathcal{H}^t(\varepsilon) = \mathcal{H}^t(y)$ for all t. The unique reduced-form representation of $\{y_t\}$ consistent with these economic restrictions is

$$y_t = by_{t+1} + \sigma_u \varepsilon_t.$$

Therefore, we have $\hat{\varphi}(\lambda) = \sigma_u/(1 - be^{i\lambda})$. $\diamond$

While the restriction to autoregressive models in the reduced-form step is perfectly compatible with non-invertible structural shocks, the restriction that $\mathcal{H}_t(\varepsilon) = \mathcal{H}_t(y)$ is not.[16] However, this is 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. Of course, many papers since Hansen and Sargent (1991) have presented examples in which this is the case. The correct interpretation of these examples is not that they represent situations in which Sims' proposal is inapplicable; rather, they illustrate situations in which the restriction $\mathcal{H}_t(\varepsilon) = \mathcal{H}_t(y)$ should be replaced by something else. Therefore, we conclude that this restriction is not a necessary feature of Sims' proposal, and should only be imposed if it is implied by the structural model (or class of models) being entertained.

The fact that invertibility is not necessary for adopting Sims' semi-structural approach also helps to clarify the relationship between semi-structural analysis and other empirical approaches. For example, in their recent article on identification using external instruments, Stock and Watson (2018) provide a detailed description of "local projections instrumental variables" (LP-IV) analysis, which makes use of external sources of information about economic shocks to study how those shocks affect other macroeconomic variables. In describing the advantages of adopting this approach, they explain that "a major appeal of LP-IV is that the direct regression approach does not explicitly assume invertibility" (p.27).

But as we have argued, semi-structural analysis in the spirit of Sims (1980) does not need to assume invertibility either. Therefore, this really isn't a good reason to prefer one approach over the other. Of course, there may be *other* dimensions along which one approach has important advantages over the other. For instance, the LP-IV approach does not require a complete characterization of the spectral density of observables, which makes it more robust to potential misspecification of the reduced-form statistical model. Conversely, the semi-structural approach allows the researcher to answer a broader range of questions, and it does not require the researcher to have access to an external instrument. The point is that while there may be many good

[16] This restriction often appears in an equivalent form, $\varepsilon_t = \Theta(y_t - P(y_t|\mathcal{H}_{t-1}(y))$ for all t, where Θ is a non-singular matrix, and $P(y_t|\mathcal{H}_{t-1}(y))$ denotes the orthogonal projection of y_t onto the space $\mathcal{H}_{t-1}(y)$. See, for example, Fernández-Villaverde et al. (2007) equation (3), or Stock and Watson (2018) equation (17).

reasons to choose one approach over another, the invertibility or non-invertibility of the structural shocks of interest is not one of them.

4.3 Simple Economic Example

To tie together all the discussion so far, 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 7. An econometrician tries to recover labor income shocks $\{\varepsilon_t\}$ from observations of 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.[17] Combining the previous two equations with the definition of surplus income, it follows that

$$y_t \equiv s_t - s_{t-1} = \frac{1}{R}\varepsilon_t - \varepsilon_{t-1}. \tag{16}$$

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,

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

[17]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 $T \to \infty$.

More formally, the spectral characteristic linking the shocks to observables in equation (16) is

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

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 apply Thereom (2), 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 $\psi(\lambda) = \varphi(\lambda)^{-1}$. By Theorem (2), the income shocks are not invertible because the Fourier coefficient of $\psi(\lambda)\gamma(\lambda)$ for $s = -1$ does not vanish. In particular, for $R > 1$,

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

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:[18]

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

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

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

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

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

where $\gamma(\lambda)$ is the spectral characteristic from any version of Wold's decomposition of $\{y_t\}$. Since the Fourier coefficeints of the function $\hat{\gamma}(\lambda)$ vanish for $s < 0$, it follows that the Fourier coefficients of $\gamma(\lambda)^*$ 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 $\hat{\varphi}(\lambda) = R^{-1} - e^{-i\lambda}$. $\diamond$

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 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 noise shocks are one example, and we discuss these at length in the following section.

5 Noise shocks

The macroeconomic literature on noise shocks considers situations in which the beliefs of economic agents fluctuate for reasons entirely unrelated to the underlying economic fundamentals. Agents' beliefs fluctuate in this way because they receive imperfect signals about fundamentals, and must solve a signal extraction problem to form expectations about underlying outcomes. At the time that they make their decisions, the agents are unable to determine whether changes in their signals are due to actual fundamental developments or just unrelated noise. As a result, noise shocks can generate rational fluctuations in their expectations (and therefore also their actions) that nevertheless turn out to be incorrect after the fact.

We might say that the failure of non-invertibility in models with noise shocks is more severe than in other contexts, such as the permanent-income model we considered in previous sections. This is because, even if an econometrician has exactly the same date-t information as economic agents, he would still be unable to recover the structural shocks from the history of observables. If he could, then the agents could as well, which means they would be able to distinguish fundamental shocks from noise shocks, and would never respond to the latter. But then there would be no non-fundamental fluctuations in beliefs.

This line of reasoning, originally due to Blanchard et al. (2013), has lead a number of researchers to conclude that semi-structural methods cannot be applied to models with noise shocks.[19] The usual suggestion is that to make progress the econometrician must rely more heavily on his theoretical model by adopting a fully structural empirical approach. However, these conclusions rest 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.

In this section we describe how semi-structural methods can be applied to recover noise and fundamental shocks. We first describe a simple bivariate model of consumption determination taken from Blanchard et al. (2013) with noise and fundamental shocks. Then we explain how to perform a semi-structural analysis of these shocks. We verify our results through a Monte Carlo simulation study. Then we apply exactly the same empirical procedure on an actual sample of U.S. data to quantify the importance of non-TFP fluctuations in aggregate consumption from 1984:Q1-2016:Q4.

Example 8. *Model*: At each date, consumption is equal to agents' long-run forecast of total factor productivity,

$$c_t = \lim_{j \to \infty} E_t[a_{t+j}]. \tag{17}$$

This forecast is made conditional on the current and past history of productivity and

[19] 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).

signals about future productivity, a_τ and s_τ for $\tau \leq t$. Productivity is a random walk,

$$a_t = a_{t-1} + \sigma_a \varepsilon_t^a, \tag{18}$$

and the signal about future productivity is given by

$$s_t = \left(\frac{1-\rho}{1+\rho}\right) \sum_{j=-\infty}^{\infty} \rho^{|j|} a_{t-j} + v_t. \tag{19}$$

The parameter $\rho \in (0, 1)$ 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-fundamental noise, and is assumed to follow a law of motion of the form

$$v_t = 2\rho v_{t-1} - \rho^2 v_{t-2} + \sigma_v \varepsilon_t^v - (\beta + \bar{\beta})\sigma_v \varepsilon_{t-1}^v + \beta\bar{\beta}\sigma_v \varepsilon_{t-2}^v. \tag{20}$$

The vector of fundamental 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.[20]

Because $\{a_t\}$ is not stationary in the wide sense, we need to clarify the precise meaning of the forecast in equation (17). Following the discussion in Remark (3), we let $\{\tilde{\xi}_t(\theta)\}$ for $|\theta| < 1$ denote the stationary counterpart to any process $\{\xi_t\}$ that is stationary only after suitable differencing, and $\mathcal{H}(\tilde{\xi})$ the Hilbert space generated by its values. In this example, by letting $q_t \equiv (a_t, s_t)'$ we can understand

$$E_t[a_{t+j}] \equiv \lim_{\theta \to 1^-} E_t[\tilde{a}_{t+j}(\theta)],$$

where the conditional expectation on the right side is the linear projection of $\tilde{a}_{t+j}(\theta)$ onto $\mathcal{H}_t(\tilde{q})$. To illustrate, in the case that the signal is completely redundant ($\rho = 0$), this would mean that

$$E_t[a_{t+j}] = \lim_{\theta \to 1^-} \theta^j \tilde{a}_t(\theta) = a_t.$$

[20] Blanchard et al. (2013) write the information structure in this alternative but observationally equivalent way. For more details on the mapping from their representation to the noise representation presented above, including the nonlinear parameter restriction, see Chahrour and Jurado (2017).

Recoverability: We can now show that the structural shocks $\{\varepsilon_t\}$ are recoverable with respect to $\{y_t\}$, where $y_t \equiv (a_t, c_t)'$. To see this, first note that it is sufficient to establish these results with respect to $\{q_t\}$, since $\mathcal{H}_t(\tilde{y}) = \mathcal{H}_t(\tilde{q})$ whenever the signal is not redundant. According to equations (18) and (19), $\{\tilde{q}_t(\theta)\}$ can be obtained from $\{\varepsilon_t\}$ by a linear transformation with spectral characteristic

$$\varphi(\lambda;\theta) = \frac{1}{1-\theta e^{-i\lambda}} \begin{bmatrix} \sigma_a & 0 \\ \\ \dfrac{(1-\rho)^2\sigma_a}{|1-\rho e^{-i\lambda}|^2} & \dfrac{(1-\beta e^{-i\lambda})(1-\bar{\beta} e^{-i\lambda})(1-e^{-i\lambda})\sigma_v}{(1-\rho e^{-i\lambda})^2} \end{bmatrix}.$$

Here we have used the fact that for any integer s,

$$\frac{1}{2\pi}\int_{-\pi}^{\pi} e^{i\lambda s}\frac{(1-\rho)^2}{|1-\rho e^{-i\lambda}|^2}d\lambda = \left(\frac{1-\rho}{1+\rho}\right)\rho^{|s|}.$$

It is easy to see that $\varphi(\lambda;\theta)$ has full rank for almost all $\lambda \in [-\pi,\pi]$ and $\theta \in [0,1)$ whenever $\sigma_a, \sigma_v > 0$. By Theorem (1), this means that the structural shocks are recoverable with respect to $\{\tilde{y}_t(\theta)\}$ for almost all θ. Using the terminology introduced in Remark (3), we conclude that the shocks are recoverable from $\{y_t\}$.

Structural step: Now we illustrate how semi-structural methods can be applied to recover the noise and fundamental shocks from observations of productivity and consumption. As in Example (2), we first suppose that the econometrician has some characterization of the spectral density of the stationary observable process $\{\Delta y_t\}$, $f_{\Delta y}(\lambda)$. The structural step involves factoring the spectral density as

$$f_{\Delta y}(\lambda) = \frac{1}{2\pi}\hat{\varphi}(\lambda)\hat{\varphi}(\lambda)^*, \tag{21}$$

where the factor $\hat{\varphi}(\lambda)$ is defined by a set of theoretical restrictions that are sufficient to correctly identify the structural shocks in the model. One such set is[21]

(i) $\mathcal{H}_t(\varepsilon^a) = \mathcal{H}_t(\Delta a), \quad E[\Delta a_t \varepsilon^a_t] \geq 0$

(ii) $\mathcal{H}_t(\varepsilon^v) = \mathcal{H}_t(\zeta), \quad E[\zeta_t \varepsilon^v_t] \geq 0,$

where $\{\zeta_t\}$ is the unique process composed of the orthogonal projections of Δc_t onto the space $\mathcal{H}(y)$ for each t. The first restriction says that the fundamental shock comes

[21] These restrictions have also been used by Forni et al. (2017b), although their empirical application does not impose them in the same way we do in the following subsection.

from the unique Wold decomposition of $\{\Delta a_t\}$ with $E[\Delta a_t \varepsilon_t^a] \geq 0$. The second says that the noise shock captures the fluctuations in current consumption growth that are orthogonal to productivity growth across all time periods.

These restrictions imply that $\hat{\varphi}(\lambda)$ has a lower-triangular form

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

Alternatively, in terms of the associated moving average representation, that

$$\begin{bmatrix} \Delta a_t \\ \Delta c_t \end{bmatrix} = \cdots + \underbrace{\begin{bmatrix} 0 & 0 \\ * & 0 \end{bmatrix}}_{\hat{\varphi}_{-1}} \begin{bmatrix} \varepsilon_{t+1}^a \\ \varepsilon_{t+1}^v \end{bmatrix} + \underbrace{\begin{bmatrix} * & 0 \\ * & * \end{bmatrix}}_{\hat{\varphi}_0} \begin{bmatrix} \varepsilon_t^a \\ \varepsilon_t^v \end{bmatrix} + \underbrace{\begin{bmatrix} * & 0 \\ * & * \end{bmatrix}}_{\hat{\varphi}_1} \begin{bmatrix} \varepsilon_{t-1}^a \\ \varepsilon_{t-1}^v \end{bmatrix} + \cdots,$$

where $\{\hat{\varphi}_s\}$ are the sequence of Fourier coefficients associated with $\hat{\varphi}(\lambda)$.

To obtain the factor $\hat{\varphi}(\lambda)$, we can write equation (21) out more explicitly, using equation (22), as

$$\begin{bmatrix} \hat{f}_{\Delta a}(\lambda) & \hat{f}_{\Delta a \Delta c}(\lambda) \\ \hat{f}_{\Delta c \Delta a}(\lambda) & \hat{f}_{\Delta c}(\lambda) \end{bmatrix} = \frac{1}{2\pi} \begin{bmatrix} |\hat{\varphi}_{11}(\lambda)|^2 & \hat{\varphi}_{11}(\lambda)\overline{\hat{\varphi}_{21}(\lambda)} \\ \hat{\varphi}_{11}(\lambda)\hat{\varphi}_{21}(\lambda) & |\hat{\varphi}_{22}(\lambda)|^2 + |\hat{\varphi}_{21}(\lambda)|^2 \end{bmatrix}. \tag{23}$$

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 (23) uniquely determines $\hat{\varphi}_{21}(\lambda)$ as a function of $\hat{f}_{\Delta c \Delta a}(\lambda)$ and $\hat{\varphi}_{11}(\lambda)$, the first of which is given and the second of which has already been determined from the upper-left equation. The lower-right equation in (23) implies that

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

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

Reduced-form step: Lastly, we need to describe the statistical model the econometrician uses to construct his estimate of the spectral density. Of course, there are many possibilities. To be consistent with the model we use in the following subsections, let us suppose the econometrician uses an autoregressive model of the form

$$\sum_{s=0}^{p} b_s y_{t-s} = u_t, \tag{24}$$

where $\{u_t\}$ is a two-dimensional uncorrelated reduced-form shock process with zero mean and mutually uncorrelated values, and p is chosen to maximize some measure of empirical fit. If we define

$$b(\lambda) \equiv \sum_{s=0}^{p} b_s e^{-i\lambda s},$$

then the assumption that $\{y_t\}$ is difference stationary implies that the rows of $(1 - e^{-i\lambda})b(\lambda)$ must each be square integrable. Therefore, the econometrician's estimate of the spectral density of $\{\Delta y_t\}$ is

$$f_{\Delta y}(\lambda) = \frac{1}{2\pi}(1 - e^{-i\lambda})b(\lambda)b(\lambda)^*(1 - e^{i\lambda}).$$

With this estimate, he can proceed to perform the factorization described in the structural step. ◇

5.1 A Monte Carlo Study

To demonstrate how semi-structural methods can be applied in practice to models with noise shocks, we perform a Monte Carlo exercise using the model from Example (3). The exercise entails simulating data on consumption and productivity from the model, and placing ourselves in the shoes 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 noise shocks and the effects of a noise shock on consumption from that sample. To do so, he imposes only the structural restrictions (i) and (ii) from Example (8).

In practice, we simulate $N = 1000$ samples of $T = 275$ observations of consumption and productivity from the model. The structural parameters are set to

$$\rho = 0.8910, \quad \sigma_a = 0.6700, \quad \sigma_v = 0.9937, \quad \text{and} \quad \beta = 0.7833 - 0.1525i,$$

which correspond to the same parameters chosen by Blanchard et al. (2013). The reduced-form model is an unrestricted vector autoregression of the type in equation (24). We fit the model to the data using the multivariate algorithm of Morf et al. (1978), and the lag length is chosen to minimize the information criterion proposed in Hannan and Quinn (1979).

The left panel of Figure (1) plots the true impulse response of consumption to a noise shock that increases consumption by one unit on impact, together with 95%

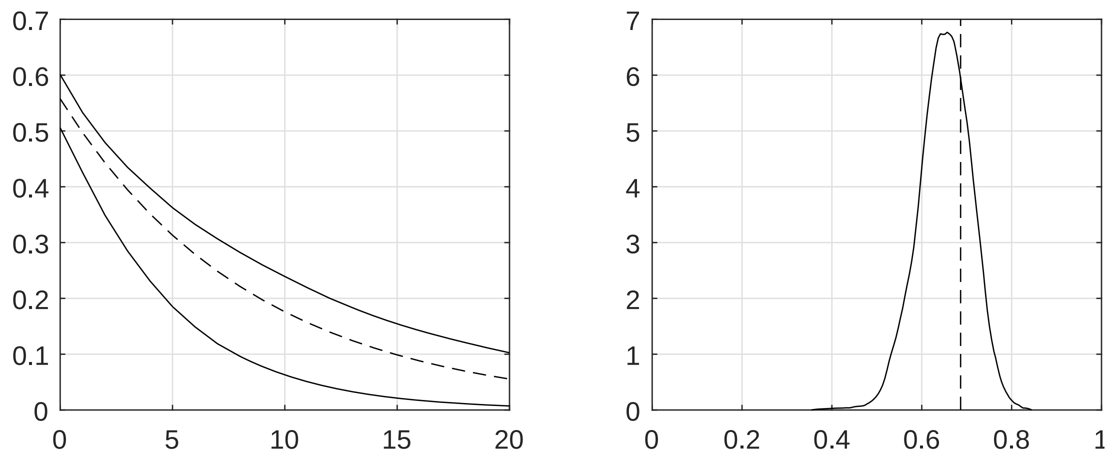

Figure 1: Semi-structural analysis of data simulated from a model with noise shocks. *Left*: the dashed line is the true impulse response of consumption to a unit noise shock, while solid lines are 95% bands from the distribution of point estimates from each of $N = 1000$ samples of length $T = 275$. *Right*: the dashed line is the true contribution of noise shocks over business-cycle frequencies (6 to 32 quarters), and solid line is the distribution of point estimates over all simulated samples.

bands constructed from the point estimates across the N different samples. The true response of consumption is one of geometric decay; initially consumption increases due to positive expectations about future productivity, but over time those effects die out as people come to realize that their expectations had only responded to noise. In the long run, the effect of noise shocks on consumption converges to zero. The figure indicates that structural VAR analysis does a good job capturing the response of consumption to a noise shock, even for samples of $T = 275$ observations. Not surprisingly, increasing the sample size increases the accuracy of our estimates.

The right panel of Figure (1) plots the share of the variance in consumption explained by noise shocks over business cycle frequencies (6 to 32 quarters). The vertical dashed line is the true noise share (0.69), while the solid line is the histogram of point estimates from each of the N different samples. Again, the structural VAR procedure evidently delivers accurate estimates of the importance of noise shocks. Based on the distribution of point estimates, it appears that the estimated importance of noise shocks does exhibit some slight downward bias due to the fact that the sample is finite. A slight downward bias in this estimate would only strengthen the conclusions we reach in the next section.

5.2 Application to U.S. Data

In this subsection, we apply the same semi-structural procedure used in our Monte Carlo study 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 seires LNU00000000Q) and productivity by the natural logarithm of utilization adjusted total factor productivity (Federal Reserve Bank of San Francisco). 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 model from Example (3), productivity is the only fundamental process, and agents have 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 plausible that consumption is driven by fundamentals other than productivity, by sunspots, or even by non-rational fluctuations in people's beliefs. Therefore, 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 (2). The left panel plots the estimated impulse response of consumption to a noise shock that increases consumption by one unit on impact. The response is hump-shaped, increasing for six quarters after the shock, and then slowly decaying back toward zero. The effect of noise shocks is also highly persistent; even after 20 quarters the response is still statistically different from zero. To the extent that these shocks do represent rational mistakes due to imperfect signals, the high persistence means that it takes a while for people to recognize their errors.

The right panel of Figure (2) plots the share of the variance in consumption explained by noise shocks over business cycle frequencies (6 to 32 quarters). The vertical dashed line is our point estimate (0.86), while the solid line is the histogram of point estimates across $N = 1000$ bootstrap samples. The point estimate indicates that productivity only explains 14% of the variation in consumption. Evidently a large majority of consumption fluctuations are not due to productivity shocks.

Cochrane (1994) reaches a similar conclusion. Using semi-structural methods, he

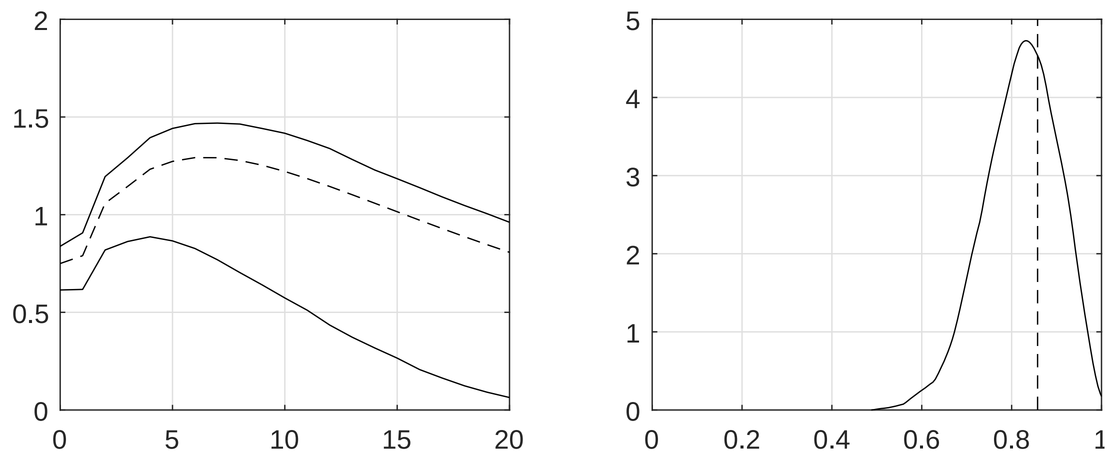

Figure 2: Semi-structural analysis of quarterly U.S. consumption and total factor productivity from 1948:Q1 to 2016:Q4. *Left*: response of consumption to unit noise shock. The dashed line is the point estimate, and the solid lines are 95% bootstrap confidence bands. *Right*: share of consumption variance due to noise shocks over business-cycle frequencies (6 to 32 quarters). The dashed line is the point estimate (0.86) and the solid line is the distribution of bootstrap estimates.

argues that the bulk of economic fluctuations is not due to productivity shocks (or a number of other shocks including those due to monetary policy, oil prices, and credit). But, he does not control for the possibility that fluctuations might be due to *future* changes in productivity to which people respond in advance. Indeed, he suggests that fundamentals might matter mainly in this way. Here we provide evidence to the contrary, at least in the case of total factor productivity. While people's beliefs about future productivity may be moving around a lot, it appears either that those movements are mostly unrelated to subsequent changes in productivity, or that people's beliefs about future productivity do not matter very much for their current actions.

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 aren't really difficulties at all. Nothing in the original empirical strategy of Sims (1980) 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 empirical methods to find greater applicability across a broader class of economically interesting models.

There are a number of practical issues that we have not addressed in this paper. Foremost among them is probably the task of characterizing precisely what constitutes a "good" reduced-form model. Undoubtedly this will vary on a case-by-case basis, but perhaps it is possible to say something about which reduced-form models are likely to deliver better or worse approximations to the relevant features of the spectral density. Such guidance could be helpful for "fine-tuning" one's empirical strategy. A solution would likely involve relying on additional theoretical restrictions to rule out certain types of reduced-form models and not others.

Our application to data on U.S. consumption and productivity also invites a more comprehensive investigation. How important are other fundamentals, like monetary policy shocks, oil price shocks, credit shocks, or government spending shocks? What about other macroeconomic variables of interest like output, inflation, or unemployment? The empirical procedure we used in this paper can be helpful for determining the importance of a any set of observable fundamental processes. Since our main purpose in this paper is to clarify the difference between invertibility and recoverability, we reserve such an investigation for future research.

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.

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 (2017). News or Noise? The Missing Link. Working paper, Duke University.

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.

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.

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.

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.

Stock, J. H. and M. W. Watson (2018). Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments. Working paper, Harvard University.

Appendix

Proof of Lemma (1)**.** First, we observe that $\mathcal{H}(\xi)$ is isomorphic to $\mathcal{L}^2(F_\xi)$.[22] This can be seen by defining a correspondence between elements $h \in \mathcal{H}(\xi)$ of the form

$$h = \int \psi(\lambda)\Phi_\xi(d\lambda), \tag{25}$$

where

$$\int |\psi_k(\lambda)|^2 F_{\xi,kk}(d\lambda) < \infty, \quad k = 1, \dots, n_\xi, \tag{26}$$

and the vector functions $\psi(\lambda) \in \mathcal{L}^2(F_\xi)$ which occur in the representation (25). This correspondence is linear, since $h_1 \leftrightarrow \psi_1$ and $h_2 \leftrightarrow \psi_2$ implies

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

[22] 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.

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 (25) coincides with $\mathcal{H}(\xi)$, and the closed linear manifold spanned by elements $\psi(\lambda)$ of the form (26) 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.[23]

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 $k = 1, \ldots, n_\eta$,

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

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

[23] See Rozanov (1967), Ch. 1.

Proof of Theorem (1).[24] 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[25]

$$\Phi_y(d\lambda) = \varphi(\lambda)\Phi_\varepsilon(d\lambda). \tag{27}$$

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

$$\psi(\lambda)\varphi(\lambda) = I_{n_\varepsilon}. \tag{28}$$

Combining equations (27) and (28), 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 $k = 1, \dots, n_\varepsilon$, equations (27) and (28) 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. 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 $k = 1, \dots, n_y$ 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 $\mathcal{H}(\varepsilon) \subseteq \mathcal{H}(y)$. □

[24] This proof comes from Rozanov (1967), Ch. 1.

[25] More precisely, equation (27) means that $\Phi_y(\Delta) = \int_\Delta \psi(\lambda)\Phi_\varepsilon(d\lambda)$ for any Borel set Δ.

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

$$h = \sum_{j=0}^{\infty} \alpha_j w_{t-j} \tag{29}$$

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),

$$\varepsilon_t = \int e^{i\lambda t} \psi(\lambda) \Phi_y(d\lambda) = \int e^{i\lambda t} \psi(\lambda) \delta(\lambda) \Phi_w(d\lambda) \tag{30}$$

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 (30), we can see that the elements of ε_t have a representation of the form (29) 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

$$\varphi(\lambda)^* = \psi(\lambda)\varphi(\lambda)\varphi(\lambda)^* \tag{31}$$

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 (31). 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. □

Proof of Theorem (5)**.** Because it is the spectral characteristic associated with a reduced-form representation, the matrix function $\hat{\varphi}(\lambda)$ has full column rank for almost all λ. The fact that $\hat{\varphi}(\lambda) = \varphi(\lambda)$ under Assumptions (1) and (2) means that $\varphi(\lambda)$ is full column rank for almost all λ under these assumptions. By Theorem (1), it follows that the structural shocks are recoverable. □

Proof of Theorem (6)**.** Let $\{\hat{\varepsilon}_t\}$ denote the shock process from the reduced-form representation of $\{y_t\}$ associated with the spectral measure $\hat{\varphi}(\lambda)$. Let $\hat{\psi}(\lambda)$ denote any $r_y \times n_y$ matrix function that satisfies $\hat{\psi}(\lambda)\hat{\varphi}(\lambda) = I_{r_y}$ for almost all λ. Such a matrix function always exists because, by definition, $\hat{\varphi}(\lambda)$ is full column rank for almost all λ. Under Assumptions (1) and (2), it follows that

$$\begin{aligned}\hat{\varepsilon}_t &= \int e^{i\lambda t} \hat{\psi}(\lambda)\varphi(\lambda)\Phi_\varepsilon(d\lambda) \\ &= \int e^{i\lambda t} \left[\hat{\psi}(\lambda)\varphi(\lambda)\delta_k^* \Phi_{\varepsilon,k}(d\lambda) + \sum_{l \neq k} \hat{\psi}(\lambda)\varphi(\lambda)\delta_l^* \Phi_{\varepsilon,l}(d\lambda) \right].\end{aligned}$$

Using the hypothesis that $\varphi(\lambda)\delta_k^* = \hat{\varphi}(\lambda)\delta_k^*$ for almost all λ,

$$E[\hat{\varepsilon}_t \overline{\varepsilon}_{k,t}] = \frac{1}{2\pi} \int \hat{\psi}(\lambda)\varphi(\lambda)\delta_k^* d\lambda = \delta_k^*,$$

which means that $\hat{\varepsilon}_{k,t} = \varepsilon_{k,t}$ for all k, t. The recoverability of $\{\varepsilon_{k,t}\}$ then follows from the recoverability of $\{\hat{\varepsilon}_{k,t}\}$. □

Appendix for Online Publication

A 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

$$y_t = \int_{-\infty}^{\infty} \varphi(\lambda) \Phi_\varepsilon(d\lambda). \tag{32}$$

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 $-\pi, \pi$.[26] 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.,

$$E[\Phi_{\varepsilon,k}(\mathit{\Delta}_1)\overline{\Phi_{\varepsilon,l}(\mathit{\Delta}_2)}] = 0$$

if $k \neq l$, for any Borel sets $\mathit{\Delta}_1$ and $\mathit{\Delta}_2$ 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 shocks to be a collection of uncorrelated random measures $\varepsilon_k(dt)$, $k = 1, \ldots, n_\varepsilon$, defined as the Fourier transforms of the measures $\Phi_{\varepsilon,k}(d\lambda)$, $k = 1, \ldots, n_\varepsilon$. That is, we set

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

for any semi-interval $\mathit{\Delta} = (t_1, t_2]$, 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*

$$y_t = \int e^{i\lambda t} \varphi(\lambda) \Phi_\varepsilon(d\lambda) \quad \textit{for all } t, \tag{33}$$

where $E|\varepsilon_k(dt)|^2] = dt$ for all k, and $E[\varepsilon_k(\mathit{\Delta}_1)\overline{\varepsilon_l(\mathit{\Delta}_2)}] = 0$ for $k \neq l$ and any Borel sets $\mathit{\Delta}_1$ and $\mathit{\Delta}_2$.

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

Example 9. As a special case of the model in equation (33), 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 $k = 1, \ldots, n_\varepsilon$. 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

$$\tilde{\eta}_t = \int e^{i\lambda t} \tilde{\psi}(\lambda) \Phi_y(d\lambda), \tag{34}$$

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 (33) 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*

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

for almost all λ, 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 (34) 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*

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

for almost all λ, 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 7. 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 $s < 0$.