Identifying News Shocks from Forecasts

Jonathan J. Adams* Philip Barrett†

March 29, 2023

EARLY DRAFT

Link to Most Current Version

Abstract

We propose a method to identify the anticipated components of macroeconomic shocks in a structural VAR: we include empirical forecasts about each time series in the VAR, which introduces enough linear restrictions to identify each structural shock and to further decompose each one into "news" and "surprise" shocks. We estimate our VAR on US time series using forecast data from the SPF, Federal Reserve, and asset prices. The fiscal stimulus and interest rate shocks that we identify have typical effects that comport with existing evidence. In our news-surprise decomposition, we find that news contributes to a third of US business cycle volatility, where the effect of fiscal shocks is mostly anticipated, while the effect of monetary policy shocks is mostly unexpected. Finally, we use the news structure of the shocks to estimate counterfactual policy rules, and compare the ability of fiscal and monetary policy to moderate business cycles.

JEL-Codes: C32, E32, E52, E62

Keywords: Identification, Structural shocks, SVAR, News, Fiscal multiplier, Fiscal policy, Monetary policy, Policy counterfactuals

*University of Florida, email: adamsjonathan@ufl.edu

†Research Department, International Monetary Fund, email: pbarrett@imf.org.

The views expressed herein are those of the authors and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

1 Introduction

The effects of anticipated macroeconomic shocks differ from when the shocks are unexpected. Is it possible to isolate the effects of news from surprises in general settings? Estimating these different effects is crucial for drawing conclusions, especially regarding the effectiveness of policy. In this paper, we introduce a method to separately identify the anticipated and unanticipated components of macroeconomic shocks.

Our strategy is to include data on forecasts about the macroeconomic time series in a vector autoregression (VAR). Forecasts are valuable because they reveal information about the future that is not otherwise revealed by the macroeconomic time series alone. We modify a standard structural VAR (SVAR) driven by a series of structural shocks, by assuming that each shock has an anticipated component – the "news" – and an unanticipated component – the "surprise". If each news dimension is sufficiently relevant, we prove that including a forecast about *each* time series in the VAR identifies the effect of the news and surprise components of every structural shock.

Our method is not only useful for isolating news from surprise: it is a method to identify structural shocks themselves. Structural VARs typically assume that shocks are mutually orthogonal in order to identify them from reduced form innovations in the observed time series. If their news and surprise components are also mutually orthogonal, then our method identifies the *entire set* of structural shocks, including their news and surprise components. Thus our method is an alternative to the large variety of other strategies for identifying the set of structural shocks in VARs.[1]

We apply our method by estimating a VAR on US time series. We take data on forecasts from the Survey of Professional Forecasters (SPF), the Federal Reserve's Greenbook forecasts, and also construct some expectations from asset prices. In our VAR, we estimate a variety of structural shocks that resemble well-understood objects, including shocks to fiscal and monetary policy. Our estimated shocks have realistic effects, including fiscal multipliers that match other estimates in the literature, quantitatively realistic effects of monetary policy shocks that resemble those implied by high-frequency-identified instruments. Crucially, we can decompose each shock into the news and surprise components. For example, we find that the effects of fiscal shocks on output are largely anticipated, and the news component

[1] A classic approach is to make assumptions about the causal ordering of shocks within a period, and apply a Cholesky decomposition to the variance matrix (Sims, 1980). Other linear restrictions can identify the structural shocks by making assumptions about long-run effects (Shapiro and Watson, 1988), restrictions on the signs of shocks (Uhlig, 2005) or outside evidence on the magnitude of short-run effects (Blanchard and Perotti, 2002). Recently, attention has been focused on identifying the set of structural shocks using higher order moments and heteroskedasticity. Examples with dynamic heteroskedasticity include Sentana and Fiorentini (2001), Rigobon (2003), Lanne et al. (2010), and Lewis (2021). Lütkepohl and Netšunajev (2017) reviews this literature further. Other papers lean on non-Gaussianity more generally including Hyvärinen et al. (2010) and Gouriéroux et al. (2017).

implies much a larger government spending multiplier than the surprise component, echoing the findings in Ramey (2011). In contrast, the effects of monetary policy shocks are mostly surprises.

We find a large role for news in explaining business cycles: half of output volatility is due to news shocks. This echoes the findings of a large literature studying the relevance of news shocks for the macroeconomy. Many of these papers focus on news about technology[2] but we join a modest group studying news about policy shocks, discussed below. Indeed, many papers follow a conceptually similar approach to ours by including a forecast in their VAR to isolate surprises or news about the forecasted variable.[3] However, including a single forecast identifies a specific news shock only if there is a single structural shock that is anticipated. Otherwise, what might appear to be news about a shock such as fiscal policy also includes news about shocks to supply, demand, and so forth. This is the main advantage of our approach relative to existing VAR studies of news: by including forecasts about every time series, we can distinguish the effects of news to different structural shocks. And we find that conflation of news about multiple shocks is a nontrivial concern, as the news component of nearly all shocks is relevant for at least one time series.

A valuable advantage of decomposing shocks into news and surprise is the ability to estimate the effects of counterfactual policies. Wolf and McKay (2022) demonstrate that, under some assumptions, impulse response functions to news about shocks at different horizons are sufficient to construct counterfactual impulse response functions under alternative policy rules. We implement their approach using our identification of impulse responses to news and surprise shocks and conduct several counterfactual experiments. We find that fiscal policy can be effective at stabilizing output over the business cycle, but with costs: taxes and inflation become more volatile. And current fiscal policy is already somewhat stabilizing; when we consider a counterfactual with fixed government spending, real activity and inflation are both more volatile. There are some shocks that fiscal policy is not effective at moderating, but monetary policy is more effective at moderating precisely these shocks, suggesting a role for fiscal and monetary coordination. We come to similar conclusions as Wolf and McKay when considering counterfactual monetary policy. The best counterfactual monetary policy rules that we can construct are less effective at stabilizing output than fis-

[2]Examples include Beaudry and Portier (2006), Barsky and Sims (2012), Schmitt-Grohé and Uribe (2012), Blanchard et al. (2013), and Chahrour and Jurado (2021). The most closely related papers are those that utilize forecast data to identify news about technology: Hirose and Kurozumi (2021) include forecast data in a New Keynesian DSGE model to identify news shocks and estimate that technology news drives nearly half of output volatility; Cascaldi-Garcia (2022) uses forecast revisions of economic growth to instrument for technology news shocks, which drive $11\% - 26\%$ of output volatility depending on the horizon.

[3]Papers including forecasts to identify fiscal surprises include Ramey (2011), Auerbach and Gorodnichenko (2012), and Born et al. (2013). VAR methods using forecasts and additional structural assumptions to identify fiscal news include Caggiano et al. (2015), Ricco (2015), Ricco et al. (2016) and Forni and Gambetti (2016).

cal stimulus, while interest rate pegs do not lead to more volatile inflation and cause output to be more elastic to shocks in the short run.

We contribute to a large literature studying the effects of news about fiscal policy. Ramey (2011) uses narrative methods to identify changes in current and future government spending driven by military events, and argues the many fiscal shocks identified by structural VARs are actually anticipated. Fisher and Peters (2010) use financial returns to defense contractors to identify shocks that include news about future defense spending. Ben Zeev and Pappa (2017) apply the Barsky and Sims (2012) methodology to identify the shock dimension that contains the most news about government defense spending over a 5-year horizon. A common theme in these papers is that the fiscal multiplier due to news about government spending is large.

The revenue side of fiscal policy has received a similar treatment. Leeper et al. (2009) argue VAR-based estimates of shocks will be misleading when tax changes are anticipated. Romer and Romer (2010) use a narrative approach to construct a series of anticipated tax changes, and estimate that legislation of relatively exogenous tax increases have large contractionary effects. Mertens and Ravn (2012) decompose the Romer-Romer series into anticipated and unanticipated components, and show that they have opposite effects on output in the short run. House and Shapiro (2006) come to a similar conclusion studying tax reforms in the early 2000s. Ramey (2019) surveys additional evidence.

2 A Simple Example

We introduce our identification strategy in a simple example, before exploring the general case.

2.1 Identification in the Simple Example

Consider the univariate time series x_t is given by

$$x_t = \rho x_{t-1} + \epsilon_t + \gamma v_t \tag{1}$$

where ϵ_t is some i.i.d. fundamental shock process. However, while ϵ_t affects the variable x_t at time t, some component of it may be anticipated in advance. Assume that the innovation $\epsilon_t \sim N(0, \sigma_\epsilon)$ is the sum of two components:

$$\epsilon_t = u_t + v_{t-1}$$

where $u_t \sim N(0, \sigma_u)$ and $v_{t-1} \sim N(0, \sigma_v)$ are independent, so that $\sigma_u^2 + \sigma_v^2 = \sigma_\epsilon^2$.

With this structure, news has two effects. News revealed in period $t-1$ directly affects the realization of the shock ϵ_t. But we also allow news about the future v_t to enter equation (1) contemporaneously if $\gamma \neq 0$. This is because when individuals receive news about future shocks, the information may affect their decisions immediately. For example, if firms expect their productivity to change in the future, they will adjust their investment today. This presents additional challenges: news affects some variables with delay, but others immediately. The full multivariate model is necessary to capture these effects on other variables, but the econometric challenges are clear even when considering the simple univariate example.

An econometrician observing only the history of x_t is unable to distinguish between the "surprise shock" u_t and "news shock" v_{t-1}, because each period introduces two new shock realizations, but only one new observation of x_t.

Crucially, while the "news shock" v_{t-1} is not directly observed by the econometrician, it *is* observed by forecasters at time $t-1$. For this simple example, assume that forecasters report their rational expectation $\mathbb{E}[x_{t+1}|x_t, v_t]$:

$$f_t = \mathbb{E}[x_{t+1}|x_t, v_t] = \rho x_t + v_t \tag{2}$$

The contemporaneous forecasts include information that the econometrician cannot infer from the history of x_t alone.

Including the forecasts f_t in a VAR allows news and surprise to be jointly identified. The two-dimensional VAR is given by

$$\begin{pmatrix} f_t \\ x_t \end{pmatrix} = \mathbf{B} \begin{pmatrix} f_{t-1} \\ x_{t-1} \end{pmatrix} + \mathbf{A} \begin{pmatrix} v_t \\ u_t \end{pmatrix} \tag{3}$$

and let $w_t = \begin{pmatrix} f_t \\ x_t \end{pmatrix} - \mathbf{B} \begin{pmatrix} f_{t-1} \\ x_{t-1} \end{pmatrix}$ denote the reduced form innovations. In this case, the coefficient matrices are

$$\mathbf{B} = \begin{pmatrix} \rho & 0 \\ 1 & 0 \end{pmatrix} \qquad \mathbf{A} = \begin{pmatrix} 1+\rho\gamma & \rho \\ \gamma & 1 \end{pmatrix}$$

$\mathbf{A}$ is necessarily invertible, so the econometrician can identify both structural shocks by $\begin{pmatrix} v_t \\ u_t \end{pmatrix} = \mathbf{A}^{-1} w_t$

In this simple example, including forecasts in a VAR allows for the anticipated and surprise components of a shock to be independently identified. However, with such a simple model, their effects are so similar as to be uninteresting. News and surprise imply the same impulse response function of x_t. But when news can also have contemporaneous effects, these two shocks may imply very different dynamics.

2.2 The Importance of Accounting for News

In order to identify the dynamic effects of the fundamental ϵ_t shocks, the econometrician *must* account for news using the f_t forecasts.

What does an econometrician without forecast data estimate? This time series is an ARMA(1,1) process in which neither v_t nor u_t are independently identifiable. To see why, consider the Wold representation:

$$x_t = \rho x_{t-1} + w_t + \theta w_{t-1}$$

with the white noise innovation denoted w_t. The MA component is found by solving two equations for θ and $var(w_t)$:

$$(1+\theta^2)var(w_t) = var(w_t + \theta w_{t-1}) = (1+\gamma^2)\sigma_v^2 + \sigma_u^2$$

$$\theta var(w_t) = cov(w_t + \theta w_{t-1}, w_{t+1} + \theta w_t) = cov(\gamma v_t + v_{t-1} + u_t, \gamma v_{t+1} + v_t + u_{t+1}) = \gamma^2\sigma_v^2$$

The actual expression for w_t in terms of the underlying shocks is recursive, and thus most cleanly expressed using lag operator polynomials:

$$w_t = \frac{1}{(1+\theta L)}\left(u_t + (\gamma + L)v_t\right)$$

An econometrician estimating the ARMA(1,1) process for x_t will not recover the impulse response function (IRF) to a fundamental shock ϵ_t. Figure 1b reports these IRFs. The solid blue line is the IRF to a unit innovation in the ARMA, estimated without forecast data. The dashed red line is the IRF to an average unit increase in the fundamental shock. A unit ϵ_t may have many possible combinations of news and surprise; the plot shows the average combination. A component of this shock is anticipated, given that $\epsilon_t = u_t + v_{t-1}$, so there is a non-causal entry in the plot. The w IRF and ϵ IRF have dissimilar shapes.

Without forecast data, there is a more substantial problem than being unable to identify the underlying shocks: each component u_t, v_{t-1}, and their sum ϵ_t has a different impulse response function, none of which can be identified from the x_t time series. The one-standard-deviation IRF for u_t is:

$$\mathbb{E}[x_{t+h}|u_t = \sigma_u] = \rho^h\sigma_u$$

the shape of this IRF is easily identified by estimating the AR coefficient of the ARMA(1,1),

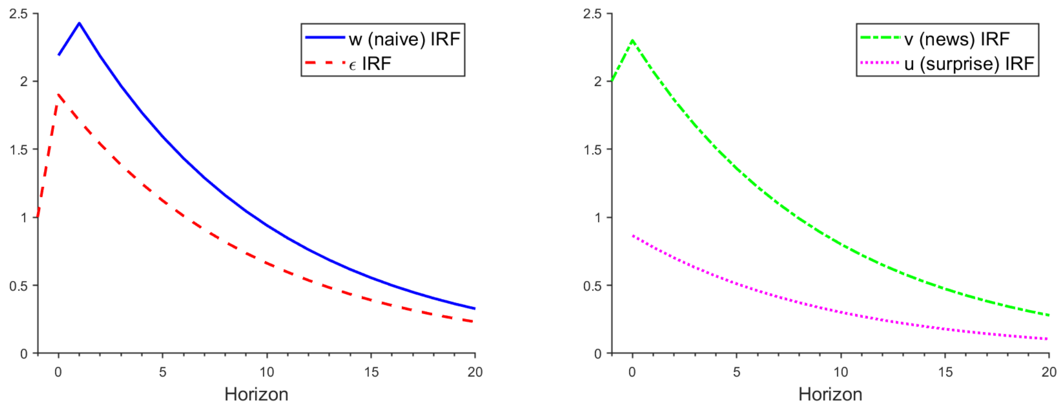

(a) Fundamental vs. Reduced Form Shocks (b) News and Surprise Components

Figure 1: Impulse Response Functions in the Simple Example

Panel 1a plots the impulse response to a unit innovation to the ARMA(1,1) Wold representation (blue solid line), and to a unit structural shock (red dashed line). Panel 1b plots the impulse response to a standard deviation innovation to the structural news component (dash-dot green line) and to the structural surprise component (dotted magenta line). The parameter values are $\rho = 0.9$, $\gamma = 4$, $\sigma_u^2 = 0.75$, and $\sigma_v^2 = 0.25$.

but the scale σ_u is not identified. Worse, the IRF for v_t is not a simple exponential decay:

$$\mathbb{E}[x_{t+h}|v_t = \sigma_v] = \begin{cases} \gamma\sigma_v & h = 0 \\ (\gamma\rho^h + \rho^{h-1})\sigma_v & h > 0 \end{cases}$$

so neither the scale nor the shape is identified, because γ and σ_v are not identified without forecast data.

Figure 1b plots the true impulse responses to the news and surprise components. The dotted magenta line is the response to a one-standard-deviation surprise shock u_t: x_t suddenly jumps then declines as if driven by an AR(1) process. The green line is the response to a one-standard-deviation news shock v_{t-1}. The anticipation effect of the news causes an immediate increase. Then, once the full shock is realized in period t, there is a further increase, followed by the same exponential decay as a surprise shock. Properly weighted, these two IRFs sum to the ϵ IRF in Figure 1a.

Including forecasts in a VAR allowed for news and surprises to be separately identified in this univariate example. But identification is always more complicated in a multivariate VAR, which typically requires additional identifying assumptions to identify structural shocks. Yet, the next section demonstrates that the lessons from the simple example generalize: including rational forecasts is enough for identification without any additional structure.

3 Identification

This section outlines the general structural VAR, provides a constructive proof of identification, describes how rational forecasts are cleaned from empirical forecasts, and derives the implied impulse response functions.

3.1 The Basic VAR Model

A standard structural VAR(m) is

$$x_t = \sum_{j=1}^{m} B_j x_{t-j} + A\epsilon_t$$

where x_t is an $n \times 1$ vector of time series, B_j are a series of $n \times n$ coefficient matrices, $\epsilon_t \sim N(0, I)$ is an $n \times 1$ vector of standard normal structural shocks, and A is an $n \times n$ matrix that determines how the structural shocks affect contemporaneous time series.

We allow the structural shocks to be partially anticipated in ways that are not directly observable to the econometrician. Specifically, we decompose the shock ϵ_t into a surprise component u_t and a news component v_{t-1} that is anticipated one period in advance:

$$\epsilon_t = u_t + v_{t-1}$$

We assume the components are orthogonal so that news does not predict surprises: $u_t \perp v_{t-1}$. The standard SVAR assumption is that each entry in the shock vector is mutually orthogonal, so that $Var(\epsilon_t) = I$. We further assume that the entries in the surprise and news components are mutually orthogonal, i.e. $Var(u_t) = D_u^2$ and $Var(v_{t-1}) = D_v^2$ where D_u and D_v are diagonal matrices.[4] The requirement that news is orthogonal to future surprises implies

$$D_u^2 + D_v^2 = I \tag{4}$$

The standard SVAR is also modified so that time series may be affected by news about future shocks v_t:

$$x_t = \sum_{j=1}^{m} B_j x_{t-j} + A\epsilon_t + Cv_t \tag{5}$$

This is not an identifying restriction. Rather, allowing for arbitrary C introduces additional flexibility and challenges. But economic theory suggests that news about future shocks can have large effects on contemporaneous decisions.

[4]Alternatively, this property is implied by assuming that the structural shocks are not just uncorrelated, but independent.

Assume that f_t is a vector of rational expectations for the corresponding time series:

$$f_t = E\left[x_{t+1}|\{x_{t-j}\}_{j=0}^{m-1}, \epsilon_t, v_t\right] \tag{6}$$

The expectation is conditional on current news v_t, so the vector f_t contains information that may not be directly observable to the econometrician.

Because f_t is the rational expectation, there exist restrictions on the relationship between f_t and x_t that are sufficient to identify all of the structural shocks. Equation (5) implies that f_t follows

$$f_t = \sum_{j=1}^{m} B_j x_{t+1-j} + A v_t \tag{7}$$

because $E\left[\epsilon_{t+1}|\{x_{t-j}\}_{j=0}^{m-1}, \epsilon_t, v_t\right] = v_t$ and $E\left[v_{t+1}|\{x_{t-j}\}_{j=0}^{m-1}, \epsilon_t, v_t\right] = 0$.

The time series x_t can be written recursively in terms of current surprises u_t and current news v_t using the SVAR structure (5) and the rational expectation (7):

$$x_t = \sum_{j=1}^{m} B_j x_{t-j} + A(u_t + v_{t-1}) + C v_t$$

$$= \sum_{j=1}^{m} B_j x_{t-j} + (f_{t-1} - \sum_{j=1}^{m} B_j x_{t-j}) + A u_t + C v_t$$

$$= f_{t-1} + A u_t + C v_t$$

The expectations f_t can similarly be written

$$f_t = B_1 x_t + \sum_{j=2}^{m} B_j x_{t+1-j} + A v_t$$

$$= B_1(f_{t-1} + A u_t + C v_t) + \sum_{j=2}^{m} B_j x_{t+1-j} + A v_t$$

Stack the expectations and time series into a single VAR($m-1$):

$$\begin{pmatrix} f_t \\ x_t \end{pmatrix} = \sum_{j=1}^{m-1} \mathbf{B}_j \begin{pmatrix} f_{t-j} \\ x_{t-j} \end{pmatrix} + \mathbf{A} \begin{pmatrix} v_t \\ u_t \end{pmatrix} \tag{8}$$

where

$$\mathbf{B}_j \equiv \begin{cases} \begin{pmatrix} B_1 & B_2 \\ I & 0 \end{pmatrix} & j = 1 \\ \begin{pmatrix} 0 & B_{j+1} \\ 0 & 0 \end{pmatrix} & j > 1 \end{cases}$$

and

$$\mathbf{A} \equiv \begin{pmatrix} B_1 C + A & B_1 A \\ C & A \end{pmatrix}$$

Estimating the VAR (8) recovers the coefficients $\{B_j\}_{j=1}^m$ and the variance matrix of forecast errors Σ, which satisfies

$$\Sigma = \mathbf{A} \begin{pmatrix} D_v^2 & 0 \\ 0 & D_u^2 \end{pmatrix} \mathbf{A}'$$

The symmetric matrix Σ has $2n^2 + n$ unique entries. B_1 is identified from the VAR, while A and C each have n^2 unknown parameters. D_u^2 and D_v^2 each have n unknowns, but the equation (4) implies n additional restrictions, enough to exactly identify the unknown parameters.

3.2 Deriving the Estimator

In this section, we introduce and prove the main identification theorem. The proof is constructive, describing how to estimate the unknown matrices given estimates from the reduced form VAR of the first coefficient matrix B_1 and the residual covariance matrix Σ.

The model must satisfy two key assumptions. First, A must be invertible: this implies that the shocks in ϵ_t have linearly independent effects on the time series. Second, D_v^2 must be invertible: each shock must have a nontrivial news component. However, we do not require that D_u^2 is invertible, i.e. some shocks can be fully anticipated.

Theorem 1 *If A and D_v^2 are full rank, then A, C, D_u^2 and D_v^2 are determined (up to sign and column order) by Σ and B_1.*

Proof. Subdivide the matrix $\Sigma = \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix}$ into $n \times n$ blocks. The off-diagonal submatrices satisfy $\Sigma_{12} = \Sigma_{21}'$, so the three remaining submatrices are given by

$$\begin{pmatrix} \Sigma_{11} & \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix} = \begin{pmatrix} (B_1 C + A) D_v^2 (B_1 C + A)' + B_1 A D_u^2 A' B_1' & \\ C D_v^2 (B_1 C + A)' + A D_u^2 A' B_1' & C D_v^2 C' + A D_u^2 A' \end{pmatrix}$$

Define the $n \times n$ matrices ϕ and ψ by

$$\phi \equiv \Sigma_{11} - B_1\Sigma_{21} - \Sigma_{21}'B_1' + B_1\Sigma_{22}B_1'$$

$$= AD_v^2A'$$

$$\psi \equiv \Sigma_{22} - (\Sigma_{21} - \Sigma_{22}B_1')\phi^{-1}(\Sigma_{21} - \Sigma_{22}B_1')'$$

$$= CD_v^2C' + AD_u^2A' - CD_v^2A'(AD_v^2A')^{-1}AD_v^2C'$$

$$= AD_u^2A'$$

Equation (4) implies

$$\phi + \psi = AA'$$

The singular value decomposition (SVD) of $\phi + \psi$ gives unitary matrix U and diagonal matrix Λ^2 such that

$$\phi + \psi = U\Lambda^2U'$$

and

$$A = U\Lambda V'$$

for some unitary V. Then the SVD of $\Lambda^{-1}U'\phi U\Lambda^{-1}$ gives the matrices V and D_v^2 from

$$\Lambda^{-1}U'\phi U\Lambda^{-1} = V'D_v^2V$$

This gives the matrices $A = U\Lambda V'$ and $D_u^2 = I - D_v^2$. Then the final matrix C is found from

$$C = (\Sigma_{21} - \Sigma_{22}B_1')(D_v^2A')^{-1}$$

■

The application of the singular value decomposition makes it clear that the shocks are only identified up to column order; the SVD is only unique up to reordering of the singular values. Choosing an order for the singular values implies an ordering of the shocks in ϵ_t. Moreover, our method only determines the variances of the shocks D_u^2 and D_v^2, so the shock signs are also indeterminate.

3.3 Forecast Cleaning

In practice, empirical forecasts $\tilde{f}_t$ may not correspond to the rational expectation (7). For example, there is extensive evidence that surveyed expectations feature predictable biases.[5]

[5]Notable examples include Souleles (2004), Greenwood and Shleifer (2014) Coibion and Gorodnichenko (2015), and Bordalo et al. (2020), among many others.

Therefore it is necessary to "clean" any empirical forecasts in order to transform them into rational expectations. To be a rational expectation, the cleaned forecast's errors must be orthogonal to m lags of the time series x_t, of the empirical forecasts $\tilde{f}_t$, and any other data z_t in the information set.

To construct the rational expectation f_t, we run the VAR(k) with $k \geq m$:

$$\begin{pmatrix} \tilde{f}_t \\ z_t \\ x_t \end{pmatrix} = \sum_{j=1}^{k} \mathbf{G}_j \begin{pmatrix} \tilde{f}_{t-j} \\ z_{t-j} \\ x_{t-j} \end{pmatrix} + \upsilon_t$$

where υ_t is a reduced form error.

Let $\mathbf{G}_{x,j}$ denote the final n rows of $\mathbf{G}_j$. The cleaned rational forecast f_t is given by

$$f_t = \sum_{j=1}^{k} \mathbf{G}_{x,j} \begin{pmatrix} \tilde{f}_{t+1-j} \\ z_{t+1-j} \\ x_{t+1-j} \end{pmatrix} \tag{9}$$

which is the best linear forecast of x_{t+1} conditional on the information set spanned by lags of measured forecasts $\tilde{f}_t$, the time series x_t, and other regressors z_t.

Under some assumptions, this cleaning procedure recovers the true rational expectation. We model empirical forecasts $\tilde{f}_t$ as linear deviations from the rational forecast f_t. The deviations may depend on lags of the rational forecast f_t, the time series x_t, observable confounders z_t, fundamental surprises u_t, or fundamental news v_t:

$$\tilde{f}_t = \sum_{j=0}^{k} \left(H_j^f f_{t-j} + H_j^x x_{t-j} + H_j^z z_{t-j} + H_j^u u_{t-j} + H_j^v v_{t-j} \right)$$

or in terms of lag operator polynomials

$$\tilde{f}_t = H^f(L) f_t + H^x(L) x_t + H^z(L) z_t + H^u(L) u_t + H^v(L) v_t \tag{10}$$

Theorem 2 *If $H^f(L)$ is causally invertible, then the rational forecast f_t is given by equation* (9).

Proof: Appendix A.1

This approach makes two strong assumptions: the additional confounding terms z_t are all observable, and $H^f(L)$ is invertible. In particular, if aggregate forecasts reflect publicly available information, the observability assumption is a reasonable one. But – as with any regression – it will be essential to include all of the relevant controls in the forecast cleaning.

What if the assumptions are broken, so that forecasts are affected by some unobserved confounders beyond z_t? In these cases we can still clean the forecast and identify shocks under looser assumptions. But the interpretation of a news shock changes. Appendix A.2 considers this case.

3.4 Impulse Response Functions in the Presence of News

This section describes the impulse response functions implied by the structural VAR.

The horizon h impulse response $\phi_u(h)$ to a surprise u_t is standard:

$$\phi_u(h) = B^h A$$

$\phi_u(h)$ is a matrix, so that the entry in row i and column j captures the horizon h response of time series i to shock j.

The impulse responses to news have an additional term, because the news shock v_{t-1} first affects the period $t-1$ time series through the news channel, and then again in period t when the full shock is realized. The corresponding impulse response matrix is:

$$\phi_v(h) = \begin{cases} C & h = 0 \\ B^h C + B^{h-1} A & h > 0 \end{cases}$$

The impulse response functions are related to conditional expectations by:

$$\mathbb{E}[x_{t+h}|u_t] = \phi_u(h)u_t \qquad \mathbb{E}[x_{t+h}|v_t] = \phi_v(h)v_t$$

The fundamental shock $\epsilon_t = u_t + v_{t-1}$ is the sum of the surprise and news components. We calculate the IRF to a unit ϵ_t shock as the response to an *average* ϵ_t realization:

$$\phi_\epsilon(h) = \mathbb{E}[x_{t+h}|\epsilon_t = 1]$$

$$= \mathbb{E}\left[\mathbb{E}[x_{t+h}|u_t] + \mathbb{E}[x_{t+h}|v_{t-1}]|\epsilon_t = 1\right]$$

$$= \mathbb{E}\left[\phi_u(h)u_t + \phi_v(h+1)v_{t-1}|\epsilon_t = 1\right]$$

$$= \phi_u(h)\mathbb{E}\left[u_t|\epsilon_t = 1\right] + \phi_v(h+1)\mathbb{E}\left[v_{t-1}|\epsilon_t = 1\right]$$

$$= \phi_u(h)D_u^2 + \phi_v(h+1)D_v^2$$

where D_u^2 and D_v^2 are the diagonal matrices of shock variances.

Accordingly, for each shock i, a unit impulse to ϵ_t^i is the sum of a $Var(u_t^i)$ impulse to u_t^i and a $Var(v_{t-1}^i)$ impulse to v_{t-1}^i. Because of the news timing, the impulse response to ϵ_t is non-causal: it can affect time series in period $t-1$. Correctly accounting for the timing, the impulse response matrix is:

$$\phi_\epsilon(h) = \phi_u(h)D_u^2 + \phi_v(h+1)D_v^2$$

$$= \begin{cases} CD_v^2 & h = -1 \\ B^{h+1}CD_v^2 + B^h A(D_u^2 + D_v^2) & h \geq 0 \end{cases}$$

3.5 Generalizations and Alternatives

Our main approach applies to a broad class of dynamic models. But it still includes some restrictions that can be further relaxed.

Thus far, we have assumed that news occurs one period in advance. But news might realistically have longer horizons. For example, Mertens and Ravn (2012) estimate the effects of tax changes with announcements measured up to 16 quarters in advance of the policy change.

It is possible to account for additional news horizons by incorporating data on additional forecasts. Appendix B derives the SVAR restrictions in this case. Some additional horizons may be feasible, but the data requirements grow rapidly: we show that the VAR is potentially identified when news occurs at h different horizons by including forecasts at each of the h additional horizons. This may be possible for some variables – in particular interest rates and inflation – but many variables do not have widely available forecast data beyond a year in advance. For example, the Survey of Professional Forecasts only reports expectations over $0 - 4$ quarter horizons.

We have also assumed thus far that the econometrician has data on all relevant state variables in the economy. That is, they observe the entire vector x_t and the associated forecasts in the SVAR equation (5). But what if a critical time series is missing from the data? Appendix D derives the appropriate SVAR restrictions when some state variables are unobserved. News and noise shocks may still be identified, but the problem is computationally more intensive; we do not have an analytical solution for the implied decomposition of the variance matrix Σ.

4 Application to Fiscal and Monetary Policy Shocks

We apply our structural VAR method to data on US time series. We identify clear fiscal shocks and monetary policy shocks, estimate the implied multipliers, and study the general effects of news versus surprises.

4.1 Data

Our main source of forecast data is the Survey of Professional Forecasters (SPF), which is currently run by the Federal Reserve Bank of Philadelphia. The survey is administered quarterly to roughly 40 anonymous forecasters since 1968. We take the median reported values as our measure of forecasts.

Some variables are not available in the SPF for the entire sample, so we turn to other sources. In particular, the SPF only collects estimates on real government consumption and investment since 1981:III, so before this period we draw from the Federal Reserve's official forecasts reported in the Greenbook for every FOMC meeting. These values are not collected in publicly available datasets for all periods, so when necessary, we transcribe them from the original Greenbooks. For each quarter, we take the most recent estimate. We also use the Greenbook forecasts for Federal budget receipts and surpluses. For these variables, we use the dataset collected by Croushore and van Norden (2018), which we extend to 2016:IV by transcribing from the most recently released Greenbooks.

For interest rates, we measure forecasts directly from the yield curve. We use this measure because the SPF only provides forecasts for a limited number of interest rate horizons, and only since 1981:III. Where r_t^h denotes the return from time t to $t+h$, we calculate the forecast $\mathbb{E}[r_{t+1}^h]$ by

$$\mathbb{E}[r_{t+1}^h] = r_t^{h+1} - r_t^1$$

This is known to be a biased forecast, as the yield curve incorporates liquidity and risk premia as well as expectations. Yet while the yield curve-implied forecasts do not exactly match the SPF forecasts, they track each other very closely; for 3-month T-bills, the correlation coefficient is 0.996.

Finally, we use 3-month-ahead futures contracts to measure forecasts for oil prices and exchange rates. Covered interest rate parity predicts that the implied forecasted growth rates should track 3-month interest rates closely, but not exactly; deviations depend on expected costs of holding oil or interest rate differences across countries, respectively.

Table 1 reports the time series that we use. We transform the variables in three different ways. For NIPA variables and federal budget variables, we follow Ramey (2016) and divide by an estimated quadratic trend in real GDP. This transformation allows fiscal multipliers to be read directly from the impulse response functions. For the price level as measured by the GDP deflator, we take log differences and annualize to calculate the inflation rate. For other variables that grow regularly (e.g. housing starts), we take logs, but we leave in levels those variables that are not clearly nonstationary (unemployment, interest and exchange rates). Finally, we remove a quadratic trend and linear seasonal factors from all variables.

Variable	Date range	Forecast Source
Baseline Specification		
Real GDP	1968:IV - 2022:II	SPF
Federal tax receipts	1968:IV - 2016:IV	Fed Greenbooks
Real government spending	1968:IV - 2022:II	Fed Greenbooks for 1968:IV - 1981:II SPF for 1981:III - 2022:II
GDP deflator	1968:IV - 2022:II	SPF
3-month Treasury rate	1968:IV - 2022:II	Yield curve
Housing starts	1968:IV - 2022:II	SPF
Additional Variables		
Unemployment Rate	1968:IV - 2022:II	SPF
Industrial production	1968:IV - 2022:II	SPF
Federal budget surpluses	1968:IV - 2016:IV	Fed Greenbooks
USD/CAD exchange rate	1968:IV - 2022:II	Futures contracts
Real oil price	1983:I - 2022:II	Futures contracts
1, 2, 3, 4, and 5-year Treasury rates	1968:IV - 2022:II	Yield curve

Table 1: List of Variables

Our baselines specification appears above the break in Table 1. We include output, government spending, taxes, short term interest rates, and inflation so that we might identify shocks that reflect fiscal and monetary policy, which have well-understood effects on these variables. We also include housing starts as a second measure of real activity; housing starts have SPF forecasts that cover our entire sample, and aggregate forward-looking decisions that may be informative about news. Figure 2 plots these detrended and deseasonalized time series.

Figure 2 plots our baseline time series and their associated forecasts. In constructing the forecast series f_t we aim to satisfy three objectives. The first objective is plausibility: that our forecasts plausibly reflect all information about outcomes x_{t+1} at time t. The second objective is that we do not overfit to the data. The third is the forecasts must satisfy the identifying assumption: that forecasts contain all the information already available to the VAR structure, formalized in equation (6).

To meet these objectives we proceed in in two steps, based on the methodology in Section 3.3. We start by constructing a vector of variables z_t which aims to include as much as possible of the information available at time t about relevant future outcomes. To do this without overfitting, we a construct three machine learning models separately for each of the six variables in the baseline VAR: an elastic net, a regression tree, and a simple linear projection. Each model predicts one-period-ahead outcomes using up to eight lags of

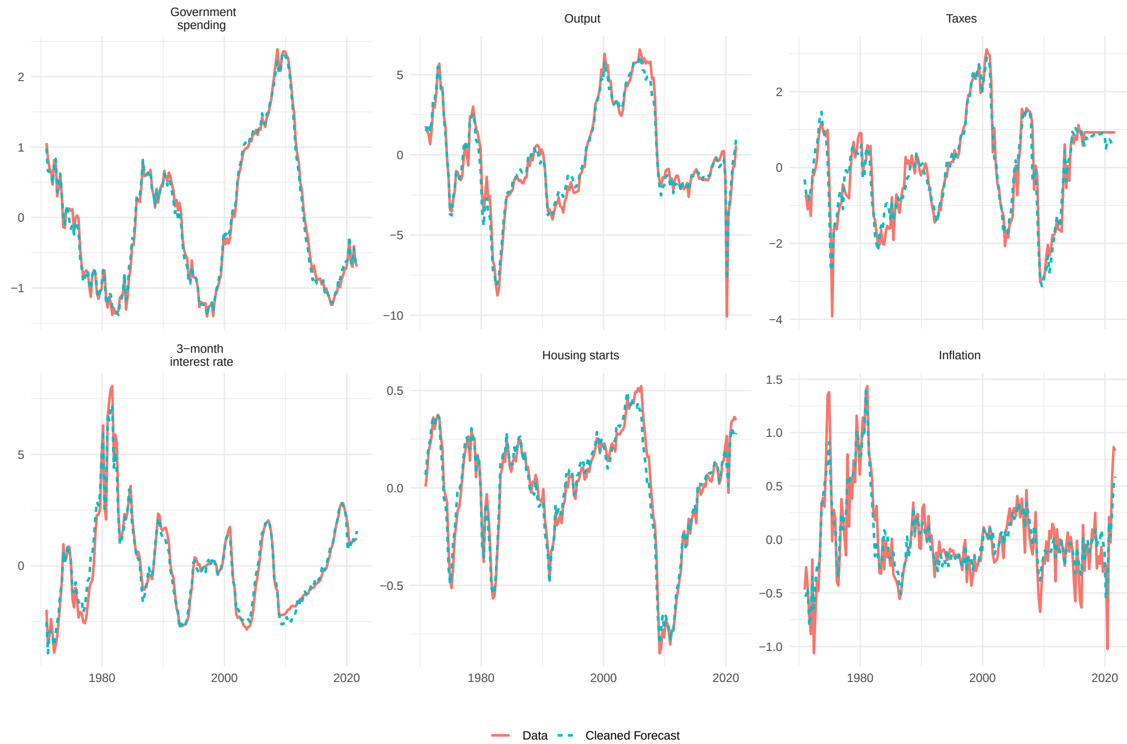

Figure 2: Baseline Time Series and Forecasts

The solid red line plots our baseline time series. Government spending, output, and federal taxes are real, deflated by the GDP deflator, and expressed relative to a quadratic real GDP trend. Housing starts are the natural log, and all data series are deseasonalized and detrended. The source of forecast data is the SPF for all baseline series, except the Federal Reserve's Greenbook is used for government spending before 1981:III, and for taxes, while the Treasury forecast is derived from the yield curve. Forecasts are cleaned to be rational in sample.

both data and outcomes for all 16 variables in Table 1, some 256 possible predictors. We use rolling cross-validation to select tuning parameters and then pick the model with the lowest out-of-sample average RMSE individually for each of the six variables. The fitted predictions thus embody plausible forecasts of x_t robust to overfitting. These, we label z_t. And so the N entries of z_t are the machine learning predictions for each of the elements of x_{t+1}. We then include these z_t in the cleaning process described in 3.3.

The advantage of this approach is that if there is a variable not in the VAR specification contains reliable information about future outcomes, this will be included in the constructed forecast f_t. For example, if lagged oil prices – a variable not in our baseline VAR – happen to be a robust predictor of inflation, then the machine learning models will include them. And so the relevant entry of z_t will contain the component of inflation that can be explained by oil prices. If this information is supplementary to the information in the lags of the data and the empirical forecasts, $(x_t, \ldots, x_{t-m}, \tilde{f}_t, \ldots, \tilde{f}_{t-m})$, then the cleaned forecast f_t will put weight on it. Likewise, if the empirical forecasts $\tilde{f}_t$ happen to embody all the information available about future outcomes, this method would allow f_t to fully reflect that.

One disadvantage of this method is that there remains some risk of overfitting. This arises because we clean the forecasts after cross-validating, and so there may be spurious reliance on the variables in the VAR. However, this is mitigated by the relatively short lag length and limited specification of the baseline VAR. Moreover, this reflects a deeper issue, that the well-known bias-variance tradeoff in forecasting means that our objective of not overfitting is not always compatible with the identifying assumption in equation (6). Yet our approach aims to limit the extent of this problem by using the machine learning forecasts as a bottleneck, limiting information about future outcomes to the same dimension as the data itself.

4.2 Measuring uncertainty over point estimates

To compute confidence intervals for various statistics, including impulse responses, we bootstrap the model. To do this, we create a large number, N^{sim}, of simulations of the data and unbiased forecasts using the data generating process, each of 202 periods – the same as our sample. For each simulation, we re-estimate the reduced-form VAR and denote the reduced form coefficients for the jth simulation by $(B_1^{(j)}, \ldots, B_m^{(j)}, \Sigma^{(j)})$. We then apply the identification process to produce estimated structural parameters $A^{(j)}, C^{(j)}, D_u^{(j)}, D_v^{(j)}$ using the algorithm outlined in Theorem 1. This givens us a simulated distribution of estimates which should reflect sampling uncertainty under the null hypothesis that the point estimates are consistent.

One difficulty is that the simulated structural matrices are only unique up to sign and re-ordering of the shocks. For example, if shock number 1 in the point estimate A happens to be a demand shock, there is no guarantee that the same shock is in column 1 of $A^{(j)}$ is the same shock. Depending on the ordering of components of the singular value decomposition, a completely different shock may be ordered first. Moreover, because the identification relies on a second-order statistic – the variance-covariance matrix – the identification is not unique up to sign. Multiplying the same column in the A and C matrices by -1 gives the same time series properties, just with the interpretation of what constitutes positive and negative shocks reversed.

Thus for each simulation, we search over all possible combinations of re-orderings and sign flips to find that which minimizes the square difference to the point estimates for A and C. With $N = 6$ variables, this is already very large, with 2^N possible sign flips and $N!$ possible reorderings, giving $2^N \times N! = 46,080$ possible combinations in total. This ordering procedure minimizes a continuous loss function, satisfying the requirements for

Lewis (2021) Theorem 4: our labeling method does not affect the asymptotic distribution of the structural matrices (and so neither the implied impulse response functions). We can thus use the sample of structural parameters so created to calculate the distributions of model statistics as required.

In general, the resulting confidence intervals are quite wide. However, it should be noted that these reflect a broad range of sources of uncertainty, not always included in other approaches. First, because we re-estimate the shock variance matrices $D_u^{(j)}, D_v^{(j)}$ for each simulation, our impulse responses show not just the uncertainty over how a given shock propagates, but also that due to the uncertainty over the size of each shock. This is particularly important when thinking about impulse responses. Second, sampling variation means that the reordering and re-signing of the shocks is imperfect – variation due to one shock may be mistakenly attributed to another. One feature of this is that confidence intervals often exhibit a degree of symmetry around zero, reflecting to difficulty in consistently signing the shocks. Third, the identification method itself is intrinsically nonlinear, with small changes in reduced form coefficients sometimes leading to large changes in the structural coefficients. Finally, because our bootstrap technique matches the observed sample length, we include variation appropriate to small sample, and do not rely on large-sample approximations.

As a way to check our estimated sample uncertainty from the bootstrap, we implement several alternative methods to compute the distribution of coefficient, each of which shut down some of these sources of uncertainty. One alternative is to draw reduced form samples from the asymptotic distribution of the VAR coefficients (see (Hamilton, 1994) for details). This isolates down the small-sample part of the sampling variability. Another is to use the delta method, which effectively linearizes the mapping from reduced form to structural coefficients, thus compressing the tails of the sampling distributions. Finally, we also investigate the role of uncertainty over structural variances D_u, D_v by presenting results where the size of the shock is held fixed in the estimation. We discuss these methods further further when presenting relevant statistics.

4.3 Shock Labeling

Our identification scheme recovers the structural shocks, but it does not tell us what they are. For convenience, and to help with interpretation, we devise a labeling scheme for the shocks, giving them each a name. We base this on the impulse responses to the average combination of news and surprise shocks. This is because these average shocks are the closest analogue of estimates which do not make a general distinction between the news and surprise components of macroeconomic shock. This allows us to label the shocks in line with the past work. We initially do this qualitatively, proposing a labelling scheme based on the signs of the impulse responses. For two of the shocks – monetary and fiscal policy shocks – we check that the quantitative responses they those estimated elsewhere.

4.3.1 Average impulse response functions

We base our shock labeling scheme on the average impulse responses. That is, the ϵ_t responses described in Section 3.4. These are shown in Figure 3, which we calculate as described in Section 3.4. The response to each structural shock is weighted average of the response to news and surprise components of the shock. For all variables, the response is measured as the percentage deviation from trend associated with a unit structural shock. Dotted and dashed lines show the central 60 and 90 percent of the bootstrapped distribution of outcomes respectively. We assign labels to each shock initially based on the signs and significance of the impulse responses.

The first shock features an immediate and statistically significant contraction in government tax revenues and a prolonged and statistically significant increase in government spending, albeit somewhat delayed. At the same time, output and real activity (as measured by housing starts) increases. This we label as a fiscal stimulus shock. In Section 4.3.2 we verify that the magnitude of these responses are consistent with tax and spending multipliers in the literature.

The second shock features a clear, statistically significant, and immediate increase in short term interest rates. This is followed by a decline in output over the next year or so and then a subsequent decline in inflation, although not always strongly statistically significant. In Section 4.3.3 we again verify this shock, by comparing to estimated monetary policy shocks in the literature.

The third and fourth shocks we label as demand and supply respectively. In the case of the former, the output response is immediate and statistically significant, with a more long-lasting increase in inflation and a delayed interest rate response. In contrast to the fiscal shock, spending goes down and taxes go up, consistent with a aggregate expansion not driven by the public sector. In the case of the latter, we base our labeling on the markedly opposing responses of inflation and output on impact.

We leave the final two shocks unlabeled. This is not to say that one could not make a case for a structural interpretation of either. In particular, the second unlabeled shock appears much like our monetary policy shock. However, these shocks both fail some of the quantitative validation tests below. And so we remain silent on the interpretation of these.

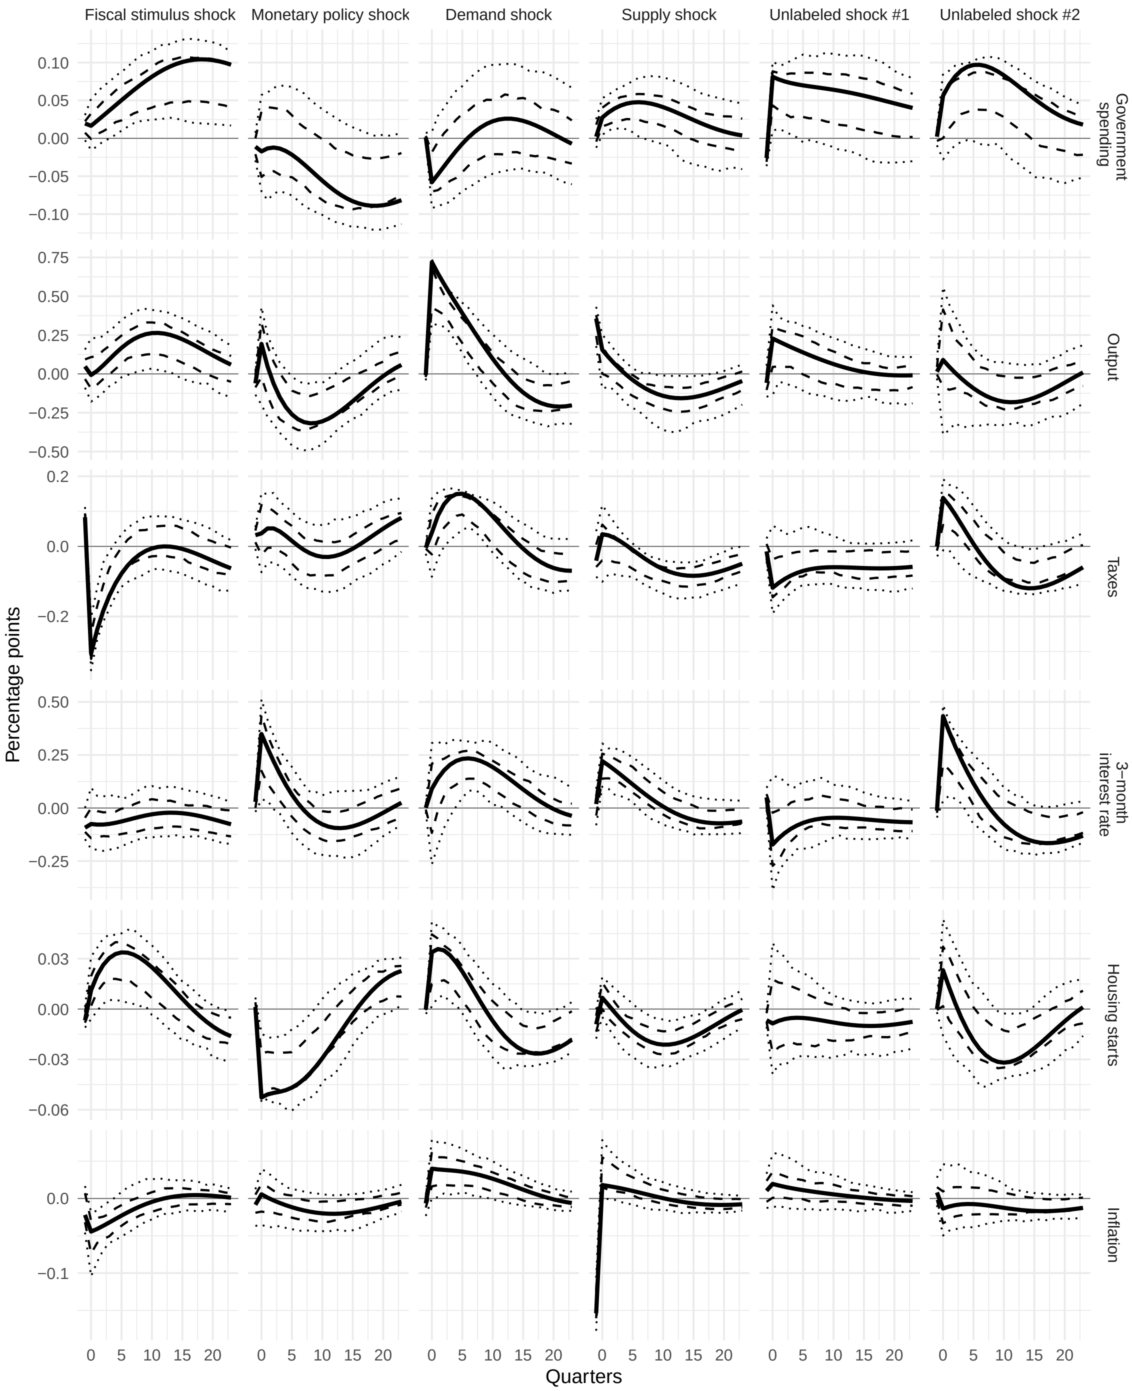

Figure 3: Estimated average impulse responses to structural shocks

The impulse response functions are plotted to an average one standard deviation structural shock, calculated as in Section 3.4. The solid line is the point estimate and the dashed and dotted lines show respectively the $5^{th} - 95^{th}$ and $20^{th} - 80^{th}$ percentile ranges from a bootstrap simulation with $N^{sim} = 500$ replications. For government consumption, output, and taxes, units are percentage points relative to trend lagged output. For inflation, interest rates, and housing starts, units are annualized percentage points relative to own-variable lagged trend.

4.3.2 Validating the Fiscal Policy Shock

To corroborate our interpretation of the first shock as a fiscal policy changes, we show that the responses are consistent with tax and spending multipliers estimated the literature.

Typically, the $h-$period fiscal multiplier in response to a fiscal policy shock is defined as the ratio of the cumulative change in output relative to the cumulative change the relevant fiscal variable (either taxes or spending).[6] That is, the multipliers are:

$$\mu_G^h = \frac{\sum_{s=0}^h \mathbb{E}_t \Delta Y_{t+s}}{\sum_{s=0}^h \mathbb{E}_t \Delta G_{t+s}} \qquad \mu_T^h = \frac{\sum_{s=0}^h \mathbb{E}_t \Delta Y_{t+s}}{\sum_{s=0}^h \mathbb{E}_t \Delta T_{t+s}}$$

where Y_t and G_t are output and government spending relative to trend GDP. So an increase in government spending over h periods totalling 1 percent of trend GDP thus leads to an increase in cumulative output over the same period equivalent to μ_G^h percent of trend GDP.

As we estimate a more general fiscal shock, which includes both tax and spending changes, we cannot compute these multipliers individually. However, we can do this exercise in reverse. That is, taking as given estimates of multipliers from the literature, we can compute the output response that would be implied by the tax and spending profiles. So for fixed values of μ_T^h, μ_G^h we can compute:

$$\mu_Y^h = \mu_G^h \sum_{s=0}^h \mathbb{E}_t \Delta G_{t+s} + \mu_T^h \sum_{s=0}^h \mathbb{E}_t \Delta T_{t+s} \tag{11}$$

If we have identified a fiscal shock, and the multipliers estimated in the literature are correct, then this quantity should be close to our cumulative estimated output response, $\sum_{s=0}^h \mathbb{E}_t \Delta Y_{t+s}$. This fact allows us to construct a test of whether of our fiscal shock labeling is consistent with the estimates in the literature. We substitute values from several papers for the tax and spending multipliers into equation (11) and replace the conditional expectations for changes in tax and spending with out estimated impulse responses to compute μ_Y^h at various horizons.[7]

Table 2 lists the values of the multipliers we use and their sources. Perhaps the most similar exercise is Lewis (2021), who also identifies the entire set of structural shocks and must label fiscal shocks based on estimated IRFs. To this we add results from three classic

[6] Notable papers using this definition include Mountford and Uhlig (2009), Farhi and Werning (2016), Hagedorn et al. (2019), and others mentioned in the main text. See Batini et al. (2014) or Ramey (2016) for an overview. Other definitions of multipliers are sometimes used; for example, Blanchard and Perotti (2002) measure the multiplier using the peak output response, while Leeper et al. (2017) use real interest rates to discount future quantities.

[7] This is a benefit of scaling these variables relative to trend GDP prior to estimating our VARs (see Section 4.1). It means that the impulse responses are already in the appropriate units.

papers: Blanchard and Perotti (2002), Ramey (2016), and Romer and Romer (2010).[8] As the latter two estimate only spending and tax multipliers separately, we combine them. To these, we add the well-known estimates of Caldara and Kamps (2017) who use two approaches to estimate dynamic tax and spending multipliers. We also consider two recent estimates of the spending multiplier – Ricco (2015) and Ben Zeev and Pappa (2017) – again supplementing them with tax multipliers from Romer and Romer (2010).

The individual points in Figure 4 the show the corresponding literature-consistent output responses, μ_Y^h, for each of these estimates. This is compared to our estimated cumulative output response, both point estimates (solid lines) and bootstrapped confidence intervals (dashed and dotted). The agreement with the Lewis (2021) estimates is remarkably close. Indeed, we should emphasize that the solid line and the points are from entirely separate calculations. Ex ante, there is nothing which necessarily says that these should line up. This close agreement suggests that is shock very similar to the fiscal shock identified by Lewis (2021), validating our interpretation. The remaining estimates are generally within or close to the confidence intervals for our estimates, especially at shorter horizons, and so further buttressing further our interpretation. Of course, the implied output response using the Blanchard and Perotti (2002) multipliers is a little large. However, this reflects the fact that they simply find multipliers which are much larger than those measured in more recent work.

4.3.3 Validating the Monetary Policy Shock

Here we validate our claim that the second shock in Figure 3 can be interpreted as a monetary policy shock. Our approach is to again to show that the average impulse responses look quantitatively similar to those in the literature.

Specifically, we utilize the monetary policy shocks identified by Bauer and Swanson (2022), who isolate the unanticipated component of monetary policy changes using high frequency variation around both FOMC announcements and speeches by the Fed chair, orthogonalized to asset prices. In order to compare like with like, we aggregate the Bauer-Swanson shocks at the quarterly frequency and estimate their effects in a VAR with our baseline data. Figure 5 compares the impulse responses from their shocks against ours. The high-frequency shocks have very similar quantitative effects on real activity, taxes, and spending. The impact on interest rates is a little different, although the asymmetry in the bootstrap means that for most periods, the responses are within the 60 percent media-centered confidence interval. Effects on inflation, however, are somewhat stronger.

[8] Blanchard and Perotti (2002) and Romer and Romer (2010) do not report their estimates as cumulative multipliers, so in order to compare with the other studies, we use the multipliers re-estimated by Lewis (2021) using Blanchard and Perotti's method, and the multipliers re-estimated by Favero and Giavazzi (2012) using Romer and Romer's method.

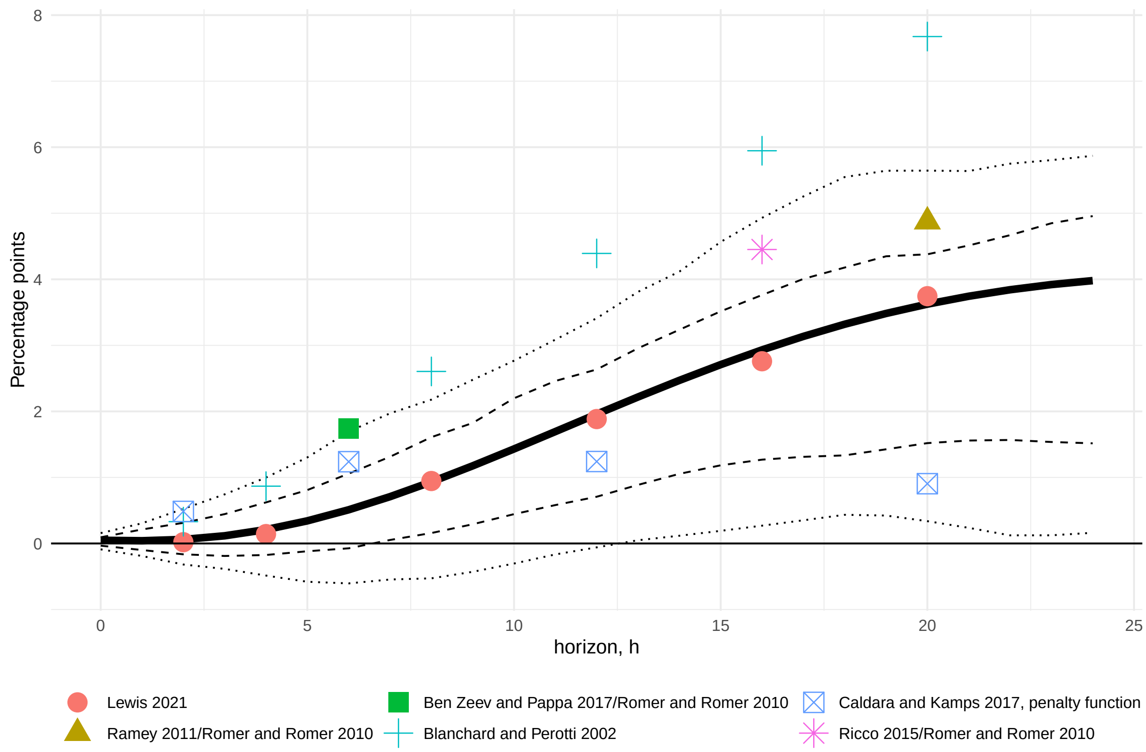

Figure 4: Cumulative output response: estimated versus implied by multipliers from the literature

The solid line is the point estimate and the dashed and dotted lines show respectively the $5^{th} - 95^{th}$ and $20^{th} - 80^{th}$ percentile ranges from a bootstrap simulation with $N^{sim} = 500$ replications. The points show the cumulative output responses, μ_Y^h, implied by our estimated tax and spending responses if the multipliers were those in the literature, summarized in Table 2.

Source	h	μ_G^h	μ_T^h
Lewis 2021	2	0.56	0.03
	4	0.57	-0.09
	8	0.57	-0.71
	12	0.64	-1.33
	16	0.76	-1.77
	20	0.87	-2.06
Blanchard and Perotti 2002	2	0.61	-0.64
	4	0.60	-1.01
	8	0.56	-2.28
	12	0.60	-3.63
	16	0.70	-4.69
	20	0.80	-5.41
Ramey 2011/Romer and Romer 2010	20	1.20	-2.60
Caldara and Kamps 2017, penalty function	2	0.05	-1.05
	6	0.35	-1.20
	12	0.55	-0.80
	20	0.25	-0.45
Ben Zeev and Pappa 2017/Romer and Romer 2010	6	2.40	-1.25
Ricco 2015/Romer and Romer 2010	16	1.50	-2.60

Table 2: Tax and spending multipliers from the literature

Table 2 shows the values of the tax and spending multipliers used to calculate μ_Y^h, the implied cumulative output response from the tax and spending responses for the fiscal shock. Where a pair of papers is cited, the former is used to calculate the spending multiplier, μ_G^h, and the latter the tax multiplier, μ_T^h. The cumulative Blanchard and Perotti (2002) multipliers are those reported by Lewis (2021), and the cumulative Romer and Romer (2010) multipliers are those reported by Favero and Giavazzi (2012).

4.4 The Importance of News Versus Surprises in Macroeconomic Fluctuations

Having calculated and labeled the structural shocks, ad validated the interpretation of the two policy shocks, we now decompose each into its news and surprise components

Figure 6 splits out the point estimate impulse responses to the average shocks into their news and surprise components. This shows that the importance of news and surprises varies considerably across different shocks. For instance, supply and demand shocks are driven more by news and surprises respectively. This accords with the common view of demand shocks as relatively fast-moving and harder to predict. Likewise, fiscal policy appears on average a larger surprise component than monetary policy, for which surprises seem generally to be more important.

To investigate this issue in a little more depth, we construct an explicit variance decomposition for all the variables and shocks in our model. It is relatively straightforward to

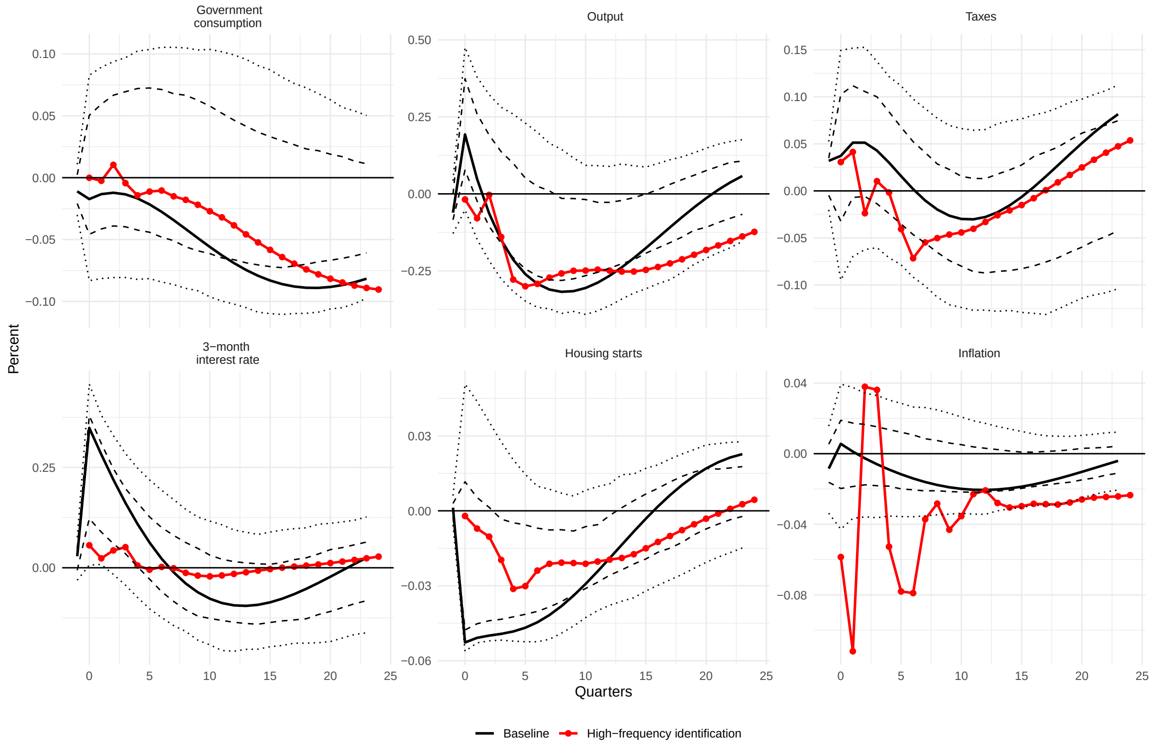

Figure 5: Estimated IRFs to Monetary Shocks

Figure shows estimated impulse responses to a monetary policy shock from our baseline compared to those computed using the high-frequency monetary policy shocks of Bauer and Swanson (2022). To match samples and specification, the line labelled "high-frequency identification" line reports the results from estimating a four-lag VAR with the same variables and coverage as our baseline model, extended to including the Bauer and Swanson (2022) shocks and where the impulse responses are computed from a Cholesky decomposition with the monetary shock ordered first. The solid line labeled "Baseline" is the point estimate and the dashed and dotted lines show respectively the $5^{th} - 95^{th}$ and $20^{th} - 80^{th}$ percentile ranges from a bootstrap simulation with $N^{sim} = 500$ replications.

show that the $h-$step ahead forecast error variance can be written as the sum of contributions from the news and surprise components of each of the structural shocks. In Appendix C, we work out this decomposition for the general case. But when $M = 1$, this becomes:

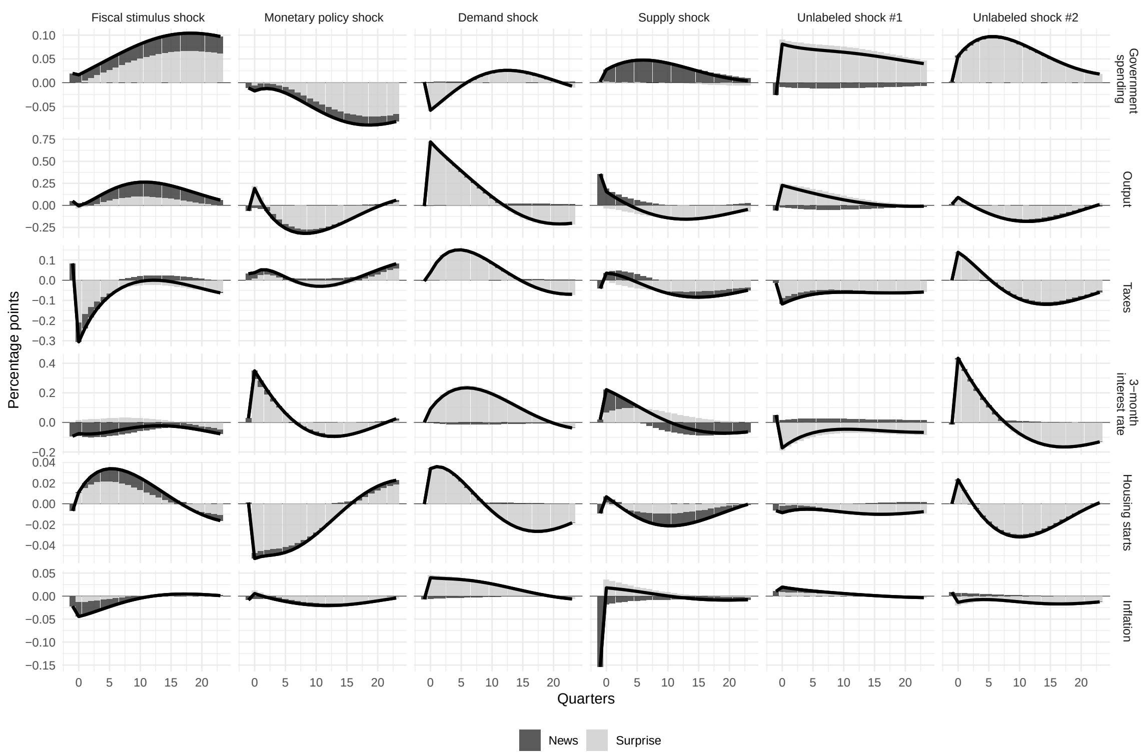

Figure 6: Estimated IRFs to Structural Shocks

The impulse response functions are plotted to an average unit structural shock, calculated as in Section 3.4. The dark and light gray bars capture the relative contribution of news and surprise to an average structural shock. For government consumption, output, and taxes, units are percentage points relative to trend lagged output. For inflation, interest rates, and housing starts, units are annualized percentage points relative to own-variable lagged trend.

$$MSE_t x_{t+h} = \sum_{j=1}^{N} \left(\sum_{s=1}^{h} B^{h-s} (A_j A_j')(B')^{h-s} \right) \sigma_{u,j}^2$$
$$+ \sum_{j=1}^{N} \left(\sum_{s=1}^{h} 1_{h>1} B^{h-s-1} \left(A_j A_j' + B(C_j A_j') + (A_j C_j') B' \right.\right.$$
$$\left.\left. + B(C_j C_j') B' \right) (B')^{h-s-1} + (C_j C_j') \right) \sigma_{v,j}^2 \tag{12}$$

where A_j and C_j are the jth columns of matrices A and C respectively and $1_{h>1}$ is an indicator function that is 1 if $h > 1$ and 0 otherwise. Note that because this is linear in the variances of each of the news and surprise shocks (the $\sigma_{u,j}^2$ and $\sigma_{v,j}^2$), this can be interpreted as an additive decomposition of the total variance with each term representing the contribution from each shock.

Table 3 reports this variance decomposition. For most variables, both news and surprises play an important role. In general, news seems to account for a smaller share of variance, but it is still not trivial – around 25 to 35 percent. This generally large role for news is consistent with broad themes int he literature. Empirical studies of news following Beaudry and Portier (2006) and Barsky and Sims (2011) broadly find large roles for news to explain business cycles. These types of papers associate news with forecast errors about technology; with our identification strategy, we can go further and find news associated with the entire set of structural shocks. One important exception is that inflation. There, news matters more, accounting for around 67 percent of fluctuations 6 years ahead. This is principally driven by news about demand – consistent with the idea that inflation is driven by forward-looking agents responding to changes in aggregate demand. Indeed, demand shocks are the largest or second largest contributor to the variation in all variables, something that seems reasonable for a business-cycle frequency analysis if policy shocks are generally stabilizing.

The relative importance of news is much less stable across shocks, however. It is hard to summarise the relative importance of news for each shock – different variables have different variances, and so it is not obvious how to create a summary measure of the importance of news for each shock across a range of outcomes. But in general, fiscal shocks are the most "newsy" likely reflecting the long lags in implementing fiscal policies. The minority share of news in supply shocks might seem a little curious given its apparent importance in Figure 6. However, this is largely due to the fact that new about supply shocks typically have offsetting effect at different horizons, limiting their impact as a long-run driver of macroeconomic fluctuations.

Some variable-shock-specific points are also worth highlighting here. Notably, monetary

Variable	Type	Fiscal stimulus	Mon. policy	Demand	Supply	Unlabeled #1	Unlabeled #2	Total
Gov. spending	News	7.9	7.8	6.0	5.2	3.0	0.6	30.5
	Surprise	17.9	10.6	2.0	0.2	22.5	16.3	69.5
	Total	25.7	18.5	8.0	5.3	25.5	16.9	100.0
Output	News	8.7	2.8	16.7	3.2	2.5	2.7	36.5
	Surprise	2.0	9.6	28.1	15.3	5.5	3.0	63.5
	Total	10.7	12.4	44.8	18.5	8.0	5.6	100.0
Taxes	News	5.1	5.2	7.5	1.8	2.2	3.9	25.7
	Surprise	19.4	2.1	14.6	17.3	7.6	13.4	74.3
	Total	24.5	7.3	22.0	19.1	9.8	17.3	100.0
3-month interest rate	News	6.4	2.1	16.9	4.8	2.2	1.9	34.4
	Surprise	0.9	8.2	16.6	11.2	7.3	21.4	65.6
	Total	7.3	10.2	33.5	16.1	9.6	23.3	100.0
Housing starts	News	4.8	4.0	5.4	2.7	1.6	2.5	21.0
	Surprise	10.6	30.5	17.5	5.6	2.6	12.2	79.0
	Total	15.4	34.6	22.9	8.3	4.1	14.7	100.0
Inflation	News	6.0	2.4	33.6	19.3	1.7	4.8	67.8
	Surprise	0.8	2.2	10.8	14.7	0.4	3.2	32.2
	Total	6.8	4.6	44.4	34.0	2.1	8.0	100.0

Table 3: Forecast error variance decomposition, 24 quarters ahead

The forecast error decomposition shows for each variable in percent the fraction of the overall forecast error variance attributable to each shock, split into the news and surprise components. Totals are shown in the right hand column. The news and surprise components sum to 100 for each variable. The "Avg. share" row is calculated as the unweighted average by shock of the share of variance for each variable due to news or surprise for each variable. It is thus a rough summary measure of the relative importance of news and surprise for a given shock across all variables.

policy shocks only drive a small amount of the variance in interest rates. This would be true if monetary policymakers generally adhere to a policy rule which responds to other shocks. The same is not true for fiscal variables, which are predominantly driven by policy changes and, in the case of taxes, supply and demand. Housing starts, a particularly forward-looking measure of real activity, are most affected by monetary policy and demand shocks.

While news is important for all time series, it is heterogeneous by the type of shock. For example, the effects of fiscal stimulus shocks are disproportionately due to news, especially for the output response.. In contrast, the effects of monetary policy shocks are almost exclusively due to surprises. Other shocks are mixed.

4.5 Policy shocks: News versus surprise

We now turn to the differential impact of policy news and surprise shocks, something that seems likely to be of interest to policymakers. Figure 7 presents separately the impulse responses to one-standard-deviation news and surprise innovations for the policy shocks, as defined in Section 3.4. The interpretation of a news shock is that in period 1, it is revealed that there will be a surprise shock in period 2. The news impulse therefore combines both the anticipation of the policy change in period 2 and its realized impact.

The news and surprise components of the fiscal shock are broadly similar. In both cases, increases in spending and reductions in taxes lead to an increase in output which peaks three to four years later. If anything, the policy impact of a news shock is a little more front-loaded. Upon announcement of the expansion starting in period 2, output rises immediately. Taxes also increase by a similar amount. This make some intuitive sense – if an announcement of a tax cut tomorrow likely expands activity today, contemporaneous revenues will likely rise. Government spending also increases immediately, suggesting an endogenous anticipation effect of government expenditure to expectations of a future expansion (for example, if an announcement of new projects tomorrow leads to preparatory work today). The peak response of output also arrives a little sooner for a news shock, consistent with a more front-loaded impact.

The news and surprise components of the monetary policy shock are less clear, with wider confidence intervals. However, despite similar impacts on the interest rate, the impact of a news shock on activity is more immediate. Indeed, the well known "liquidity effect" – whereby activity and inflation increase temporarily on impact of a monetary policy tightening appears to be a feature only of surprises and not of news shocks.

5 Counterfactual Policy

This section applies the Wolf and McKay (2022) method to study counterfactual policy rules.

5.1 Method

One of the key observations in Wolf and McKay (2022) is that in a world where news shocks matter, policymakers are able to pursue their goals not just through their current actions but also through news about their future actions. They exploit this insight to address a long-standing critique of the usefulness of VARs for computing purely empirical policy counterfactuals: that they are subject to the Lucas critique (Lucas Jr, 1976).

For intuition, imagine that one were to able to perfectly identify the impact of a monetary policy shock using a VAR and wanted to understand what would have happened if policy had followed a different rule, one that perfectly stabilized inflation. One possibility, pioneered by Sims and Zha (2006), would be to use the estimated impulse responses for inflation from the monetary shock to compute the sequence of policy innovations which would have stabilized inflation period-by-period. The challenge to this approach is that the policy realized ex post is inconsistent with agents' expectations.

Wolf and McKay (2022) show that identification of news shocks is sufficient to overcome this challenge in a relatively large class of commonly used macro models. The intuition is that policymakers can implement a different rule not just through a surprise today but by also communicating their future actions as news shocks. As a result, agents' ex ante

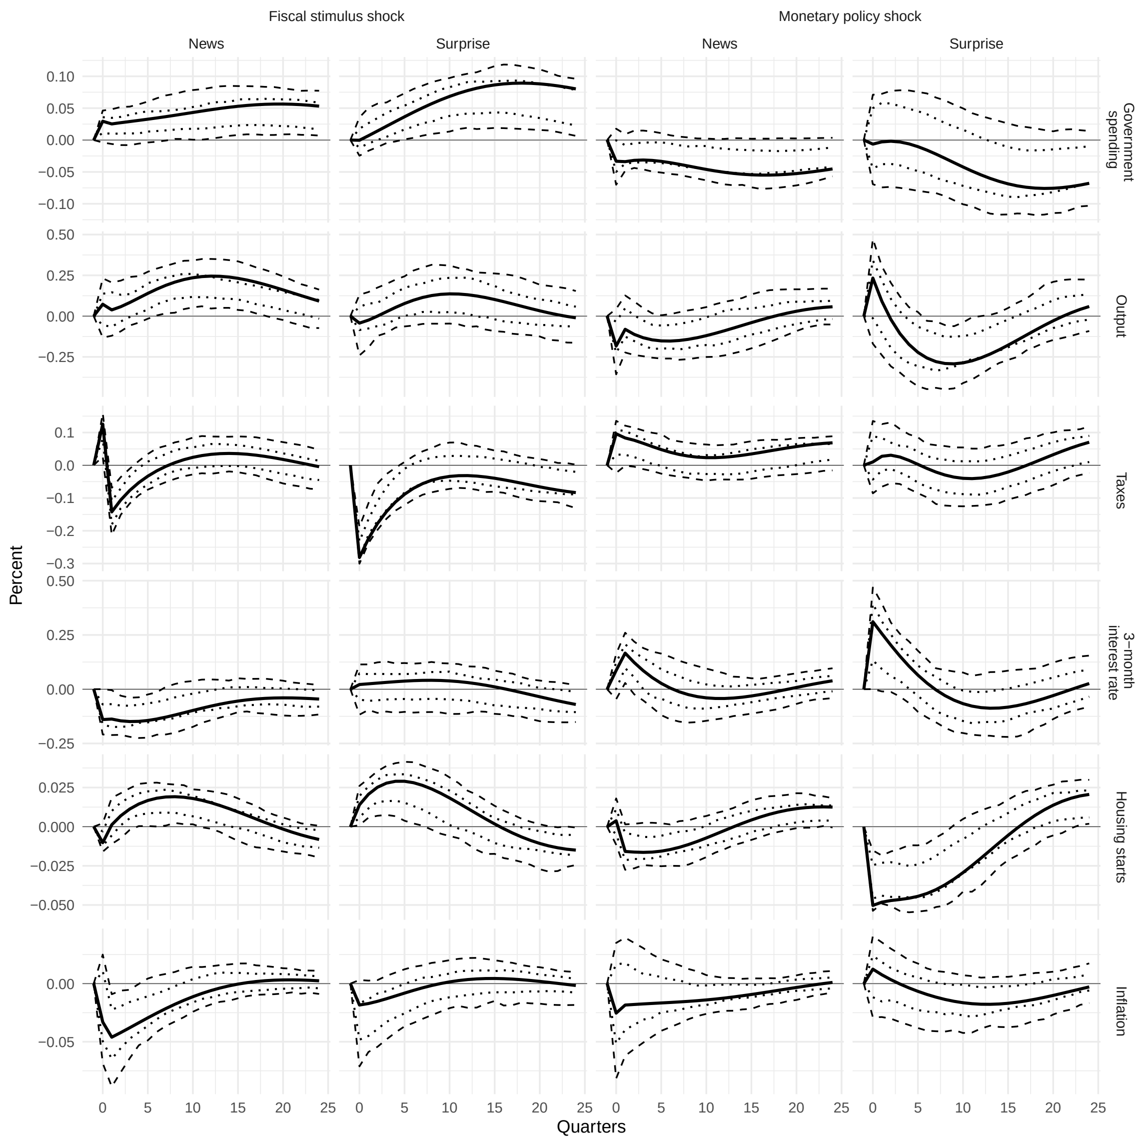

Figure 7: Estimated IRFs to Policy Shocks

The impulse response functions are to one-standard-deviation news and surprise components of each fiscal shock, calculated as in Section 3.4. The solid line is the point estimate and the dashed and dotted lines show respectively the $5^{th} - 95^{th}$ and $20^{th} - 80^{th}$ percentile ranges from a bootstrap simulation with $N^{sim} = 500$ replications. For government consumption, output, and taxes, units are percentage points relative to trend lagged output. For inflation, interest rates, and housing starts, units are annualized percentage points relative to own-variable lagged trend.

beliefs are then consistent with the ex post policy rule. This in turn means that policy counterfactuals can be estimated in three steps: 1) identifying news shocks, 2) compute the sequence of news and surprises which would implement the counterfactual policy, 3) use the estimated impulse responses to calculate the responses of the macroeconomy to that rule. So far, this paper has been about the first of these. We now turn to the remaining ones.

To apply this to our setting, we start by classifying our estimated shocks as either policy shocks (the fiscal stimulus and monetary policy shocks) or as others (demand, supply, and the unlabeled shock). We then consider one-at-a-time the problem of the policymakers in control of each policy shock, assuming that they wish to minimize some loss function.

Specifically, assume that the policymaker controls both the surprise and the news for shock g, denoted u_t^g and v_t^g. We denote the vectors of non-policy shocks by u_t^{-g} and v_t^{-g}. We consider linear policy counterfactuals which can be written as:

$$\begin{bmatrix} u_t^g \\ v_t^g \end{bmatrix} = \alpha \begin{bmatrix} u_t^{-g} \\ v_t^{-g} \end{bmatrix} \tag{13}$$

where α is a $2 \times 2(n-1)$ matrix recording how the policymaker responds to the other structural shocks.

Let the impulse responses to surprise and news under this rule be denoted by $\psi_u(h)$ and $\psi_v(h)$. Then:

$$\begin{bmatrix} \psi_u(h) \\ \psi_v(h) \end{bmatrix} = \begin{bmatrix} \phi_u^{-g}(h) \\ \phi_v^{-g}(h) \end{bmatrix} + \alpha \begin{bmatrix} \phi_u^{g}(h) \\ \phi_v^{g}(h) \end{bmatrix}$$

We then assume that the policymaker aims to minimize a linear loss function:

$$\min ||x_t F||$$

for some matrix F. This loss function could be a direct loss due to macroeconomic fluctuations (e.g. departures from an inflation target) or it could be deviations from a specific policy rule (e.g. a Taylor rule). In either case, we follow Wolf and McKay (2022) by computing α to minimize this loss. A sufficient condition for this is to minimize the loss function on the impulse responses, as these are just the building blocks of the linear model. We thus rewrite the problem as:

$$\min \left|\left| \begin{bmatrix} \psi_u(h) \\ \psi_v(h) \end{bmatrix} F \right|\right| = \min \left|\left| \begin{bmatrix} \phi_u^{-g}(h) \\ \phi_v^{-g}(h) \end{bmatrix} F + \alpha \begin{bmatrix} \phi_u^{g}(h) \\ \phi_v^{g}(h) \end{bmatrix} F \right|\right|$$

When the metric $||\cdot||$ is a sum of squares, this can be solved by estimating α from the regression:

$$\begin{bmatrix} \phi_u^{-g}(h) \\ \phi_v^{-g}(h) \end{bmatrix} F = -\alpha \begin{bmatrix} \phi_u^{g}(h) \\ \phi_v^{g}(h) \end{bmatrix} F + \epsilon_h \tag{14}$$

5.2 Counterfactual Exercises

We study two types of counterfactual policies: active policies which aim to moderate business cycles, and passive policies which hold policy instruments fixed. In both cases, we compare and contrast fiscal and monetary policy.

5.2.1 Output Stabilization

In this section, we study how different policy instruments can be used for output stabilization. For each policy instrument, we select the linear combination of news and surprise shocks that minimize the variance in detrended output. This gives a different policy response for each of the remaining 10 shocks (for each policy instrument there are 5 remaining structural shocks, each with a news and noise component.) In the plots that follow, we only plot the counterfactual impulse responses to the non-policy structural shocks: supply, demand, and the unlabeled shocks.

Figure 8 shows our results when fiscal stimulus is used to moderate business cycles. The path of government spending changes most from the baseline in three cases: news about demand is ordinarily expansionary, so government spending must contract to moderate output, supply surprises are ordinarily contractionary, so government spending must increase after these surprises, and similarly for surprises to the second unlabeled shock. In most cases, government spending is able to stabilize output to have nearly no response. The main exceptions are demand surprises and supply news, which have IRF shapes that cannot be moderated bya combination of fiscal policy news and surprises. However, there are costs to this stabilization. In some cases, inflation or interest rates become more volatile, particularly for the shocks to which the fiscal variables must adjust the most. Focusing on output stabilization also does not perfectly moderate the other measure of real activity; housing starts become somewhat less volatile, but not by nearly as much as output, and certainly not after supply news shocks.

Figure 9 shows our results when monetary policy is used to stabilize output. Monetary policy is more effective than fiscal stimulus in response to some shocks, in particular demand suprrises and supply news, but less effective in repsonse to others. This suggests a clear role for monetary and fiscal coordination. The general prescription is for monetary policy to tighten relative to the baseline after the expansionary surprise shocks, as conventional wisdom suggests. In most cases, using monetary policy to stabilize output does not lead to substantially more inflation.

5.2.2 Passive Policies

All policy shocks move over the business cycle and lead to expansions and contractions in the other time series. In this section, we consider counterfactuals where the policy instruments are as fixed as possible.

Figure 10 plots the counterfactual impulse responses when government spending is nearly fixed. We are able to construct combinations of shocks to get close to zero government spending response in most cases except the supply news and unlabeled surprises, which are small contributors to business cycle volatility. When government spending is fixed, most time series are much more volatile, output in particular.Most noticeably, unlabeled surprises are much more expansionary and inflationary. And demand surprises are more expansionary, suggesting that the baseline government spending policy is actively attenuating the largest driver of business cycles (Table 3).

Figure 11 plots our results when monetary policy is selected to approximate an interest rate peg as well as possible. The Wolf and McKay (2022) method assumes that counterfactual equilibria exist and are unique; this famously is not the case for interest rate pegs in New Keynesian models, so there may be other possible equilibria when such a policy is implemented. Similar to Wolf and McKay, we find that an interest rate peg does not substantially affect inflation dynamics. However, we find larger changes to output impulse responses. Still, Wolf and McKay specifically study counterfactual responses to the Ben Zeev and Khan (2015) news shock about investment costs, which does not have a clear analog in our VAR.

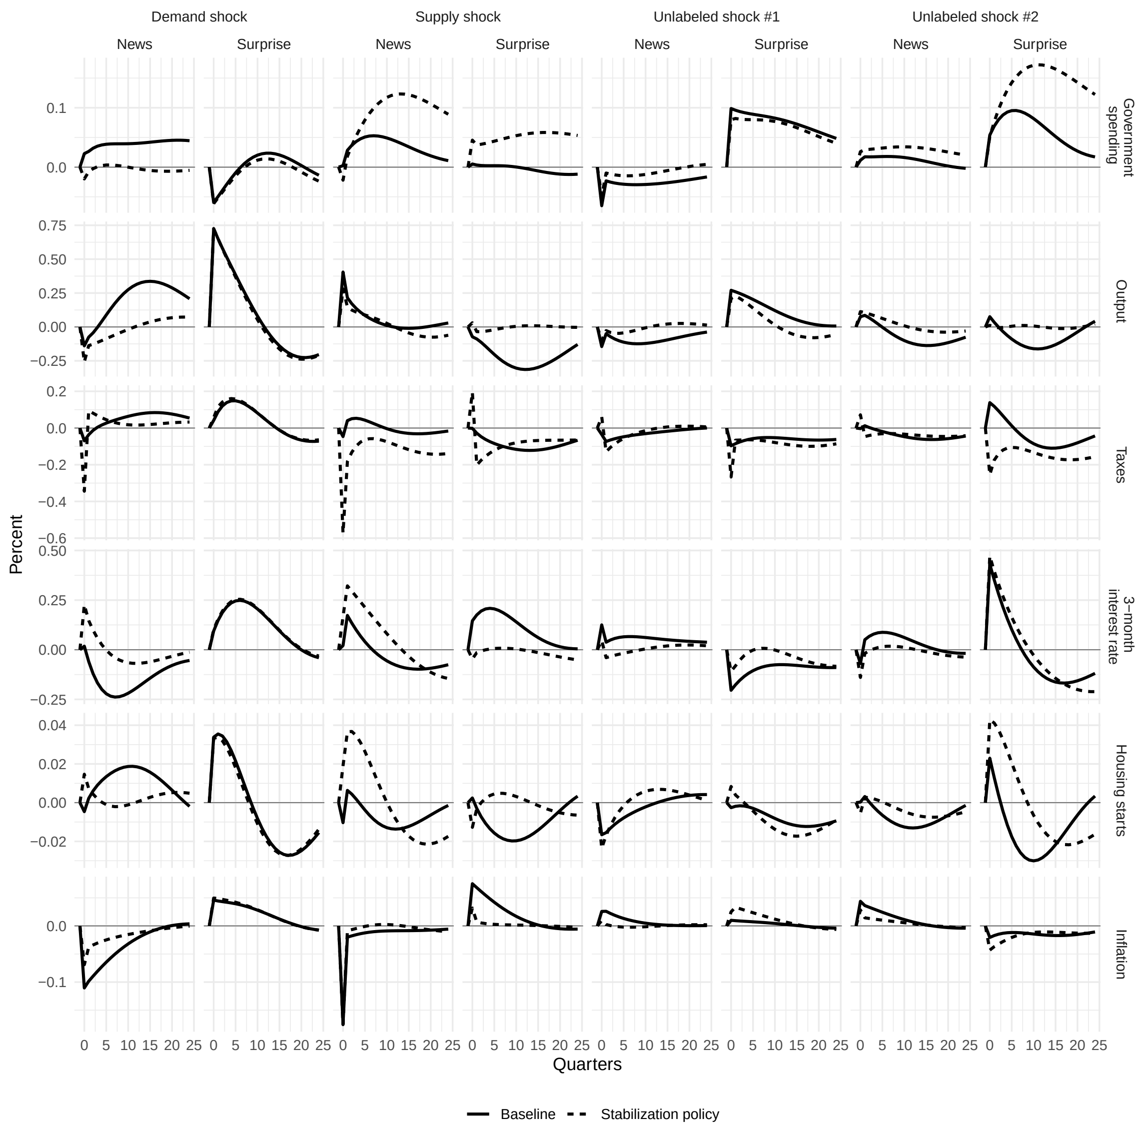

.

Figure 8: Output Stabilization with Government Spending

Baseline time series impulse responses to news and surprise components of the four non-policy structural shocks under the prevailing baseline rule are plotted as solid lines. The best feasible approximations using government spending for output targeting are plotted as dashed lines, computed following equation (14)

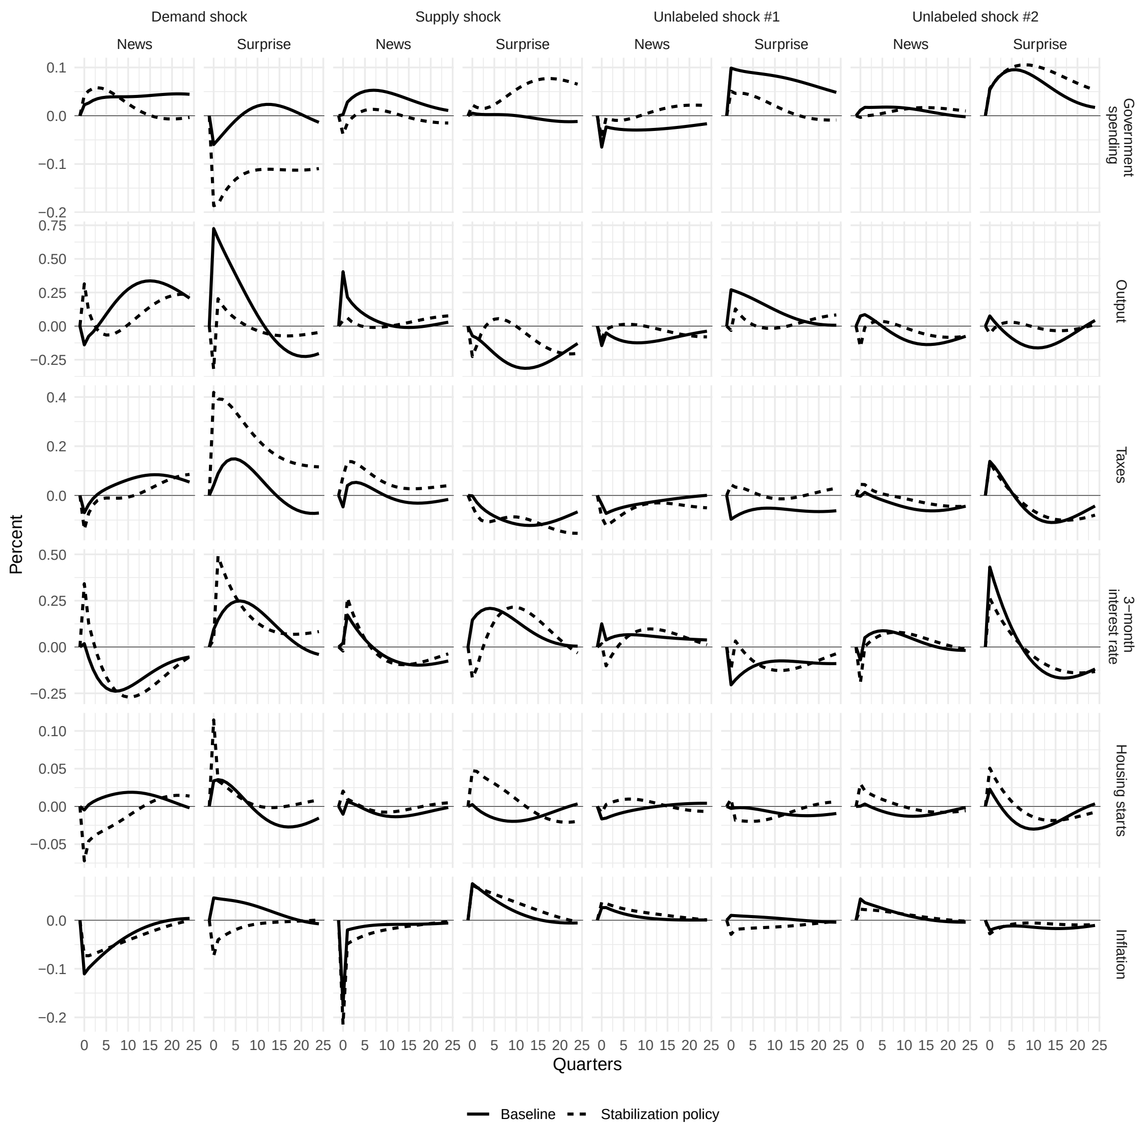

.

Figure 9: Output Stabilization with Monetary Policy

Baseline time series impulse responses to news and surprise components of the four non-policy structural shocks under the prevailing baseline rule are plotted as solid lines. The best feasible approximations using monetary policy for output targeting are plotted as dashed lines, computed following equation (14)

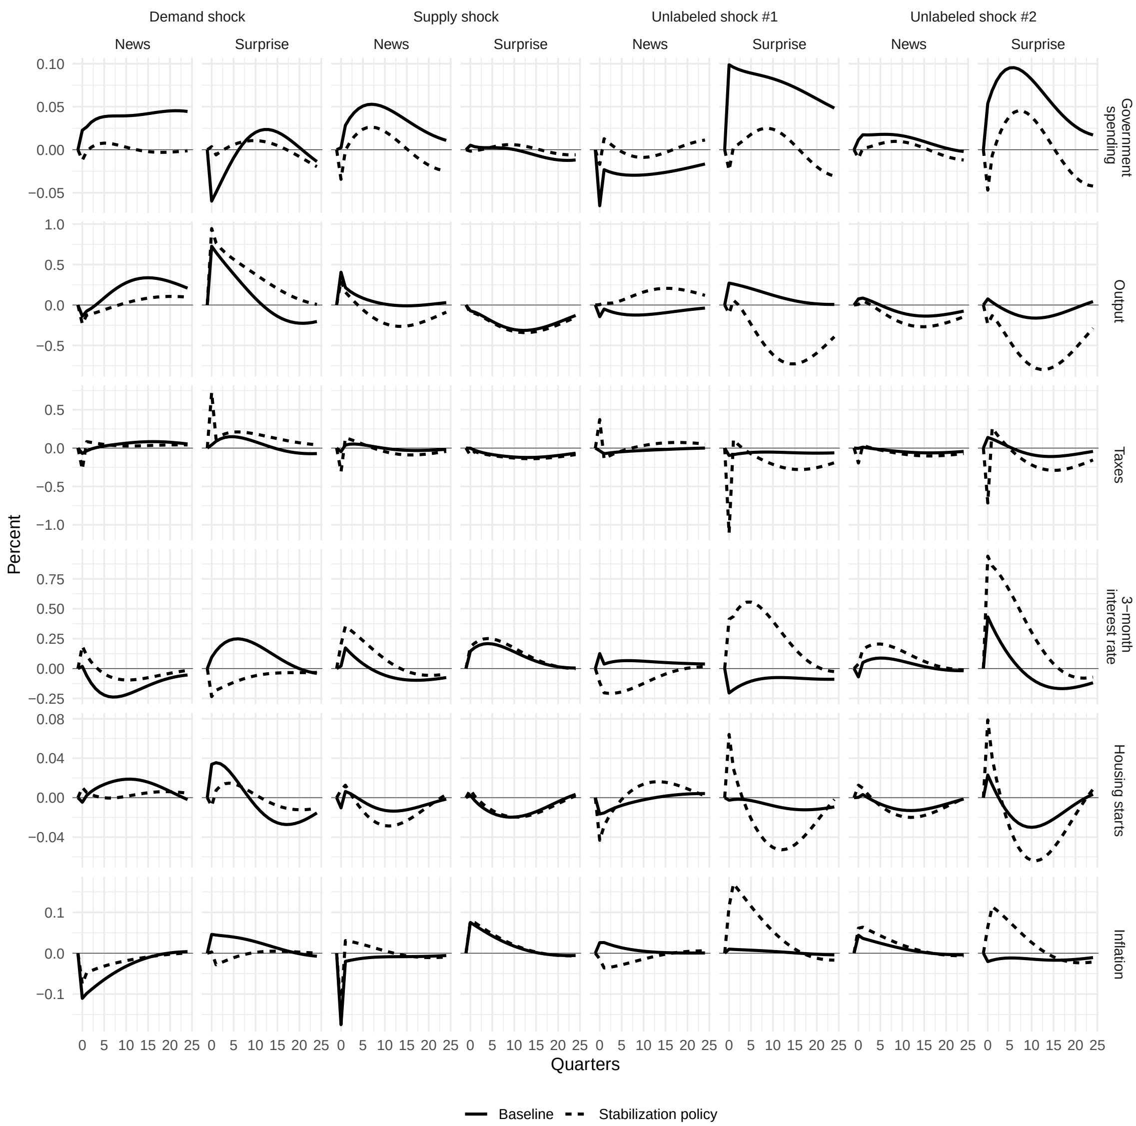

.

Figure 10: Fixed Government Spending

Baseline time series impulse responses to news and surprise components of the four non-policy structural shocks under the prevailing baseline rule are plotted as solid lines. The best feasible approximations to fixed government spending are plotted as dashed lines, computed following equation (14)

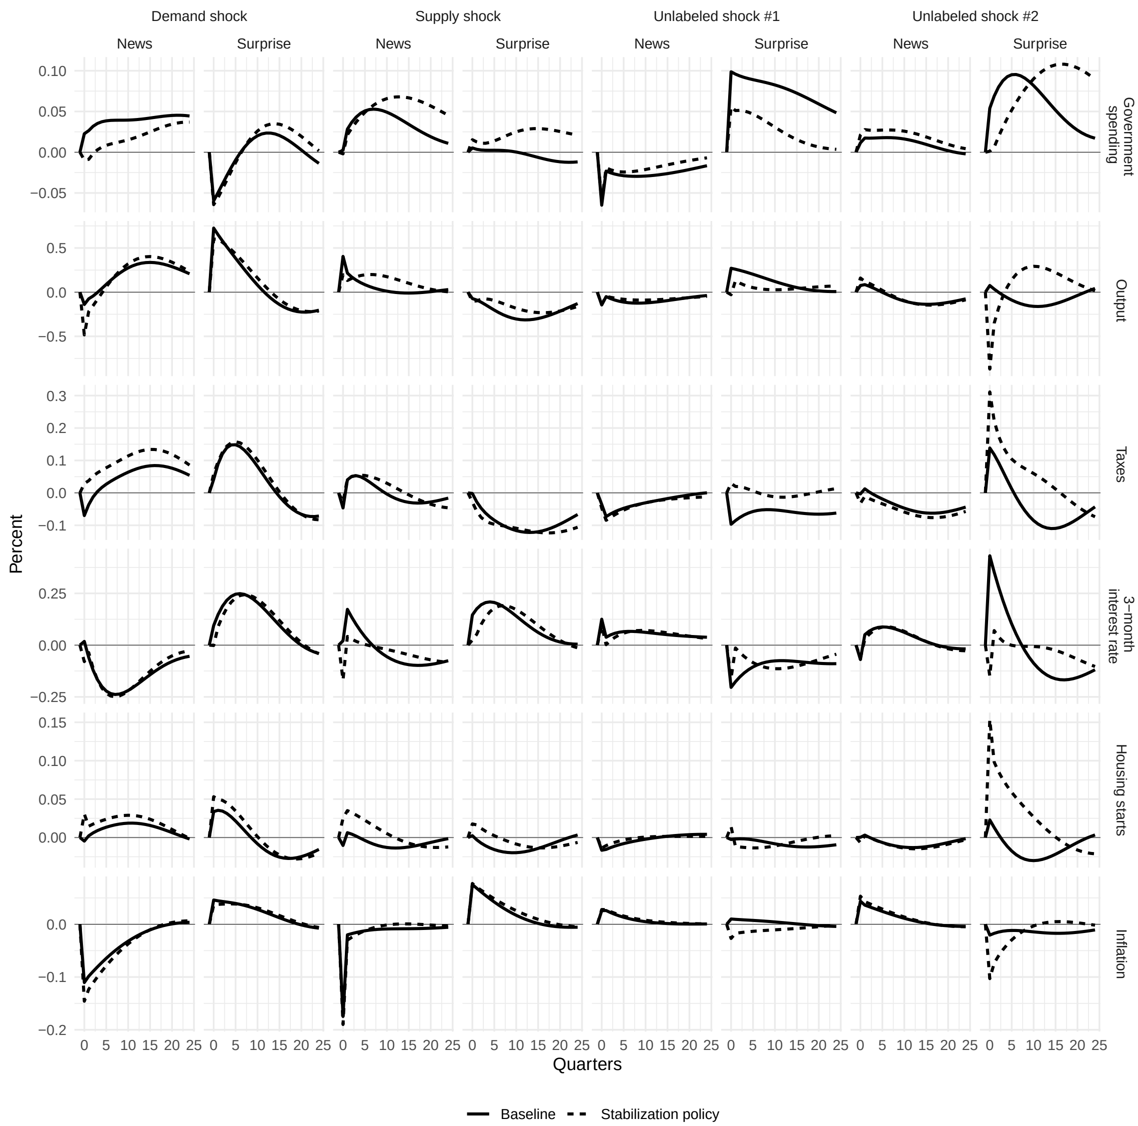

.

Figure 11: Monetary Policy with an Interest Rate Peg

Baseline time series impulse responses to news and surprise components of the four non-policy structural shocks under the prevailing baseline rule are plotted as solid lines. The best feasible approximations using monetary policy to enforce an interest rate peg are plotted as dashed lines, computed following equation (14)

References

Auerbach, Alan J. and Yuriy Gorodnichenko, "Fiscal Multipliers in Recession and Expansion," in "Fiscal Policy after the Financial Crisis," University of Chicago Press, February 2012, pp. 63–98.

Barsky, Robert B. and Eric R. Sims, "News shocks and business cycles," *Journal of Monetary Economics*, April 2011, *58* (3), 273–289.

_ **and** _ , "Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence," *American Economic Review*, June 2012, *102* (4), 1343–1377.

Batini, Nicoletta, Luc Eyraud, Lorenzo Forni, and Anke Weber, *Fiscal multipliers: Size, determinants, and use in macroeconomic projections*, International Monetary Fund, 2014.

Bauer, Michael D. and Eric T. Swanson, "A Reassessment of Monetary Policy Surprises and High-Frequency Identification," April 2022.

Beaudry, Paul and Franck Portier, "Stock Prices, News, and Economic Fluctuations," *American Economic Review*, September 2006, *96* (4), 1293–1307.

Blanchard, Olivier and Roberto Perotti, "An Empirical Characterization of the Dynamic Effects of Changes in Government Spending and Taxes on Output*," *The Quarterly Journal of Economics*, November 2002, *117* (4), 1329–1368.

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

Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer, "Overreaction in Macroeconomic Expectations," *American Economic Review*, September 2020, *110* (9), 2748–2782.

Born, Benjamin, Falko Juessen, and Gernot J. Müller, "Exchange rate regimes and fiscal multipliers," *Journal of Economic Dynamics and Control*, February 2013, *37* (2), 446–465.

Caggiano, Giovanni, Efrem Castelnuovo, Valentina Colombo, and Gabriela Nodari, "Estimating Fiscal Multipliers: News from a Non-linear World," *The Economic Journal*, May 2015, *125* (584), 746–776.

Caldara, Dario and Christophe Kamps, "The Analytics of SVARs: A Unified Framework to Measure Fiscal Multipliers," *The Review of Economic Studies*, July 2017, *84* (3), 1015–1040.

Cascaldi-Garcia, Danilo, "Forecast revisions as instruments for news shocks," *International Finance Discussion Paper*, 2022, (1341).

Chahrour, Ryan and Kyle Jurado, "News or Noise? The Missing Link," *American Economic Review*, July 2018, *108* (7), 1702–1736.

_ **and** _ , "Recoverability and Expectations-Driven Fluctuations," *The Review of Economic Studies*, March 2021, p. rdab010.

Coibion, Olivier and Yuriy Gorodnichenko, "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," *American Economic Review*, August 2015, *105* (8), 2644–2678.

Croushore, Dean and Simon van Norden, "Fiscal Forecasts at the FOMC: Evidence from the Greenbooks," *The Review of Economics and Statistics*, December 2018, *100* (5), 933–945.

Farhi, E. and I. Werning, "Fiscal Multipliers: Liquidity traps and currency unions," in John B. Taylor and Harald Uhlig, eds., *Handbook of Macroeconomics*, Vol. 2, Elsevier, January 2016, pp. 2417–2492.

Favero, Carlo and Francesco Giavazzi, "Measuring Tax Multipliers: The Narrative Method in Fiscal VARs," *American Economic Journal: Economic Policy*, May 2012, *4* (2), 69–94.

Fisher, Jonas D.M. and Ryan Peters, "Using Stock Returns to Identify Government Spending Shocks," *The Economic Journal*, May 2010, *120* (544), 414–436.

Forni, Mario and Luca Gambetti, "Government spending shocks in open economy VARs," *Journal of International Economics*, March 2016, *99*, 68–84.

Gouriéroux, Christian, Alain Monfort, and Jean-Paul Renne, "Statistical inference for independent component analysis: Application to structural VAR models," *Journal of Econometrics*, January 2017, *196* (1), 111–126.

Greenwood, Robin and Andrei Shleifer, "Expectations of Returns and Expected Returns," *The Review of Financial Studies*, March 2014, *27* (3), 714–746.

Hagedorn, Marcus, Iourii Manovskii, and Kurt Mitman, "The Fiscal Multiplier," February 2019.

Hamilton, James Douglas, *Time Series Analysis*, Princeton University Press, 1994.

Hirose, Yasuo and Takushi Kurozumi, "Identifying News Shocks with Forecast Data," *Macroeconomic Dynamics*, September 2021, *25* (6), 1442–1471. Publisher: Cambridge University Press.

House, Christopher L. and Matthew D. Shapiro, "Phased-In Tax Cuts and Economic Activity," *American Economic Review*, December 2006, *96* (5), 1835–1849.

Hyvärinen, Aapo, Kun Zhang, Shohei Shimizu, and Patrik O. Hoyer, "Estimation of a Structural Vector Autoregression Model Using Non-Gaussianity," *Journal of Machine Learning Research*, 2010, *11* (5), 1709–1731.

Jr, Robert E Lucas, "Econometric policy evaluation: A critique," in "Carnegie-Rochester conference series on public policy," Vol. 1 North-Holland 1976, pp. 19–46.

Lanne, Markku, Helmut Lütkepohl, and Katarzyna Maciejowska, "Structural vector autoregressions with Markov switching," *Journal of Economic Dynamics and Control*, February 2010, *34* (2), 121–131.

Leeper, Eric M., Nora Traum, and Todd B. Walker, "Clearing Up the Fiscal Multiplier Morass," *American Economic Review*, August 2017, *107* (8), 2409–2454.

_ , Todd B. Walker, and Shu-Chun Susan Yang, "Government Investment and Fiscal Stimulus in the Short and Long Runs," July 2009.

Lewis, Daniel J, "Identifying Shocks via Time-Varying Volatility," *The Review of Economic Studies*, November 2021, *88* (6), 3086–3124.

Lütkepohl, Helmut and Aleksei Netšunajev, "Structural vector autoregressions with heteroskedasticity: A review of different volatility models," *Econometrics and Statistics*, January 2017, *1*, 2–18.

Mertens, Karel and Morten O. Ravn, "Empirical Evidence on the Aggregate Effects of Anticipated and Unanticipated US Tax Policy Shocks," *American Economic Journal: Economic Policy*, May 2012, *4* (2), 145–181.

Mountford, Andrew and Harald Uhlig, "What are the effects of fiscal policy shocks?," *Journal of Applied Econometrics*, 2009, *24* (6), 960–992. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/jae.1079.

Ramey, V. A., "Chapter 2 - Macroeconomic Shocks and Their Propagation," in John B. Taylor and Harald Uhlig, eds., *Handbook of Macroeconomics*, Vol. 2, Elsevier, January 2016, pp. 71–162.

Ramey, Valerie A., "Identifying Government Spending Shocks: It's all in the Timing*," *The Quarterly Journal of Economics*, February 2011, *126* (1), 1–50.

_ , "Ten Years after the Financial Crisis: What Have We Learned from the Renaissance in Fiscal Research?," *Journal of Economic Perspectives*, May 2019, *33* (2), 89–114.

Ricco, Giovanni, "A new identification of fiscal shocks based on the information flow," 2015. Publisher: ECB Working Paper.

_ , Giovanni Callegari, and Jacopo Cimadomo, "Signals from the government: Policy disagreement and the transmission of fiscal shocks," *Journal of Monetary Economics*, September 2016, *82*, 107–118.

Rigobon, Roberto, "Identification Through Heteroskedasticity," *The Review of Economics and Statistics*, November 2003, *85* (4), 777–792.

Romer, Christina D. and David H. Romer, "The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks," *American Economic Review*, June 2010, *100* (3), 763–801.

Schmitt-Grohé, Stephanie and Martín Uribe, "What's News in Business Cycles," *Econometrica*, 2012, *80* (6), 2733–2764. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA8050.

Sentana, Enrique and Gabriele Fiorentini, "Identification, estimation and testing of conditionally heteroskedastic factor models," *Journal of Econometrics*, June 2001, *102* (2), 143–164.

Shapiro, Matthew D. and Mark W. Watson, "Sources of Business Cycle Fluctuations," *NBER Macroeconomics Annual*, January 1988, *3*, 111–148. Publisher: The University of Chicago Press.

Sims, Christopher A., "Macroeconomics and Reality," *Econometrica*, 1980, *48* (1), 1–48. Publisher: [Wiley, Econometric Society].

_ **and Tao Zha**, "Does Monetary Policy Generate Recessions?," *Macroeconomic Dynamics*, April 2006, *10* (2), 231–272. Publisher: Cambridge University Press.

Souleles, Nicholas S., "Expectations, Heterogeneous Forecast Errors, and Consumption: Micro Evidence from the Michigan Consumer Sentiment Surveys," *Journal of Money, Credit and Banking*, 2004, *36* (1), 39–72. Publisher: [Wiley, Ohio State University Press].

Uhlig, Harald, "What are the effects of monetary policy on output? Results from an agnostic identification procedure," *Journal of Monetary Economics*, March 2005, *52* (2), 381–419.

Wolf, Christian K. and Alisdair McKay, "What Can Time-Series Regressions Tell Us About Policy Counterfactuals?," August 2022.

Zeev, Nadav Ben and Evi Pappa, "Chronicle of A War Foretold: The Macroeconomic Effects of Anticipated Defence Spending Shocks," *The Economic Journal*, August 2017, *127* (603), 1568–1597.

_ **and Hashmat Khan**, "Investment-Specific News Shocks and U.S. Business Cycles," *Journal of Money, Credit and Banking*, 2015, *47* (7), 1443–1464. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/jmcb.12250.

A Forecast Cleaning Properties

A.1 Proof of Theorem 2

Proof. Equation (10) and the causal invertibility assumption imply that we can write the rational expectation as

$$f_t = H^f(L)^{-1}\tilde{f}_t - H^f(L)^{-1}H^x(L)x_t - H^f(L)^{-1}H^z(L)z_t - H^f(L)^{-1}H^u(L)u_t - H^f(L)^{-1}H^v(L)v_t$$

Lags of u_t and v_t can be written in terms of current and past rational forecasts and observables, per equation (8). Denote these representations with the invertible lag operator polynomials $u_t = M^u_x(L)x_t + M^u_f(L)f_t$ and $v_t = M^v_x(L)x_t + M^v_f(L)f_t$. The rational expectation becomes:

$$f_t = H^f(L)^{-1}\tilde{f}_t - H^f(L)^{-1}H^x(L)x_t - H^f(L)^{-1}H^z(L)z_t - M^u_x(L)x_t - M^u_f(L)f_t - M^v_x(L)x_t - M^v_f(L)f_t$$

$$= (I + M^u_f(L) + M^v_f(L))^{-1}\left(H^f(L)^{-1}\tilde{f}_t - (H^f(L)^{-1}H^x(L) + M^u_x(L) + M^v_x(L))x_t - H^f(L)^{-1}H^z(L)z_t\right)$$

which we simplify by defining the causal lag operator polynomials $\psi^{\tilde{f}}$, ψ^x, and ψ^z to collect coefficients, allowing us to write the rational expectation as

$$f_t = \psi^{\tilde{f}}(L)\tilde{f}_t + \psi^x(L)x_t + \psi^z(L)z_t \tag{15}$$

Consider the relationship between x_{t+1} and the lagged observables:

$$x_{t+1} = f_t + Au_{t+1} + Cv_{t+1}$$

$$= \psi^{\tilde{f}}(L)\tilde{f}_t + \psi^x(L)x_t + \psi^z(L)z_t + Au_{t+1} + Cv_{t+1}$$

u_{t+1} and v_{t+1} are orthogonal to current and past observables, so forecasting x_{t+1} by regressing on lags of $\tilde{f}_t$, x_t, and z_t recovers the rational expectation:

$$E[x_{t+1}|\{\tilde{f}_{t-j}, x_{t-j}, z_{t-j}\}_{j=0}^{\infty}] = E[f_t + Au_{t+1} + Cv_{t+1}|\{\tilde{f}_{t-j}, x_{t-j}, z_{t-j}\}_{j=0}^{\infty}]$$

$$= E[f_t|\{\tilde{f}_{t-j}, x_{t-j}, z_{t-j}\}_{j=0}^{\infty}]$$

which is given by equation (9). ■

A.2 Noisy Forecast Cleaning

When the conditions of Theorem 2 are not satisfied, the interpretation of our forecast cleaning becomes weaker, but still useful.

Instead of an ideal rational expectation conditional on all information in available to forecasters, our cleaned forecasts are the best unbiased forecasts given the observable time series and reported forecasts. The interpretation of news must change as well. Instead of the component of structural shocks that is anticipated by forecasters, news is now the component that can be forecasted by the VAR.

First, we modify equation (5) so that the structural VAR depends on expectations of future shocks $E_t[\epsilon_{t+1}]$ in general rather than the news component v_t explicitly. This expectation may include noise shocks or other confounders in addition to the structural v_t:

$$x_t = \sum_{j=1}^{m} B_j x_{t-j} + A\epsilon_t + CE_t[\epsilon_{t+1}]$$

Next modify equation (10) so that forecasts are now given by

$$\tilde{f}_t = H^x(L)x_t + H^z(L)z_t + H^u(L)u_t + H^v(L)v_t + H^\zeta(L)\zeta_t$$

Now the empirical forecasts $\tilde{f}_t$ are not deviations from some ideal rational expectation. Rather, they are just some linear combination of observables, structural shocks, and the noise shocks ζ_t.

The component of forecasts excluding the observable terms is

$$\xi_t \equiv H^u(L)u_t + H^v(L)v_t + H^\zeta(L)\zeta_t$$

Let $H^\xi(L)w_t^\xi$ denote the Wold decomposition of ξ_t, with w_t^ξ white noise. Forecasting x_{t+1} gives the cleaned forecast:

$$f_t = E[x_{t+1}|\Omega] = \sum_{j=1}^{m} B_j x_{t+1-j} + AE[\epsilon_{t+1}|\Omega]$$

$$= \sum_{j=1}^{m} B_j x_{t+1-j} + AE[\epsilon_{t+1}|\{\xi_{t-j}\}_{j=0}^{\infty}] = \sum_{j=1}^{m} B_j x_{t+1-j} + AE[\epsilon_{t+1}|w_t^\xi]$$

so we define our reduced form news $\tilde{v}_t$ as

$$\tilde{v}_t \equiv E[\epsilon_{t+1}|w_t^\xi]$$

$$= D_v H_0^{v\prime}\Sigma_{w^\xi}^{-1} w_t^\xi$$

where H_0^v is the contemporaneous coefficient matrix in the $H^v(L)$ polynomial.

$\tilde{v}_t$ enters the structural VAR in the same way as the true news shock v_t. So when can we identify it using the method derived in Section 3? When the dimensions of $\tilde{v}_t$ are orthogonal, i.e. when $H_0^{v\prime}\Sigma_{w^\xi}^{-1}$ is diagonal. What does this mean? The fundamental shock ϵ_{t+1}^i to dimension i is associated one-for-one with a noise shock ζ_t^i to that dimension. Noise shocks to different dimensions cannot co-vary.

Does this imply agents cannot receive signals about different fundamentals with correlated noise? No. For example, GDP can still be a noisy signal about both productivity and labor supply. Rather, the condition requires that the noise shocks can be separated into orthogonal noise for each fundamental shock. News-noise equivalence (Chahrour and Jurado, 2018) implies that this condition is equivalent to the structural assumption that news shocks are mutually orthogonal.

B Additional News Horizons

Our baseline method considers 1-period-ahead news. But sometimes shocks are anticipated even further in advance. In this appendix, we describe how to generalize our method to account for news at multiple horizons by including additional forecasts in the VAR.

We define some new notation decomposing structural shocks into their anticipated components over many horizons, similar to Wolf and McKay (2022):

$$\epsilon_t = \nu_{t|t} + \nu_{t|t-1} + \nu_{t|t-2} + ... + \nu_{t|t-k}$$

The shock vector ϵ_t depends on news shocks $\nu_{t|t-j}$ received at each horizon j in the past, up to k total horizons. Mapping to our original one-period-ahead notation, the first two horizons of news were written as $\nu_{t|t} = u_t$ and $\nu_{t|t-1} = v_{t-1}$.

To generalize equation (5), assume that the linear model is:

$$\begin{aligned} x_t &= \sum_{j=1}^{m} B_j x_{t-j} + \sum_{i=0}^{k} A_i \mathbb{E}_t[\epsilon_{t+i}] \\ &= \sum_{j=1}^{m} B_j x_{t-j} + \sum_{i=0}^{k} A_i \left(\sum_{\ell=0}^{k-i} \nu_{t+i|t-\ell} \right) \end{aligned}$$

And suppose you have data on rational forecasts up to horizon k:

$$f_t^i \equiv \mathbb{E}_t[x_{t+i}]$$

Stack the expectations and time series into a single VAR$(m-1)$:

$$\begin{pmatrix} f_t^k \\ \vdots \\ f_t^1 \\ x_t \end{pmatrix} = \sum_{j=1}^{m-1} \mathbf{B}_j \begin{pmatrix} f_{t-j}^k \\ \vdots \\ f_{t-j}^1 \\ x_{t-j} \end{pmatrix} + \mathbf{A} \begin{pmatrix} \nu_{t+k|t} \\ \vdots \\ \nu_{t+1|t} \\ \nu_{t|t} \end{pmatrix} \tag{16}$$

where

$$\mathbf{B}_j \equiv \begin{cases} \begin{pmatrix} B_1 & ... & B_{k-1} & B_k & B_{k+1} \\ \vdots & ... & \vdots & \vdots & \vdots \\ 0 & ... & I & 0 & 0 \\ 0 & ... & 0 & I & 0 \end{pmatrix} & j = 1 \\ \begin{pmatrix} 0 & ... & 0 & 0 & B_{k+j} \\ \vdots & ... & \vdots & \vdots & \vdots \\ 0 & ... & 0 & 0 & 0 \\ 0 & ... & 0 & 0 & 0 \end{pmatrix} & j > 1 \end{cases}$$

which is the generalization of equation (8).

The coefficients in $\mathbf{A}$ are determined by how new shocks affect the forecast updates:

$$\mathbb{E}_t[x_t] - \mathbb{E}_{t-1}[x_t] = \sum_{i=0}^{k} A_i \nu_{t+i|t}$$

$$\mathbb{E}_t[x_{t+1}] - \mathbb{E}_{t-1}[x_{t+1}] = B_1(\mathbb{E}_t[x_t] - \mathbb{E}_{t-1}[x_t]) + \sum_{i=0}^{k-1} A_i \nu_{t+1+i|t}$$

$$\vdots$$

$$\mathbb{E}_t[x_{t+\ell}] - \mathbb{E}_{t-1}[x_{t+\ell}] = \sum_{j=1}^{\ell} B_j(\mathbb{E}_t[x_{t+\ell-j}] - \mathbb{E}_{t-1}[x_{t+\ell-j}]) + \sum_{i=0}^{k-\ell} A_i \nu_{t+\ell+i|t}$$

which implies

$$\mathbf{A}\begin{pmatrix} \nu_{t+k|t} \\ \vdots \\ \nu_{t+1|t} \\ \nu_{t|t} \end{pmatrix} = ...$$

$$\begin{pmatrix} \vdots & ... & \vdots & \vdots \\ A_{k-2} + B_1(A_{k-1} + B_1 A_k) + B_2 A_k & ... & B_1(A_0 + B_1 A_1) + B2A_1 & B_1^2 B_1 A_0 + B_2 A_0 \\ A_{k-1} + B_1 A_k & ... & A_0 + B_1 A_1 & B_1 A_0 \\ A_k & ... & A_1 & A_0 \end{pmatrix}\begin{pmatrix} \nu_{t+k|t} \\ \vdots \\ \nu_{t+1|t} \\ \nu_{t|t} \end{pmatrix}$$

Our baseline method with one-period-ahead news was exactly identified (so long as invertibility conditions were met). With longer horizons, the matrix $\mathbf{A}$ is overidentified, so additional forecast horizons can be useful to help discipline estimation. The matrices $A_0, A_1, ...A_k$ have $(k+1)n^2$ unknowns, and the variance of each news shock $Var(\nu_{t+j|t})$ adds an additional $(k+1)n$ unknowns. The covariance matrix Σ of residuals from the VAR has up to $\frac{((k+1)n)^2}{2} + \frac{(k+1)n}{2}$ independent entries. Finally, variance adding up gives n additional restrictions:

$$I_n = \sum_{j=0}^{k} Var(\nu_{t+j|t})$$

When do the number of independent entries and restrictions exceed the number of unknowns? When the number of news horizons satisfy $k > 1$:

$$\frac{((k+1)n)^2}{2} + \frac{(k+1)n}{2} + n > (k+1)n^2 + (k+1)n$$

$$((k+1)n)^2 + (k+1)n > 2(k+1)n^2 + 2kn$$

$$(k-1)n^2 + (k-1)n > 0$$

which holds with equality for $k = 1$.

C Variance Decomposition Derivation

Restating equation (5)

$$\begin{aligned} x_t &= \sum_{j=1}^{m} B_j x_{t-j} + A\epsilon_t + Cv_t \\ &= \sum_{j=1}^{m} B_j x_{t-j} + Au_t + Av_{t-1} + Cv_t \end{aligned}$$

Letting X_t be the appropriately stacked vector of m lags of x_t. Then:

$$X_t = \hat{B}X_{t-1} + \hat{A}u_t + \hat{A}v_{t-1} + \hat{C}v_t$$

Where $\hat{B}$ concatenates the B_j and adds the lag matrix at the bottom, and $\hat{A}$ and $\hat{C}$ add a bunch of zeros in the extra rows.

Then the $h-$period forecast error is:

$$X_{t+h} - \mathbb{E}_t X_{t+h} = \begin{cases} \hat{A}u_{t+1} + \hat{C}v_{t+1} & h = 1 \\ \sum_{s=1}^{h} \hat{B}^{h-s}\hat{A}u_{t+s} + \sum_{s=1}^{h-1} \hat{B}^{h-s-1}\left(\hat{A} + \hat{B}\hat{C}\right) v_{t+h} + \hat{C}v_{t+h} & h > 1 \end{cases}$$

And the corresponding error variance for the forecast is:

$$MSE_t X_{t+h} = \begin{cases} \hat{A}D_u(\hat{A})' + \hat{C}D_v(\hat{C})' & h = 1 \\ \sum_{s=1}^{h} \hat{B}^{h-s}\hat{A}D_u^2\hat{A}'(\hat{B}')^{h-s} & \\ \qquad + \sum_{s=1}^{h-1} \hat{B}^{h-s-1}\left(\hat{A} + \hat{B}\hat{C}\right) D_v^2 \left(\hat{A} + \hat{B}\hat{C}\right)' (\hat{B}')^{h-s} + \hat{C}D_v^2\hat{C}' & h > 1 \end{cases}$$

And the hs-period-ahead variance due to the jth shock has contemporaneous and news components given by:

$$\text{Surprise} = \sigma_{u,j}^2 \sum_{s=1}^{h} \hat{B}^{h-s}(\hat{A}_j\hat{A}_j')(\hat{B}')^{h-s}$$

$$\text{News} = \begin{cases} \sigma_{v,j}^2(\hat{C}_j\hat{C}_j')' & h = 1 \\ \sigma_{v,j}^2(\hat{C}_j\hat{C}_j') & \\ \quad + \sum_{s=1}^{h} \hat{B}^{h-s-1}\left(\hat{A}_j\hat{A}_j' + \hat{B}(\hat{C}_j\hat{A}_j') + (\hat{A}_j\hat{C}_j')\hat{B}' + \hat{B}(\hat{C}_j\hat{C}_j')\hat{B}'\right)(\hat{B}')^{h-s-1} & h > 1 \end{cases}$$

Where $\hat{A}_j$ etc. are the jth column of the corresponding matrix

D Hidden States

Our identification method requires that the structural VAR in equation (5) is the true data generating process. But what if there are hidden states in the economy that do not appear in the data? In this section, we generalize the method to allow for this possibility.

Again suppose that the state vector x_t follows equation (5), but has some dimensions that are not directly observed. Instead, the data vector y_t is determined by the observation equation

$$y_t = x_t + Gu_t + Gv_{t-1} + Hv_t \tag{17}$$

Without loss of generality, we can normalize the hidden states to obey equations (5) and (17).

Observations are related to forecasts by

$$y_t = f_{t-1} + (A + G)u_t + (C + H)v_t$$

while the forecasts $f_t = \mathbb{E}_t[y_{t+1}]$ are now given by

$$f_t = \mathbb{E}_t[x_{t+1}] + Gv_t = \sum_{j=1}^{m} B_j x_{t+1-j} + (A + G)v_t$$

$$= \sum_{j=1}^{m} B_j(y_{t+1-j} - Gu_{t+1-j} - Gv_{t-j} - Hv_{t+1-j}) + (A + G)v_t$$

$$= B_1(y_t - Gu_t - Gv_{t-1} - Hv_t) + \sum_{j=2}^{m} B_j(y_{t+1-j} - Gu_{t+1-j} - Gv_{t-j} - Hv_{t+1-j}) + (A + G)v_t$$

$$= B_1(f_{t-1} + Au_t - Gv_{t-1} + Cv_t) + \sum_{j=2}^{m} B_j(y_{t+1-j} - Gu_{t+1-j} - Gv_{t-j} - Hv_{t+1-j}) + (A + G)v_t$$

Stack the expectations and time series into a single VARMA($m-1$,m):

$$\begin{pmatrix} f_t \\ y_t \end{pmatrix} = \sum_{j=1}^{m-1} \mathbf{B}_j \begin{pmatrix} f_{t-j} \\ y_{t-j} \end{pmatrix} + \sum_{j=0}^{m} \mathbf{A}_j \begin{pmatrix} v_{t-j} \\ u_{t-j} \end{pmatrix} \tag{18}$$

where (as before)

$$\mathbf{B}_j \equiv \begin{cases} \begin{pmatrix} B_1 & B_2 \\ I & 0 \end{pmatrix} & j = 1 \\ \begin{pmatrix} 0 & B_{j+1} \\ 0 & 0 \end{pmatrix} & j > 1 \end{cases}$$

and

$$\mathbf{A}_j \equiv \begin{cases} \begin{pmatrix} B_1C+A+G & B_1A \\ C+H & A+G \end{pmatrix} & j=0 \\ \begin{pmatrix} -B_jG-B_{j+1}H & -B_{j+1}G \\ 0 & 0 \end{pmatrix} & m>j>0 \\ \begin{pmatrix} -B_mG & 0 \\ 0 & 0 \end{pmatrix} & j=m \end{cases}$$

As in the simple VAR case, the autoregressive terms identify the B_j matrices. But now $\mathbf{A}_0$ has two additional matrices that thwart identification: G and H. Fortunately, the hidden state structure introduces additional MA terms, which allow for possible identification of G and H. We emphasize that with the structure, we only have sufficient conditions for identification – at least as many linearly independent equations as unknowns – but not a constructive proof analogous to Theorem 1. This is because our baseline method admits an analytical solution to the decomposition of the variance matrix Σ, but we have found no such analytical solution in this generalization, so estimation must use a numerical decomposition.

We use $\mathbf{A}_1$ to demonstrate identification, although these matrices are now potentially *overidentified*, so we can use even more lags to improve the statistical power when estimating G and H. The variance matrix of forecast errors is now

$$\Sigma_0 = \mathbf{A}_0 \begin{pmatrix} D_v^2 & 0 \\ 0 & D_u^2 \end{pmatrix} \mathbf{A}_0'$$

but with the MA structure, it is possible to identify the covariance matrix of any two MA components, i.e.:

$$\Sigma_{ij} = \mathbf{A}_i \begin{pmatrix} D_v^2 & 0 \\ 0 & D_u^2 \end{pmatrix} \mathbf{A}_j'$$

To calculate the $\mathbf{A}_i$ matrices, subdivide the matrix $\Sigma_{jj} \equiv \begin{pmatrix} \Sigma_{j,11} & \Sigma_{j,12} \\ \Sigma_{j,21} & \Sigma_{j,22} \end{pmatrix}$ into $n \times n$ blocks. The off-diagonal submatrices satisfy $\Sigma_{j,12} = \Sigma_{j,21}'$, so the remaining submatrices are given by

$$\begin{aligned} \Sigma_{0,11} &= (B_1C+A+G)D_v^2(B_1C+A+G)' + B_1AD_u^2A'B_1' \\ \Sigma_{0,21} &= (C+H)D_v^2(B_1C+A+G)' + (A+G)D_u^2A'B_1' \\ \Sigma_{0,22} &= (C+H)D_v^2(C+H)' + (A+G)D_u^2(A+G)' \end{aligned}$$

which correspond to the three block matrix equations that we used to identify the original VAR (Theorem 1). With two additional matrices to identify, use the covariance between MA terms:

$$\Sigma_{01} = \begin{pmatrix} -(B_1C+A+G)D_v^2(B_1G+B_2H)' - B_1AD_u^2G'B_2' & 0 \\ -(C+H)D_v^2(B_1G+B_2H)' - (A+G)D_u^2G'B_2' & 0 \end{pmatrix}$$

Which, in addition to

$$D_u^2 + D_v^2 = I$$

is as many linear restrictions as unknowns.