Recoverability*

Ryan Chahrour
Boston College

Kyle Jurado
Duke University

November 1, 2017

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, C31
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 usually 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 often entails 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 parametric 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. In a paper closely related to our empirical application, Forni et al. (2017) 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. Our contribution is to point out that it is recoverability, not invertibility, that really matters for empirical work, and to provide a simple but general condition to check whether recoverability is satisfied.

2 Recoverability Condition

This section presents our main theorem. We begin with some notation and definitions. For an arbitrary $n_\xi \times 1$ dimensional covariance stationary vector process $\{\xi_t\}$, we let $\mathcal{H}(\xi)$ denote the Hilbert space spanned by the variables $\xi_{k,t}$ for $k = 1, \ldots, n_\xi$ and $t \in \mathbb{Z}$, closed with respect to convergence in mean square. Similarly, we let $\mathcal{H}_t(\xi)$

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

denote the space spanned by these variables over all k but only up through date t. This is enough for us to define what we mean by recoverability.

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

$$\mathcal{H}_t(\eta) \subseteq \mathcal{H}_t(\xi) \quad \text{for all } t \in \mathbb{Z}.$$

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

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

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

where Φ_ξ is its associated random spectral measure. We say that a $1 \times n_\xi$ dimensional vector function $\psi(\lambda)$, defined for $\lambda \in [-\pi, \pi]$, belongs to the space $\mathcal{L}^2(F_\xi)$ if

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

In this expression, F_ξ denotes the spectral measure of $\{\xi_t\}$ and the asterisk denotes complex conjugate transposition.[6] If we define the scalar product

$$(\psi_1, \psi_2) = \int_{-\pi}^{\pi} \psi_1(\lambda) F_\xi(d\lambda) \psi_2(\lambda)^*,$$

[6]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 in $[-\pi, \pi]$.

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.

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

$$\eta_t = \int_{-\pi}^{\pi} e^{i\lambda t}\varphi(\lambda)\Phi_\xi(d\lambda) \quad \textit{for all } t \in \mathbb{Z}. \tag{1}$$

Proof. Define the operator T such that $Th = \psi$, where $h \in \mathcal{H}(\xi)$ is an arbitrary element of the form

$$h = \int_{-\pi}^{\pi} \psi(\lambda)\Phi_\xi(d\lambda), \tag{2}$$

and $\int_{-\pi}^{\pi} |\psi_k(\lambda)|^2 F_{\xi,kk}(d\lambda) < \infty$ for $k = 1, \ldots, n_\xi$. This operator can be extended to an isometric mapping of $\mathcal{H}(\xi)$ onto $\mathcal{L}^2(F_\xi)$.[7] Therefore, $h \in \mathcal{H}(\xi)$ if and only if there exists some $\psi \in \mathcal{L}^2(F_\xi)$ such that equation (2) holds. By introducing a scale factor $e^{i\lambda t}$ for each t, it follows that $\eta_{k,t} \in \mathcal{H}(\xi)$ for $k = 1, \ldots, n_\eta$ if and only if there exists some $\varphi_k(\lambda) \in \mathcal{L}^2(F_\xi)$ for $k = 1, \ldots, n_\eta$ such that

$$\eta_{k,t} = \int_{-\pi}^{\pi} e^{i\lambda t}\varphi_k(\lambda)\Phi_\xi(d\lambda) \quad \text{for } 1, \ldots, n_\eta.$$

Since $\eta_{k,t} \in \mathcal{H}(\xi)$ for all k, t if and only if $\mathcal{H}(\eta) \subseteq \mathcal{H}(\xi)$, we can define the $n_\eta \times n_\xi$ function φ by stacking the vectors φ_k one on top of the other. □

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 (1), and we will call φ 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 variables that are observable to an outside economist. We suppose that an n_y-dimensional observable process $\{y_t\}$ can be obtained from an n_ε-dimensional structural shock process $\{\varepsilon_t\}$ by a linear transformation of the form

$$y_t = \int_{-\pi}^{\pi} e^{i\lambda t}\varphi(\lambda)\Phi_\varepsilon(d\lambda). \tag{3}$$

[7] See Rozanov (1967), Chapter 1.

The process $\{y_t\}$ is covariance stationary and linearly regular, and the structural shocks are uncorrelated over time and normalized to have mean zero and an identity covariance matrix, I_{n_ε}.[8]

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

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

where x_t is an n_x-dimensional state vector. In this case, the spectral characteristic $\varphi(\lambda)$ 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.[9] $\Diamond$

By Lemma (1), the model in equation (3) says that the observables are recoverable with respect to the structural shocks. Naturally, knowledge of the inputs of the system is enough to perfectly reveal the outputs. Our question is: when can the shocks be recovered from the observables? The following theorem provides the answer.[10]

Theorem 1. *The structural shocks $\{\varepsilon_t\}$ are recoverable from the observables $\{y_t\}$ if and only if*

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

for almost all $\lambda \in [-\pi, \pi]$.

Proof. Sufficiency: $\{y_t\}$ 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[11]

$$\Phi_y(d\lambda) = \varphi(\lambda)\Phi_\varepsilon(d\lambda). \tag{6}$$

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

$$\psi(\lambda)\varphi(\lambda) = I_{n_\varepsilon}. \tag{7}$$

[8] Linear regularity means that $\bigcap_{t=-\infty}^{\infty} \mathcal{H}_t(y) = 0$.

[9] 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.

[10] The proof comes from Rozanov (1967), Chapter 1.

[11] More precisely, equation (6) means that $\Phi_y(\Delta) = \int_\Delta \psi(\lambda)\Phi_\varepsilon(d\lambda)$ for any Borel set Δ in $[-\pi, \pi]$.

Combining equations (6) and (7), we get

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

Moreover, note that the rows of ψ are elements of $\mathcal{L}^2(F_y)$ because for any $k = 1, \ldots, n_\varepsilon$, equations (6) and (7) imply that

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

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 $\psi \in \mathcal{L}^2(F_\varepsilon)$ such that $\|\psi(\lambda)\| \neq 0$ and

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

for all $\lambda \in [-\pi, \pi]$. This would mean that the element

$$\eta = \int_{-\pi}^{\pi} \psi(\lambda)\Phi_\varepsilon(d\lambda)$$

is orthogonal to $\mathcal{H}(y)$, because, for all $k = 1, \ldots, n_y$ and $t \in \mathbb{Z}$,

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

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

Before moving on, a couple of remarks are in order.

Remark 1. In the special case from Example (1), the condition in the theorem is equivalent to the condition that the matrix

$$\mathcal{R}(\lambda) \equiv \begin{bmatrix} I_{n_x} - Be^{-i\lambda} & C \\ -A & 0_{n_y \times n_\varepsilon} \end{bmatrix} \tag{8}$$

be full column rank for almost all $\lambda \in [-\pi, \pi]$. This follows from the so-called Guttman rank additivity formula. Specifying the condition in terms of $\mathcal{R}(\lambda)$ has the advantage that it does not involve any matrix inverses, and may be more efficient to check on a computer. To do so, we can draw a random number λ_u from the uniform distribution over $[-\pi, \pi]$ and check whether $\mathcal{R}(\lambda_u)$ is full column rank.

Remark 2. 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 stochastic processes.

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. Despite all the attention invertibility has received in the literature, it does not seem that conditions of this type have been articulated.[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 $\text{rank}(\varphi(\lambda)) = n_\varepsilon$ as we look for the additional restrictions that are needed.

The first step is to recall that, using Wold's decomposition theorem, it is possible to represent $\{y_t\}$ by a linear transformation of the form

$$y_t = \int_{-\pi}^{\pi} e^{i\lambda t} \delta(\lambda) \Phi_w(d\lambda), \tag{9}$$

where Φ_w is the random spectral measure associated with an uncorrelated process $\{w_t\}$ with spectral density $f_w(\lambda) = \frac{1}{2\pi} I_{n_\varepsilon}$. This uncorrelated process has the property that w_s for $s \leq t$ form a basis in $\mathcal{H}_t(y)$ at each date, so that $\mathcal{H}_t(y) = \mathcal{H}_t(w)$. This implies that $\{w_t\}$ is both invertible and recoverable from $\{y_t\}$.

Using the spectral characteristic $\delta(\lambda)$ from equation (9) and the function $\psi(\lambda)$ defined in equation (7), we can state the following result.

Theorem 2. *The structural shocks $\{\varepsilon_t\}$ are invertible from the observables $\{y_t\}$ if and only if they are recoverable and*

$$\frac{1}{2\pi} \int_{-\pi}^{\pi} e^{i\lambda j} \psi(\lambda) \delta(\lambda) d\lambda = 0$$

for all integers $j < 0$.

[12] There are places where sufficient conditions appear, however. The condition of Fernández-Villaverde et al. (2007) is one example.

Proof. The fact that w_s, $s \leq t$ forms 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

$$h = \sum_{j=0}^{\infty} \alpha_j w_{t-j} \tag{10}$$

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 equations (7) and (9),

$$\varepsilon_t = \int_{-\pi}^{\pi} e^{i\lambda t} \psi(\lambda) \Phi_y(d\lambda) = \int_{-\pi}^{\pi} e^{i\lambda t} \psi(\lambda)\delta(\lambda) \Phi_w(d\lambda). \tag{11}$$

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 $[-\pi, \pi]$. 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

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

Combining this with equation (11), we can see that the elements of ε_t have a representation of the form (10) if and only if the Fourier coefficients $\{\alpha_j\}$ vanish for negative values of j, which is the condition stated in the theorem. □

We now illustrate our recoverability condition in the context of a simple permanent income model of consumption. 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.

Example 2. An econometrician tries to recover labor income shocks $\{\varepsilon_t\}$ from observations of surplus income, $s_t \equiv y_t - c_t$, where c_t is date-t consumption and y_t is date-t labor income, which satisfies

$$y_t = \sigma \varepsilon_t, \quad \varepsilon_t \overset{\text{iid}}{\sim} \mathcal{N}(0, 1), \tag{12}$$

with $\sigma > 0$. The optimal path for consumption is a random walk

$$c_t = c_{t-1} + \left(1 - \frac{1}{R}\right) \sigma \varepsilon_t, \tag{13}$$

where $R > 1$ is the constant gross real interest rate.[13] Combining equations (12) and

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

(13) with the definition of surplus income, it follows that

$$\Delta s_t = \frac{1}{R}\sigma\varepsilon_t - \sigma\varepsilon_{t-1}, \tag{14}$$

where Δ denotes the first-difference operator, $\Delta s_t \equiv s_t - s_{t-1}$. Therefore, the change in surplus income follows a first-order moving average process.

The spectral characteristic linking the shocks to observables is

$$\varphi(\lambda) = \left(\frac{1}{R} - e^{-i\lambda}\right)\sigma. \tag{15}$$

It is easy to see that $\varphi(\lambda)$ has rank equal to 1 everywhere on $[-\pi, \pi]$ except for at $\lambda_0 = -\ln(1/R)/i$. Therefore by Theorem (1) the shocks are recoverable from observations on surplus income only.

To apply Theorem (2), we first use the spectral characteristic in equation (15) to solve for $\psi(\lambda)$ from equation (7). In this case, we simply have $\psi(\lambda) = 1/\varphi(\lambda)$. Next, the Wold representation of $\{\Delta s_t\}$ is

$$\Delta s_t = \sigma w_t - \frac{1}{R}\sigma w_{t-1},$$

which means that the spectral characteristic in equation (11) is given by

$$\theta(\lambda) = \left(1 - \frac{1}{R}e^{-i\lambda}\right)\sigma.$$

Multiplying $\psi(\lambda)$ and $\theta(\lambda)$, we have

$$\alpha(\lambda) = \frac{R - e^{-i\lambda}}{1 - Re^{-i\lambda}}. \tag{16}$$

However, the Fourier coefficient of $\alpha(\lambda)$ for $j = -1$ does not vanish, since for $R > 1$,

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

Therefore, according to Theorem (2), the shocks are not invertible. $\Diamond$

Remark 3. In this section we have focused exclusively on covariance-stationary processes; however, the discussion 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,

$$\Delta^p \xi_t = \zeta_t \tag{17}$$

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

$$\tilde{\xi}_t(\theta) \equiv \int_{-\pi}^{\pi} e^{i\lambda t} \frac{1}{(1-\theta e^{-i\lambda})^p} \Phi_\zeta(d\lambda), \tag{18}$$

which is stationary for each value of $\theta \in [0, 1)$. We can say that a process $\{\eta_t\}$ is recoverable (or invertible) with respect to $\{\xi_t\}$ whenever $\{\eta_t\}$ is recoverable (or invertible) with respect to $\{\tilde{\xi}_t(\theta)\}$ for almost all $\theta \in [0, 1)$.

3 Semi-Structural Empirical Methods

So far we have presented a condition that is necessary and sufficient to recover structural shocks from a set of observables, using complete knowledge of the structural model. That is, using knowledge of the coefficient matrices A, B, and C in the state-space system (4), or more generally, the spectral characteristic $\varphi(\lambda)$ associated with the linear transformation from shocks to observables. Given this knowledge, it is possible to use equation (7) to obtain the spectral characteristic $\psi(\lambda)$ associated with the linear transformation from observables to shocks. We refer to the process of recovering shocks in this way, using all the restrictions embedded in the structural model, as the "fully structural" approach.

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

In more detail, the semi-structural approach is made up of two steps. The first, which we call the "reduced-form" step, involves using time series methods to obtain an empirically adequate characterization of the autocovariance function (equivalently, the spectral density) of the observable processes. The goal of this step is essentially just to summarize the data. The second "structural" step involves imposing some (sub-) set of restrictions derived from economic theory, which are sufficient to re-

cover the structural shocks of interest. The goal of this step is to entertain and test hypotheses with economic content.

It should be clear that recoverability is a necessary condition for using semi-structural methods to recover economic shocks. If the shocks cannot be recovered even with the full set of structural restrictions, then there can be no hope of doing so with only a subset of those conditions. However, it should be equally clear that invertibility is not a necessary condition, either for the reduced-form model or the structural model. Both models *could* be invertible, but they could also both be non-invertible, or one could be invertible but not the other. Moreover, nothing about this approach ties it to using vector autoregressions, although that is typically the common practice.

To illustrate how semi-structural methods can be used in situations when invertibility fails but recoverability does not, we return to the simple permanent income model of the previous section. We have already seen in this example that the recoverability condition in Theorem (1) is satisfied. Now we will show how, given an estimate of the autocovariance function (equivalently, the spectral density) of surplus income changes, it is possible to recover the structural income shocks by imposing a subset of the model's theoretical restrictions. Then, we will show how to impose these restrictions in the special case that the reduced-form model is an autoregression.

Example 2 (continued). Suppose that the econometrician has an estimate of the autocovariance function of surplus income changes, summarized by its spectral density

$$\hat{f}_{\Delta s}(\lambda) = \frac{1}{2\pi} \sum_{j=-\infty}^{\infty} E[\Delta s_t \Delta s_{t-j}] e^{-i\lambda j}$$

In other words, he has completed the reduced-form step of the analysis and is ready to perform the structural step.

First, let us consider what he would do if the model were invertible, since that case should be familiar. From equation (14), the structural model says that surplus income changes are a first-order moving average process with respect to the income shocks. If it were somehow true that $R < 1$, the Fourier coefficients of $\alpha(z)$ in equation (16) would vanish for negative values of j, so this model would be invertible according to Theorem (2). In this case, he would impose the following theoretical restrictions

$$\hat{\varepsilon}_t \in \mathcal{H}_t(\Delta s) \ominus \mathcal{H}_{t-1}(\Delta s) \quad \text{for all } t \in \mathbb{Z}, \tag{19}$$

where $A \ominus B = C$ means that $A = B \oplus C$, and $\oplus$ denotes the direct sum.

This way of writing the structural restrictions may seem unusual; often the restrictions are described simply as "orthogonality conditions," and written in the form

$$E[\hat{\varepsilon}_t \Delta s_{t-j}] = 0 \quad \text{for all } j < 0. \tag{20}$$

However, there is a subtle but important difference. Equation (19) implies that the orthogonality conditions in (20) hold. But it *also* implies that $\varepsilon_t \in \mathcal{H}_t(\Delta s)$; that is, that the structural shocks are *invertible*. Without this additional restriction, the orthogonality conditions in equation (20) would not be sufficient to identify the true income shocks, even up to a scale factor. The general point is that invertibility is itself a theoretical restriction that the econometrician would need to impose as part of the structural step of his semi-structural analysis.

To find a estimated process satisfying equation (19), the econometrican would obtain the Wold representation of $\{\Delta s_t\}$ by factoring the spectral density function as

$$\hat{f}_{\Delta s}(\lambda) = \frac{1}{2\pi} |\hat{\delta}(\lambda)|^2 \tag{21}$$

where the Fourier coefficients of $\hat{\delta}(\lambda)$ vanish for negative values of j.[14] This means that $\{\Delta s_t\}$ can be related to the uncorrelated Wold shock process $\{\hat{w}_t\}$ through a one-sided moving average,

$$\Delta s_t = \sum_{j=0}^{\infty} \hat{\delta}_j \hat{w}_{t-j}. \tag{22}$$

By construction, these shocks $\{\hat{w}_t\}$ satisfy equation (19), and are unique up to a scale factor. Therefore, the econometrician sets $\hat{\varepsilon}_t = \hat{w}_t$, and under the null hypothesis that the theoretical restrictions in equation (19) are valid, he would recover the true income shocks up to a scale factor.

Unfortunately, as we saw in the previous section, the theoretical restrictions in equation (19) are not valid in the economically relevant case that $R > 1$ because the income shocks are not invertible from surplus income changes. Nevertheless, since the shocks are *recoverable*, the econometrician can proceed in a similar fashion. Let us use $\mathcal{H}^t(\Delta s)$ to denote the Hilbert space spanned by Δs_τ for $t \leq \tau < \infty$. The

[14] Procedures for doing this in practice are well known. For example, when the reduced-form model delivers an estimate $\hat{f}_{\Delta s}(\lambda)$ which is a rational function of $e^{-i\lambda}$, we can use the procedure outlined in Chapter 1 of Rozanov (1967).

econometrician can require that

$$\hat{\varepsilon}_t \in \mathcal{H}^{t+1}(\Delta s) \ominus \mathcal{H}^{t+2}(\Delta s) \quad \text{for all } t \in \mathbb{Z}. \tag{23}$$

These restrictions imply that the orthogonality conditions

$$E[\hat{\varepsilon}_t \Delta s_{t-j}] = 0 \quad \text{for all } j < -1$$

hold, as well as that $\hat{\varepsilon}_t \in \mathcal{H}^{t+1}(\Delta s)$.

To find an estimated shock process satisfying equation (23), the econometrician needs to solve a spectral factorization problem analogous to the one in equation (21). Specifically, he needs to compute the spectral factor $\hat{\varphi}(\lambda)$ such that

$$\hat{f}_{\Delta s}(\lambda) = \frac{1}{2\pi}|\hat{\varphi}(\lambda)|^2, \tag{24}$$

where now the Fourier coefficients of $\hat{\varphi}(\lambda)$ vanish for all $j < -1$. The solution to this problem can be obtained immediately from the Wold factorization in equation (21):

$$\hat{\varphi}(\lambda) = \hat{\delta}(-\lambda)e^{-i\lambda}.$$

(The additional multiplication by $e^{i\lambda}$ corresponds to a one-period time shift, since the model's timing convention says that the restrictions in equation (23) hold for $j < -1$ not $j < 0$.) Under the null hypothesis that the theoretical model is correctly specified, the econometrician will recover the structural shocks up to a scale factor.

We have shown that the structural step of the analysis involves solving a spectral factorization problem, where the constraints on that problem come from economic theory. Now we can step backward to the reduced-form step and ask what sort of spectral density estimate the econometrician might use. One possibility is that he use a standard autoregression as the reduced-form model. Under this choice, he obtains a reduced-form representation of the form

$$\sum_{j=0}^{\infty} \gamma_j \Delta s_{t-j} = u_t,$$

where $\{u_t\}$ is an uncorrelated "reduced-form" shock process with zero mean and unit variance, and the coefficients $\{\gamma_j\}$ are square-summable. Based on this representation, his spectral density estimate is given by

$$\hat{f}_{\Delta s}(\lambda) = \frac{1}{|\gamma(\lambda)|^2},$$

where $\gamma(\lambda)$ is the Fourier transform of the sequence $\{\gamma_j\}$. Using this reduced-form model, his solution for the structural factor in equation (24) is

$$\hat{\varphi}(\lambda) = \frac{e^{-i\lambda}}{\gamma(\lambda)}.$$

$\Diamond$

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 econometrician. Their date-t information set is given by the subspace $\mathcal{H}_t(\varepsilon)$, while the information set of the econometrician is given by $\mathcal{H}_t(\Delta s)$. When $R > 1$, we have shown that $\mathcal{H}_t(\Delta s) \subset \mathcal{H}_t(\varepsilon)$. If the econometrician were placed on the same informational footing as agents, 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 an econometrician were on equal footing with economic agents, he would still be unable to recover structural shocks from current and past observables. Models with noise shocks are one example, and we discuss these at length in the following section.

4 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 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.[15] 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 provided reason to doubt.

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 semi-structural methods — in particular structural VAR analysis — can be applied to recover 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 3. *Model*: At each date, consumption is equal to agents' long-run forecast of total factor productivity,

$$c_t = \lim_{j \to \infty} E_t[a_{t+j}]. \tag{25}$$

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

$$a_t = a_{t-1} + \sigma_a \varepsilon_t^a, \tag{26}$$

and the signal about future productivity is given by

$$s_t = \left(\frac{1-\rho}{1+\rho}\right) \sum_{j=-\infty}^{\infty} \rho^{|j|} a_{t-j} + v_t. \tag{27}$$

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

[15] 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).

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

$$v_t = 2\rho v_{t-1} - \rho^2 v_{t-2} + \sigma_v \varepsilon_t^v - (\beta + \bar{\beta})\sigma_v \varepsilon_{t-1}^v + \beta\bar{\beta}\sigma_v \varepsilon_{t-2}^v. \tag{28}$$

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.[16]

Because $\{a_t\}$ is not covariance stationary, we need to clarify the precise meaning of the forecast in equation (25). 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

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

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 (26) and (27), $\{\tilde{q}_t(\theta)\}$ can be obtained from $\{\varepsilon_t\}$ by a linear transformation with spectral characteristic

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

[16] 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).

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

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

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), it follows 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 an estimate of the spectral density of $\{\Delta y_t\}$, $\hat{f}_{\Delta y}(\lambda)$. The structural step involves factoring the spectral density as

$$\hat{f}_{\Delta y}(\lambda) = \frac{1}{2\pi}\hat{\varphi}(\lambda)\hat{\varphi}(\lambda)^*, \tag{29}$$

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

$$\hat{\varepsilon}^a_t \in \mathcal{H}_t(\Delta a) \ominus \mathcal{H}_{t-1}(\Delta a) \tag{30}$$

$$\hat{\varepsilon}^v_t \in \mathcal{H}_t(\Delta \hat{v}) \ominus \mathcal{H}_{t-1}(\Delta \hat{v}) \tag{31}$$

for all $t \in \mathbb{Z}$, where $\Delta\hat{v}_t$ is the orthogonal projection of Δc_t onto $\mathcal{H}(\Delta y) \ominus \mathcal{H}(\Delta a)$. Equation (30) says that the fundamental shock is the Wold innovation in productivity growth, and equation (31) says that the noise shock captures the fluctuations in current consumption growth that are orthogonal to productivity growth at all horizons.

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

$$\hat{\varphi}(\lambda) = \begin{bmatrix} \hat{\varphi}_{11}(\lambda) & 0 \\ \hat{\varphi}_{21}(\lambda) & \hat{\varphi}_{22}(\lambda) \end{bmatrix}. \tag{32}$$

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

$$\begin{bmatrix} \Delta a_t \\ \Delta c_t \end{bmatrix} = \cdots + \underbrace{\begin{bmatrix} 0 & 0 \\ * & 0 \end{bmatrix}}_{b_{-1}} \begin{bmatrix} \hat{\varepsilon}^a_{t+1} \\ \hat{\varepsilon}^v_{t+1} \end{bmatrix} + \underbrace{\begin{bmatrix} * & 0 \\ * & * \end{bmatrix}}_{b_0} \begin{bmatrix} \hat{\varepsilon}^a_{t} \\ \hat{\varepsilon}^v_{t} \end{bmatrix} + \underbrace{\begin{bmatrix} * & 0 \\ * & * \end{bmatrix}}_{b_1} \begin{bmatrix} \hat{\varepsilon}^a_{t-1} \\ \hat{\varepsilon}^v_{t-1} \end{bmatrix} + \cdots,$$

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

To obtain the factor $\hat{\varphi}(\lambda)$, we can write equation (29) out more explicitly, using equation (32), as

$$\begin{bmatrix} \hat{f}_{\Delta a}(\lambda) & \hat{f}_{\Delta a \Delta c}(\lambda) \\ \hat{f}_{\Delta c \Delta a}(\lambda) & \hat{f}_{\Delta c}(\lambda) \end{bmatrix} = \frac{1}{2\pi} \begin{bmatrix} |\hat{\varphi}_{11}(\lambda)|^2 & \hat{\varphi}_{11}(\lambda)\overline{\hat{\varphi}_{21}(\lambda)} \\ \hat{\varphi}_{11}(\lambda)\hat{\varphi}_{21}(\lambda) & |\hat{\varphi}_{22}(\lambda)|^2 + |\hat{\varphi}_{21}(\lambda)|^2 \end{bmatrix}. \tag{33}$$

The restrictions in equation (30) say that $\hat{\varphi}_{11}(\lambda)$ is nothing other than the canonical (Wold) factor of $\hat{f}_{\Delta a}(\lambda)$. This is unique and can be obtained in the usual way. The lower-left equation in (33) 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 (33) implies that

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

Together with the restrictions in equation (31), this means that $\hat{\varphi}_{22}(\lambda)$ is uniquely determined as the canonical factor of $2\pi \hat{f}_{\Delta c}(\lambda) - |\hat{\varphi}_{21}(\lambda)|^2$. 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 reduced-form model the econometrician uses to construct his esimtae of the spectral density $\hat{f}_{\Delta y}(\lambda)$. Of course, there are many possibilities. One popular proposal is to approximate the spectral density using a p-th order VAR model of the form

$$\sum_{j=0}^{p} \gamma_j y_{t-j} = u_t, \tag{34}$$

where $\{u_t\}$ is a two-dimensional uncorrelated "reduced-form" shock process with zero mean and identity covariance matrix. When this type of reduced-form model is used, the entire semi-structural strategy is usually referred to as "structural VAR analysis."

If we define

$$\gamma(\lambda) \equiv \sum_{j=0}^{p} \gamma_j e^{-i\lambda j},$$

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

$$\hat{f}_{\Delta y}(\lambda) = \frac{1}{2\pi}(1 - e^{-i\lambda})\gamma(\lambda)\gamma(\lambda)^*(1 - e^{i\lambda}).$$

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

4.1 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 relies only on the structural restrictions in equations (30) and (31).

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 (34). 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% 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.

Perhaps one puzzling aspect of this result is that it is apparently possible to identify the effects of a shock that has non-flat effects on consumption. Blanchard et al. (2013) explain that an econometrician with access to the same information consumers or less, cannot identify any shock with non-flat effects on consumption. This is because consumption is a random walk in this model, conditional on consumers' information. We agree with this result.

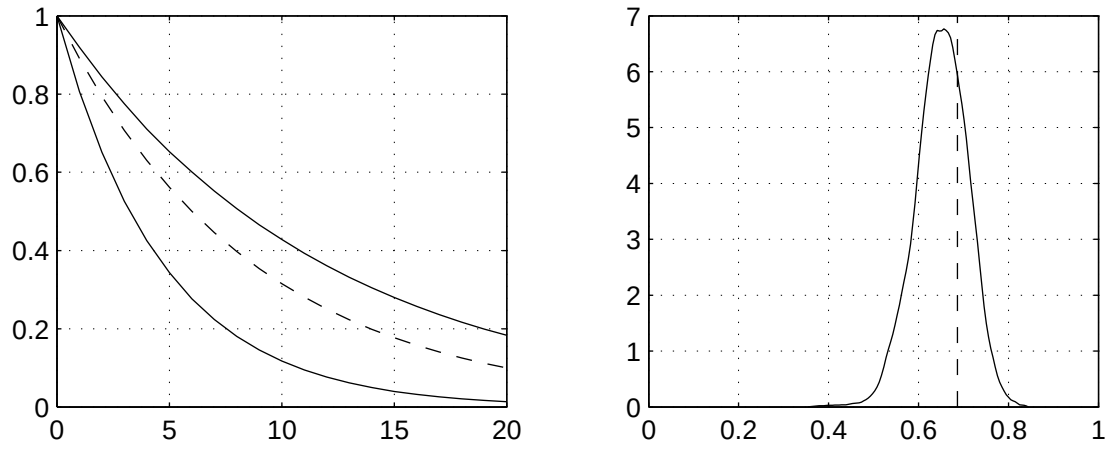

Figure 1: Strutural VAR analysis of data simulated from a model with noise shocks. *Left*: the dashed line is the true impulse response of consumption to a unit noise shock, while solid lines are 95% bands from the distribution of point estimates from each of $N = 1000$ samples of length $T = 275$. *Right*: 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.

However, the key observation is that as econometricians we always have *more* information than the agents in our models. To perform any sort of analysis on a sample of data, that data must have been realized at some date in the past. As a result, relative to agents at each date in our sample, we have access to more information about both consumption and productivity. It is by using this additional information that the we can successfully identify the effects of noise shocks.

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.

4.2 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 structural VARs, he argues

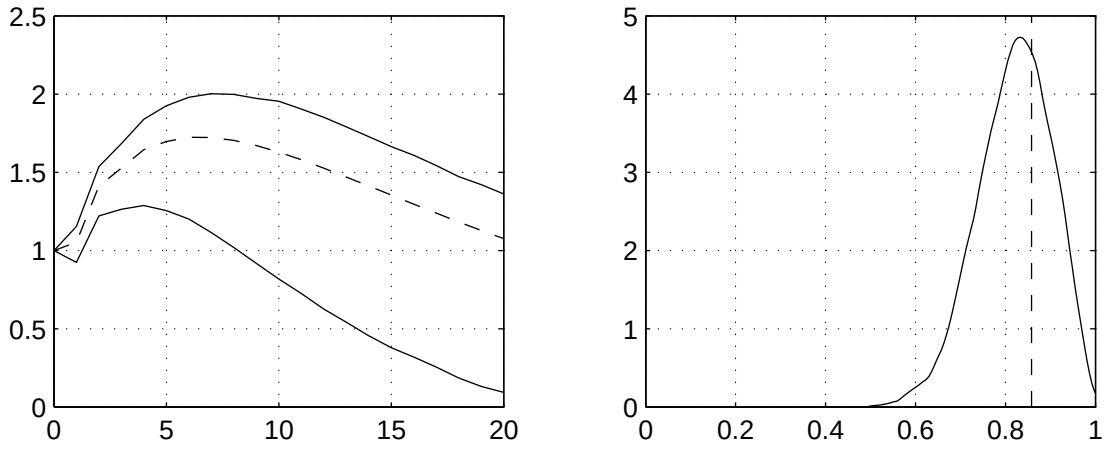

Figure 2: Structural VAR analysis of quarterly U.S. consumption and total factor productivity from 1948:Q1 to 2016:Q4. *Left*: response of consumption to unit noise shock. The dashed line is the point estimate, and the solid lines are 95% bootstrap confidence bands. *Right*: share of consumption variance 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.

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.

5 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) required 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, which have to the best of our knowledge been absent from the literature so far. 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 function. 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 save such an investigation for future research.

References

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

Beaudry, Paul and Franck Portier (2014) "News-Driven Business Cycles: Insights and

Challenges," *Journal of Economic Literature*, 52(4):993–1074.

Blanchard, Olivier J., Jean-Paul L'Huillier, and Guido Lorenzoni (2013) "News, Noise, and Fluctuations: An Empirical Exploration," *American Economic Review*, 103(7):3045–3070.

Chahrour, Ryan and Kyle Jurado (2017) "News or Noise? The Missing Link," Working Paper, Duke University.

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

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

Dupor, Bill and Jing Han (2011) "Handling Non-Invertibility: Theory and Applications," Working Paper, Ohio State University.

Fernández-Villaverde, Jesús, Juan F. Rubio-Ramírez, Thomas J. Sargent, and Mark W. Watson (2007) "ABCs (and Ds) of Understanding VARs," *American Economic Review*, 97(3):1021–1026.

Forni, Mario, Luca Gambetti, Marco Lippi, , and Luca Sala (2017) "Noisy News in Business Cycles," *American Economic Journal: Macroeconomics, Forthcoming*, 1(1):1.

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

Hannan, Edward J. and Barry 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, Lars Peter and Thomas J. Sargent (1980) "Formulating and estimating dynamic linear rational expectations models," *Journal of Economic Dynamics and Control*, 2(1):7–46.

Hansen, Lars Peter and Thomas 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, Eric M., Todd B. Walker, and Shu-Chun Susan Yang (2013) "Fiscal Foresight and Information Flows," *Econometrica*, 81(3):1115–1145.

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

Lippi, Marco and Lucrezia Reichlin (1994) "VAR analysis, nonfundamental representations, blaschke matrices," *Journal of Econometrics*, 63(1):307–325.

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

Morf, Martin, Augusto Vieira, Daniel T. L. Lee, and Thomas Kailath (1978) "Recursive Multichannel Maximum Entropy Spectral Estimation," *IEEE Transactions on Geoscience Electronics*, 16(2):85–94.

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

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

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

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

Sims, Christopher A. (1980) "Macroeconomics and Reality," *Econometrica*, 48(1):1–48.

Sims, Christopher A. and Tao Zha (2006) "Does Monetary Policy Generate Recessions?," *Macroeconomic Dynamics*, 10(2):231272.

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