

Recoverability

Ryan Chahrour

Kyle Jurado

Boston College

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

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 the first step of these methods often entails obtaining an invertible

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

³This 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-Grohe and Uribe (2012); Barsky and Sims (2012); and Blanchard et al. (2013).

⁴In fact, this is the original Sargent xent see description 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.

⁵For 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.

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. $f_t g$ is "recoverable" from $f_t g$ if

$$H(\cdot) \subset H(\cdot):$$

This says that each of the variables $x_{k,t}$ is contained in the space $H(\cdot)$. That is, each of these variables is perfectly revealed by the information contained in $f_t g$. In the Gaussian case, this can be expressed in terms of mathematical expectations as

$$x_{k,t} = E[x_{k,t} | H(\cdot)]:$$

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. $f_t g$ is "invertible" from $f_t g$ if

$$H_t(\cdot) \subset H_t(\cdot) \text{ for all } t \in \mathbb{Z}:$$

Since $H_t(\cdot) \subset H(\cdot)$, 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 $f_t g$ as

$$f_t = \int_{\mathbb{Z}} e^{i\lambda t} dF(\lambda);$$

where F is its associated random spectral measure. We say that a $1 \times n$ dimensional vector function $f(\lambda)$, defined for $\lambda \in \mathbb{Z}$, belongs to the space $L^2(F)$ if

$$\int_{\mathbb{Z}} f(\lambda) F(d\lambda) f(\lambda)^* < \infty \quad \text{or} \quad \int_{\mathbb{Z}} \sum_{k,l=1}^n f_k(\lambda) \overline{f_l(\lambda)} F_{;kl}(d\lambda) < \infty:$$

In this expression, F denotes the spectral measure of $f_t g$ and the asterisk denotes complex conjugate transposition.⁶ If we define the scalar product

$$(f_1; f_2) = \int_{\mathbb{Z}} f_1(\lambda) F(d\lambda) f_2(\lambda)^*:$$

⁶That is, $F_{;kl}(\lambda) = E[x_{k,t} \overline{x_{l,t}}]$

and do not distinguish between two vector functions that satisfy $k_1 - k_2 = 0$, then $L^2(F)$ becomes a Hilbert space. Using these definitions, the following lemma gives an alternative characterization of recoverability.

Lemma 1. $f_t g$ is recoverable from $f_t g$ if and only if there exists an $n \times n$ matrix function $'(t)$ with rows in $L^2(F)$ such that

$$\int_{\mathbb{Z}} e^{j t'}(t) (d) \text{ for all } t \in \mathbb{Z}: \quad (1)$$

Proof. Define the operator T such that $Th = n$

The process $\tilde{f}y_t g$ 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^*} .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

$$\begin{aligned} \text{(observation)} \quad y_t &= Ax_t & (4) \\ \text{(state)} \quad x_t &= Bx_{t-1} + C''_t \end{aligned}$$

where x_t is an n_x -dimensional state vector. In this case, the spectral characteristic $\gamma(\lambda)$ takes the form

$$\gamma(\lambda) = A(I_{n_x} - Be^{-i\lambda})^{-1}C \quad (5)$$

The solution to a wide class of linear (or linearized) dynamic equilibrium models can be written in this form.⁹

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 $f''_t g$ are recoverable from the observables $\tilde{f}y_t g$ if and only if*

$$\text{rank}(\gamma(\lambda)) = n^*$$

for almost all $\lambda \in [0; \pi]$.

Proof. Sufficiency: $\tilde{f}y_t g$ can be obtained from $f''_t g$ by a linear transformation with spectral characteristic $\gamma(\lambda)$. This means that the random spectral measure of $\tilde{f}y_t g$ can be decomposed as¹¹

$$\tilde{y}(d) = \gamma(\lambda) \tilde{f}(d) \quad (6)$$

Because $\gamma(\lambda)$ has constant rank n^* , there exists an $n^* \times n_y$ matrix function such that

$$\gamma(\lambda) \gamma(\lambda)' = I_{n^*} \quad (7)$$

⁸Linear regularity means that

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

$$y(d) = \alpha(d):$$

Moreover, note that the rows of α are elements of $L^2(F_y)$ because for any $k = 1, \dots, n_y$, equations (6) and (7) imply that

$$\int_{\mathcal{Z}} \alpha_k(d) F_y(d) \alpha_k(d) = \frac{1}{2} \int_{\mathcal{Z}} \alpha_k(d)' \alpha_k(d) d = 1 < 1:$$

Therefore $f''_t g$ can be obtained from $f_{y_t} g$ by a linear transformation with spectral characteristic α . By Lemma (1), it follows that the shocks are recoverable.

Necessity: To the contrary, suppose that the shocks are recoverable, so $H(\alpha) = H(y)$, but that α' has rank different than n_y on some set of positive measure. Because α' has n_y columns, its rank can never be greater than n_y . Therefore, its rank on this set must be strictly less than this.

Now we find an element in $H(\alpha)$ that is not in $H(y)$, which is a contradiction. Because $\text{rank}(\alpha') < n_y$ on some set of positive measure, there exists a $1 \times n_y$ vector function $\beta \in L^2(F_y)$ such that $\beta' \alpha \neq 0$ and

$$\beta' \alpha(d) = 0$$

for all $d \in \mathcal{Z}$. This would mean that the element

$$\beta' \alpha(d)$$

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

$$\int_{\mathcal{Z}} \beta_k(d) e^{i t d} \alpha_k(d) d = 0:$$

But this contradicts the hypothesis that $H(y) = H(\alpha)$. □

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

$$R(\alpha) = \begin{pmatrix} I_{n_x} & B e^{i \cdot} & C \\ A & 0_{n_y \times n_x} \end{pmatrix} \quad \# \quad (8)$$

be full column rank for almost all $\omega \in \Omega$. This follows from the so-called Guttman rank additivity formula. Specifying the condition in terms of $R(\omega)$ 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 ω from the uniform distribution over Ω and check whether $R(\omega)$ 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, $n_y \geq n_s$. This is intuitive; it isn't possible to recover n_s separate sources of random variation without observations of at least n_s 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 $\text{rank}(R(\omega)) = n_s$ 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 $f_{y_t|g}$ by a linear transformation of the form

$$y_t = \int_{-\infty}^{\infty} e^{i\lambda t} R(\omega) w(d\omega); \quad (9)$$

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

Using the spectral characteristic $\gamma(\lambda)$ from equation (9) and the function $\gamma(\lambda)$

defined in equation (6.017(ha.1 016828(y))TJ/F20 4a3on)8 -1.793 Td [(t)]TJ/F22 11.9552400039 0 s2Hqu

Proof. The fact that $w_s, s \leq t$ forms a basis in $H_t(y)$ at each date means that a variable h is an element of $H_t(y)$ if and only if it can be represented in the form of a series

$$h = \sum_{j=0}^{\infty} w_{t-j} \quad (10)$$

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

By equations (7) and (9),

$$h_t = \sum_{d=0}^{\infty} e^{j-t} ()_y(d) = \sum_{d=0}^{\infty} e^{j-t} () () w(d) : \quad (11)$$

The rows of $()$ are elements of $L^2(F_y)$, but they may not be square integrable

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

$$s_t = \frac{1}{R} \Delta s_t = \frac{1}{R} (s_t - s_{t-1}); \quad (14)$$

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

The spectral characteristic linking the shocks to observables is

$$f(\omega) = \frac{1}{R} e^{-i\omega}$$

for some integer $p > 0$, where $f_t g$ is a stationary process. In this case we can define a new process

$$\tilde{f}_t(\omega) = \sum_{j=0}^{\infty} e^{i\omega t - j} \frac{1}{(1 - e^{-i\omega})^p} f_{t-j} g(\omega); \quad (18)$$

which is stationary for each value of $\omega \in [0; 2\pi)$. We can say that a process $f_t g$ is recoverable (or invertible) with respect to $\tilde{f}_t g$ whenever $f_t g$ is recoverable (or invertible) with respect to $\tilde{f}_t(\omega) g$ for almost all $\omega \in [0; 2\pi)$.

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 $\gamma(\omega)$ associated with the linear transformation from shocks to observables. Given this knowledge, it is possible to use equation (7) to obtain the spectral characteristic $\tilde{\gamma}(\omega)$ 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-

where $A \oplus B = C$ 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 the re-

econometrician can require that

$$\hat{u}_t \in H^{t+1}(s) \cap H^{t+2}(s) \text{ for all } t \in \mathbb{Z}; \quad (23)$$

These restrictions imply that the orthogonality conditions

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

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

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 $\hat{u}(\omega)$ such that

$$\hat{f}_s(\omega) = \frac{1}{2} \hat{u}(\omega) \hat{u}(\omega)^* \quad (24)$$

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

$$\hat{u}(\omega) = \hat{u}(\omega) e^{-i\omega} :$$

(The additional multiplication by $e^{-i\omega}$ 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

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

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

$$\hat{f}_s(\omega) = \frac{1}{\sum_{j=0}^{\infty} \alpha_j e^{-i\omega j} \sum_{k=0}^{\infty} \alpha_k e^{i\omega k}};$$

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

$$\hat{g}(\omega) = \frac{e^{-i\omega}}{\hat{g}(\omega)}:$$

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 $H_t(\omega)$, while the information set of the econometrician is given by $H_t(s)$. When $R > 1$, we have shown that $H_t(s) \subset H_t(\omega)$. If the econometrician were placed on the same informational footing

information beyond a_t itself. The process $f_{V_t g}$ represents non-fundamental noise, and is assumed to follow a law of motion of the form

$$v_t = 2 v_{t-1} - v_{t-2} + v_t^a \quad (+) \quad v_{t-1}^v + v_{t-2}^v \quad (28)$$

The vector of fundamental and noise shocks, $u_t = (u_t^a, u_t^v)'$, is independent and identically distributed over time with zero mean and identity covariance matrix. There is also a nonlinear restriction on the parameters a , v_1 , and v_2 , which ensures that $f_{a_t g}$ can be written alternatively as the sum of a permanent component with rst-

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

$$\frac{1}{2} \int_{-\pi}^{\pi} e^{ij} \frac{(1 - e^{-ij})^2}{j^2} d = \frac{1}{1 + |j|} \quad (28)$$

It is easy to see that $\hat{f}_t(\omega)$ has full rank for almost all $\omega \in [-\pi, \pi]$ and $\omega \in [0, \pi]$ whenever $a_i v_i > 0$. By Theorem (1), this means that the structural shocks are recoverable with respect to $\hat{f}_t(\omega)g$ for almost all ω . Using the terminology introduced in Remark (3), it follows that the shocks are recoverable from $\hat{f}_t g$.

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 $\hat{f}_t g$, $\hat{F}_t(\omega)$. The structural step involves factoring the spectral density as

$$\hat{F}_t(\omega) = \frac{1}{2} \hat{\Lambda}_t(\omega) \hat{\Lambda}_t(\omega)'; \quad (29)$$

where the factor $\hat{\Lambda}_t(\omega)$ is defined by a set of theoretical restrictions that are sufficient to correctly identify the structural shocks in the model. One such set is

$$\hat{\Lambda}_t^a \geq H_t(a) \quad H_{t-1}(a) \quad (30)$$

$$\hat{\Lambda}_t^v \geq H_t(v) \quad H_{t-1}(v) \quad (31)$$

for all $t \in \mathbb{Z}$, where \hat{v}_t is the orthogonal projection of c_t onto $H(c) \perp H(a)$. 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 $\hat{\Lambda}_t(\omega)$ has a lower-triangular form

$$\hat{\Lambda}_t(\omega) = \begin{pmatrix} \hat{\Lambda}_{11}(\omega) & 0 \\ \hat{\Lambda}_{21}(\omega) & \hat{\Lambda}_{22}(\omega) \end{pmatrix}; \quad (32)$$

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

$$\begin{pmatrix} a_t \\ c_t \end{pmatrix} = \begin{pmatrix} 0 & 0 \\ \frac{0}{b_1} \{Z\} \end{pmatrix} \hat{\Lambda}_{t+1}^a + \begin{pmatrix} 0 \\ \frac{0}{b_0} \{Z\} \end{pmatrix} \hat{\Lambda}_t^a + \begin{pmatrix} 0 \\ \frac{0}{b_1} \{Z\} \end{pmatrix} \hat{\Lambda}_t^v + \begin{pmatrix} 0 \\ \frac{0}{b_1} \{Z\} \end{pmatrix} \hat{\Lambda}_{t-1}^a + \begin{pmatrix} 0 \\ \frac{0}{b_1} \{Z\} \end{pmatrix} \hat{\Lambda}_{t-1}^v; \quad (33)$$

where $\hat{f}_t g$ are the sequence of Fourier coefficients associated with $\hat{\Lambda}_t(\omega)$.

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

$$\begin{matrix} \hat{f}_a(\lambda) & \hat{f}_{ac}(\lambda) \\ \hat{f}_{ca}(\lambda) & \hat{f}_c(\lambda) \end{matrix} = \frac{1}{2} \begin{matrix} j^{\hat{h}_{11}(\lambda)^2} & \hat{h}_{11}(\lambda) \overline{\hat{h}_{21}(\lambda)} \\ \hat{h}_{11}(\lambda) \hat{h}_{21}(\lambda) & j^{\hat{h}_{22}(\lambda)^2} + j^{\hat{h}_{21}(\lambda)^2} \end{matrix} \quad (33)$$

The restrictions in equation (30) say that $\hat{h}_{11}(\lambda)$ is nothing other than the canonical (Wold) factor of $\hat{f}_a(\lambda)$. This is unique and can be obtained in the usual way. The lower-left equation in (33) uniquely determines $\hat{h}_{21}(\lambda)$ as a function of $\hat{f}_{ca}(\lambda)$ and $\hat{h}_{11}(\lambda)$, 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

$$j^{\hat{h}_{22}(\lambda)^2} = 2 \hat{f}_c(\lambda) - j^{\hat{h}_{21}(\lambda)^2}$$

Together with the restrictions in equation (31), this means that $\hat{h}_{22}(\lambda)$ is uniquely determined as the canonical factor of $2 \hat{f}_c(\lambda) - j^{\hat{h}_{21}(\lambda)^2}$. Therefore, we have shown

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

$$\alpha = 0.8910; \quad \beta = 0.6700; \quad \gamma = 0.9937; \quad \text{and} \quad \delta = 0.7833 \quad 0.1525;$$

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 resWe .95 one(v)27w81(d)r78(ou5)tialle pone to




Figure 1: Structural 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

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.

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.

Fernandez-Villaverde, Jesus, Juan F. Rubio-Ramirez, Thomas J. Sargent, and

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*