



University  
of Glasgow

Adam Smith  
Business School

# WORKING PAPER SERIES

Exploring Monetary Policy Shocks with  
Large-Scale Bayesian VARs

Dimitris Korobilis

Paper No. 2025-09  
May 2025



# Exploring Monetary Policy Shocks with Large-Scale Bayesian VARs

Dimitris Korobilis<sup>1,2,3</sup>

<sup>1</sup>*University of Glasgow, UK*

<sup>2</sup>*Center for Applied Macroeconomics and Commodity Prices (CAMP), Norway*

<sup>3</sup>*Rimini Center for Economic Analysis (RCEA), Italy*

May 9, 2025

## Abstract

I introduce a high-dimensional Bayesian vector autoregressive (BVAR) framework designed to estimate the effects of conventional monetary policy shocks. The model captures structural shocks as latent factors, enabling computationally efficient estimation in high-dimensional settings through a straightforward Gibbs sampler. By incorporating time variation in the effects of monetary policy while maintaining tractability, the methodology offers a flexible and scalable approach to empirical macroeconomic analysis using BVARs, well-suited to handle data irregularities observed in recent times. Applied to the U.S. economy, I identify monetary shocks using a combination of high-frequency surprises and sign restrictions, yielding results that are robust across a wide range of specification choices. The findings indicate that the Federal Reserve's influence on disaggregated consumer prices fluctuated significantly during the 2022–24 high-inflation period, shedding new light on the evolving dynamics of monetary policy transmission.

*Keywords:* Disaggregated consumer prices; Latent factors; High-dimensional Bayesian VAR; Time-varying parameters; Sign restrictions; High frequency data

*JEL Classification:* C11, C32, C55, E31, E52, E58, E66

---

I would like to thank Martin Bruns, Luca Gambetti, Domenico Giannone, Michele Lenza, Nicolò Maffei-Faccioli, Mirela Miescu, Ivan Petrella, Giorgio Primiceri, Barbara Rossi and Lorenza Rossi for their valuable comments and suggestions. I also thank participants at the University of Lancaster's Workshop on Empirical and Theoretical Macroeconomics, the University of East Anglia "2nd Time Series Workshop", the Collegio Carlo Alberto conference on "The Economics of Risk: Econometric Tools and Policy Implications", and seminar participants at Universities of Manchester and Paris Dauphine for their insightful feedback. Any remaining errors are solely my responsibility.

Correspondence: Professor of Econometrics, Adam Smith Business School, University of Glasgow, 2 Discovery Place, Glasgow, G11 6EY, United Kingdom; email: [dikorobilis@googlemail.com](mailto:dikorobilis@googlemail.com).

# 1 Introduction

This paper quantifies the effectiveness of conventional monetary policy shocks during the post-pandemic inflationary period using large-scale vector autoregressions (VARs). While previous studies (see, for example, [Bańbura et al., 2010](#); [Bernanke and Boivin, 2003](#); [Giannone et al., 2015](#)) have established the benefits of identifying monetary policy in data-rich environments, this research develops a flexible yet numerically robust VAR framework to address emerging estimation and identification challenges. First, after a prolonged period of near-zero interest rates, historical data may be less informative for evaluating recent monetary policy effectiveness.<sup>1</sup> Second, the COVID-19 pandemic introduced substantial outliers, disrupting macroeconomic correlations and complicating structural identification.<sup>2</sup> Third, in an increasingly globalized and financially volatile environment, enhancing the VAR with relevant predictors is crucial for statistical fit and for ensuring that structural shocks are well identified. Moving forward, addressing these data challenges requires novel statistical methods and estimation algorithms, capable of handling high-dimensional data, nonlinearities, and other complexities. A major contribution of this study is the development of such methods, allowing the integration of high-dimensional predictors to the application of robust identification techniques for structural shocks.

To tackle these challenges, I develop a high-dimensional and data-analytic approach that integrates economic identification schemes with a Bayesian shrinkage estimation methodology. Building on [Korobilis \(2022\)](#), I extend the VAR to integrate high frequency instruments, fat-tailed errors, time-varying parameters, and stochastic volatility. The approach uses “target”

---

<sup>1</sup>Studies examining conventional monetary policy (MP) shocks, such as [Arias et al. \(2019\)](#); [Read \(2024\)](#); [Schlaak et al. \(2023\)](#), rely on data ending in 2007, that is, just before the period when nominal interest rates effectively reached their zero lower bound. [Aruoba and Drechsel \(2024\)](#) perform their empirical analysis to 2008, arguing that their findings remain relevant for the 2022–2023 inflationary period, provided that the transmission mechanism of MP shocks has not changed. However, they caution that factors such as state dependency, non-linearities, and structural shifts may alter response dynamics and invalidate projections to the recent period.

<sup>2</sup>[Lenza and Primiceri \(2022\)](#) and [Carriero et al. \(2024a\)](#) develop VAR methods to accommodate COVID-19 outliers and stochastic volatility; they do not, however, address structural shock identification or inference in high-dimensional systems.

and “path” factors of high frequency surprise series from [Gürkaynak et al. \(2005\)](#) to look at how the markets react to Federal Reserve Bank’s policy announcements. The target factor describes the effects of surprise changes in the current federal funds rate target, that is, conventional shocks that are of interest in this paper. The path factor captures the future policy path which is used to identify a residual monetary policy shock.<sup>3</sup> I first show that in a small-scale VAR the use of high-frequency surprises as internal instruments for identifying monetary policy is not always robust and in line with theory (see also [Acosta, 2023](#); [Lakdawala, 2019](#), for similar arguments).<sup>4</sup>

Next, I show that a plausible solution to this lack of identification is a hybrid method that merges proxies with zero and sign restrictions to the response of key macroeconomic aggregates – namely measures of output, inflation, short/medium/long interest rates, and money supply. While sign restrictions to these variables align with conventional theoretical predictions of a monetary policy shock, they alone are also insufficient for unique identification and thus need to be combined with high-frequency instruments.<sup>5</sup> An emerging line of research, such as [Acosta \(2023\)](#), [Braun and Brüggemann \(2023\)](#), [Carriero et al. \(2024b\)](#), [Read \(2024\)](#) and [Schlaak et al. \(2023\)](#), demonstrates the advantages of combining established identification schemes, such as sign, zero and narrative restrictions, proxies, and identification via heteroskedasticity and non-Gaussianity. [Prüser \(2024\)](#) and [Chan et al. \(2022\)](#) similarly extend the VAR methodology of [Korobilis \(2022\)](#) by adapting it to identification by non-Gaussianity and heteroskedasticity, respectively. However, these

---

<sup>3</sup>There are varying views about the number and effects of future policy path and unconventional shocks, in general. [Nakamura and Steinsson \(2018\)](#) argue that the Fed has private information (what they call the “Fed information effect”), while [Bauer and Swanson \(2023\)](#) argue that both the Fed and professional forecasters respond to incoming economic news, which naturally generates a positive correlation between monetary policy surprises and forecast revisions, without requiring the Fed to have private information about the economy. [Swanson \(2021\)](#) extends [Gürkaynak et al. \(2005\)](#) to include a third, quantitative easing factor, and [Jarociński \(2024\)](#) identifies up to three unconventional monetary policy shocks.

<sup>4</sup>Several other studies have raised concerns about the use of high-frequency instruments (HFIs) in identifying monetary policy shocks, particularly regarding issues of exogeneity, temporal aggregation bias, and event window limitations; see [Bauer and Swanson \(2023\)](#), [Paccagnini and Parla \(2021\)](#) and [Casini and McCloskey \(2024\)](#), among others.

<sup>5</sup>[Braun and Brüggemann \(2023\)](#) present similar arguments in the context of VARs that combine external instruments with sign restrictions. While their algorithm is very useful, it is not computationally feasible for large VARs, particularly when allowing for time-varying shock impacts.

identification approaches are statistical in nature, and may not always ensure adequate identification of conventional monetary policy shocks. In contrast, I provide comprehensive evidence that combining high-frequency instruments with sign restrictions remains economically robust across various data and modeling assumptions, including lag length, stationarity transformations of variables, assumptions about residual shocks, the measurement of short-term interest rates, and different assumptions regarding the exogeneity of instruments.

After establishing that the hybrid identification scheme performs effectively in a small-scale VAR, the analysis examines the impact of post-pandemic conventional monetary policy on the broader economy, with a particular focus on disaggregated components of consumer inflation. The VAR model is expanded to include key financial indicators, uncertainty measures, expectations, and other predictors, ensuring that the residuals effectively capture structural monetary policy shocks. These predictors and inflation components are permitted to exhibit time-varying responses to monetary policy shocks, accommodating changes in the transmission mechanism over time, with idiosyncratic shocks modeled to incorporate stochastic volatility dynamics. Employing Bayesian shrinkage and variable selection techniques allows for the semi-automatic determination of the optimal degree of time variation within the system, adeptly adjusting to abrupt changes, such as those introduced by post-COVID-19 outliers. This approach differs from [Mumtaz and Petrova \(2023\)](#), where all parameters in their proxy VAR are time-varying. In contrast, our methodology maintains constant weights for high-frequency instruments and key macroeconomic aggregates (those subjected to sign restrictions) when loading onto structural shocks. This strategy ensures robust identification through these variables while permitting other variables of interest to have unrestricted, time-varying responses to monetary policy shocks. Consequently, the large-scale system integrates multiple identification layers, balancing the imposition of economic theory with the flexibility for data to adapt to abrupt regime shifts experienced since the global financial crisis.

The proposed VAR methods are applied to monthly macroeconomic data for the US for the period 1995-2024. Among the large number of variables, the variables of interest are disaggregate price indices and their response to conventional monetary policy shock during the post-pandemic inflationary period. Based on the definitions of the Bureau for Labor Statistics, I use eight major categories of the US consumer price index (CPI) such as food, apparel, and education, in order to pin down the heterogeneous effect of policy shocks. I find that the effects of monetary tightening vary substantially across CPI components, with core goods exhibiting rapid and pronounced disinflationary responses, while services (particularly housing) adjust more sluggishly and with long lags. Importantly, these sector-specific responses have evolved over time: in the post-pandemic period, the impact of conventional policy shocks on goods prices has intensified, likely reflecting the unwinding of pandemic-era demand and supply imbalances. Despite this heterogeneity, virtually all components eventually move in the direction of lower inflation following a contractionary shock, confirming that monetary policy remains a broadly effective tool for influencing inflation dynamics, even in a high-inflation environment.

The paper is organized as follows. The next Section presents the Bayesian VAR methodology and builds the various components that will be relevant for large-scale estimation and structural inference. Section 3 details the variables, sources and transformations used in this high-dimensional setting. In Section 4 I provide the empirical results, which are also complemented by an online supplement. Section 5 concludes the paper.

## 2 Econometric methodology

The starting point is the  $p$ -lag vector autoregression of the form

$$\mathbf{y}_t = \phi_0 + \sum_{j=1}^p \Phi_j \mathbf{y}_{t-j} + \boldsymbol{\varepsilon}_t. \quad (1)$$

This can be written more compactly as

$$\mathbf{y}_t = \mathbf{\Phi} \mathbf{x}_t + \boldsymbol{\varepsilon}_t, \quad (2)$$

where  $\mathbf{y}_t$  is an  $(n \times 1)$  vector of observed macro variables,  $\mathbf{x}_t = (1, \mathbf{y}'_{t-1}, \dots, \mathbf{y}'_{t-p})'$  a  $(k \times 1)$  vector of lags and the intercept,  $k = np + 1$ ,  $\mathbf{\Phi} = [\phi_0, \mathbf{\Phi}_0, \dots, \mathbf{\Phi}_p]$  is an  $(n \times k)$  matrix of coefficients, and  $\boldsymbol{\varepsilon}_t$  a  $(n \times 1)$  vector of disturbances distributed as  $N(\mathbf{0}_{n \times 1}, \mathbf{\Omega})$  with  $\mathbf{\Omega}$  an  $n \times n$  symmetric and positive semi-definite covariance matrix. The structural VAR (SVAR) form is derived by left-multiplying (1) with  $\mathbf{A}$  yielding

$$\mathbf{A} \mathbf{y}_t = \mathbf{B} \mathbf{x}_t + \mathbf{u}_t, \quad (3)$$

where  $\mathbf{A}^{-1} \mathbf{A}^{-1'} = \mathbf{\Omega}$  is a decomposition of the reduced-form covariance matrix,  $\mathbf{B} = \mathbf{A} \mathbf{\Phi}$  are the SVAR coefficients and  $\mathbf{u}_t \sim N(\mathbf{0}, \mathbf{I})$  are structural shocks where it holds that  $\boldsymbol{\varepsilon}_t = \mathbf{A}^{-1} \mathbf{u}_t$ .

The VAR in (1) can be estimated with standard tools, providing estimates of  $\hat{\mathbf{\Phi}}$  and  $\hat{\mathbf{\Omega}}$ . However, recovering the matrix  $\mathbf{A}$  of structural contemporaneous restrictions is not feasible without further assumptions.  $\mathbf{A}$  has  $n^2$  elements, while  $\hat{\mathbf{\Omega}}$  has  $n(n+1)/2$  elements estimated from data. Identification restrictions, usually stemming from economic theory, stylized facts, or intuition, fill the gap in providing information on the remaining  $n(n-1)/2$  coefficients. I focus on two popular approaches to identification, namely sign restrictions and high frequency instruments. The former helps identify elements of  $\mathbf{A}$  that satisfy expected signs; for example, if the response of variable  $i$  to a shock in  $j$  is positive,  $i, j = 1, \dots, n$ , then  $A_{ij} > 0$ . The latter approach uses  $m$  proxies or instruments  $m_t$  that are correlated with  $m$  shocks of interest and uncorrelated with other shocks. In this paper,  $m_t$  are high-frequency surprises in interest rate markets, so the terms proxies, instruments, and surprises may be used interchangeably.



## 2.1 A large monetary policy Bayesian VAR

In the presence of  $m$  high frequency surprises  $\mathbf{m}_t$  and large- $n$  macro variables  $\mathbf{y}_t$ , the VAR of (2) has to be augmented with the surprise data. Using the result above that  $\boldsymbol{\varepsilon}_t = \mathbf{A}^{-1}\mathbf{u}_t$ , I define the following generic representation of this augmented VAR (see also [Braun and Brüggemann, 2023](#))

$$\begin{bmatrix} \mathbf{m}_t \\ \mathbf{y}_t \end{bmatrix} = \begin{bmatrix} \boldsymbol{\Phi}^{mm} & \boldsymbol{\Phi}^{my} \\ \boldsymbol{\Phi}^{ym} & \boldsymbol{\Phi}^{yy} \end{bmatrix} \begin{bmatrix} \mathbf{m}_{t-1} \\ \mathbf{y}_{t-1} \end{bmatrix} + \begin{bmatrix} \boldsymbol{\Gamma} \\ \mathbf{A}^{-1} \end{bmatrix} \mathbf{u}_t + \begin{bmatrix} \mathbf{W}^{1/2} \\ \mathbf{0} \end{bmatrix} \boldsymbol{\eta}_t, \quad (4)$$

where  $\boldsymbol{\Phi}^{ij}$  are the autoregressive coefficients of lags of variables  $j$  in equations of variables  $i$ , with  $i, j = \mathbf{m}, \mathbf{y}$ ,  $\mathbf{u}_t \sim N(\mathbf{0}, \mathbf{I})$  still denotes the structural VAR shock, the term  $\boldsymbol{\eta}_t \sim N(\mathbf{0}, \mathbf{I})$  is an idiosyncratic disturbance with  $\boldsymbol{\eta}_t \perp \mathbf{u}_t$ ,  $\boldsymbol{\Gamma}$  is a matrix of coefficients and  $\mathbf{W}$  a diagonal matrix of variances.  $\boldsymbol{\Gamma}$  measures the strength of the instrument's correlation with the structural shock, while  $\mathbf{W}$  accounts for the possibility that the instrument is imperfect, allowing for measurement error or partial relevance. For the sake of notational simplicity the intercept term is excluded in the specification above, and I focus on the one-lag specification without loss of generality. [Mertens and Ravn \(2013\)](#) and [Jarociński and Karadi \(2020\)](#) assume that  $\boldsymbol{\Phi}^{mm} = \boldsymbol{\Phi}^{my} = \mathbf{0}$ . [Bauer and Swanson \(2023\)](#) argue that interest rate surprises might be driven by macroeconomic news which is equivalent to assuming  $\boldsymbol{\Phi}^{my} \neq \mathbf{0}$ , as long as  $\mathbf{y}$  also includes measures of macro news. Another empirically relevant question is also whether macroeconomic variables are dynamically independent of high frequency surprises, implying  $\boldsymbol{\Phi}^{ym} = \mathbf{0}$ .

In higher dimensions, when  $n$  is sufficiently large, it is expected that the VAR system does not necessarily have  $n$  structural or primitive shocks. Instead, a smaller number of common, economically meaningful, driving forces are responsible for causing disruptions in the system. For that reason I follow previous work in [Korobilis \(2022\)](#) and replace the identity  $\boldsymbol{\varepsilon}_t = \mathbf{A}^{-1}\mathbf{u}_t$  implied by equation (3) with the factor-like decomposition

$$\boldsymbol{\varepsilon}_t = \boldsymbol{\Lambda}\mathbf{f}_t + \mathbf{v}_t, \quad (5)$$



where  $\mathbf{f}_t \sim N(\mathbf{0}, \mathbf{I})$  is an  $r \times 1$  vector of factors with  $r \ll n$ ,  $\mathbf{\Lambda}$  an  $n \times r$  matrix of loadings, and  $\mathbf{v}_t \sim N(\mathbf{0}, \mathbf{\Sigma})$  is a vector of idiosyncratic disturbances where  $\mathbf{\Sigma}$  is a diagonal matrix of variances. Under this factor decomposition of the reduced-form disturbances, the VAR augmented with instruments of equation (4) can now be written as

$$\begin{bmatrix} \mathbf{m}_t \\ \mathbf{y}_t \end{bmatrix} = \begin{bmatrix} \Phi^{mm} & \Phi^{my} \\ \Phi^{ym} & \Phi^{yy} \end{bmatrix} \begin{bmatrix} \mathbf{m}_{t-1} \\ \mathbf{y}_{t-1} \end{bmatrix} + \begin{bmatrix} \mathbf{\Gamma} \\ \mathbf{\Lambda} \end{bmatrix} \mathbf{f}_t + \begin{bmatrix} \mathbf{W}^{1/2} \\ \mathbf{\Sigma}^{1/2} \end{bmatrix} \boldsymbol{\eta}_t, \quad (6)$$

where in the notation above I have used the fact that  $\mathbf{v}_t \equiv \mathbf{\Sigma}^{1/2} \boldsymbol{\eta}_t$ .

What is the economic interpretation of the model above? A smaller number of  $r$  structural shocks  $\mathbf{f}_t$  exist in the VAR. In order to see this, define  $\mathbf{y}_t^* = [\mathbf{m}_t', \mathbf{y}_t']'$  and  $\Phi = \begin{bmatrix} \Phi^{mm} & \Phi^{my} \\ \Phi^{ym} & \Phi^{yy} \end{bmatrix}$ ,

$$\mathbf{\Gamma}^* = \begin{bmatrix} \mathbf{\Gamma} \\ \mathbf{\Lambda} \end{bmatrix} \text{ and } \mathbf{W}^* = \begin{bmatrix} \mathbf{W} \\ \mathbf{\Sigma} \end{bmatrix}, \text{ and solve for the SVAR form as follows}$$

$$\mathbf{y}_t^* = \Phi \mathbf{y}_{t-1}^* + \mathbf{\Gamma}^* \mathbf{f}_t + \mathbf{W}^{*1/2} \boldsymbol{\eta}_t, \Rightarrow \quad (7)$$

$$\mathbf{A}^* \mathbf{y}_t^* = \mathbf{B} \mathbf{y}_{t-1}^* + \mathbf{f}_t + \mathbf{A}^* \mathbf{W}^{*1/2} \boldsymbol{\eta}_t, \quad (8)$$

$$\mathbf{A}^* \mathbf{y}_t^* \approx \mathbf{B} \mathbf{y}_{t-1}^* + \mathbf{f}_t, \quad (9)$$

where  $\mathbf{A}^* = (\mathbf{\Gamma}^* \mathbf{\Gamma}^*)^{-1} \mathbf{\Gamma}^*$  is an  $r \times (n + m)$  matrix,  $\mathbf{B}^* = \mathbf{A}^* \Phi$ . Compared to the standard SVAR form in equation (3) with contemporaneous relationship matrix  $\mathbf{A}$ , equation (9) is a reduced-rank SVAR with  $r$  structural relationships among the  $n$  macro variables (plus the  $m$  instruments). This SVAR formulation results from the fact that  $\mathbf{A}^* \mathbf{W}^{*1/2} \boldsymbol{\eta}_t \rightarrow 0$  as  $n \rightarrow \infty$  (Korobilis, 2022). Therefore, assuming the effect of the disturbance term  $\boldsymbol{\eta}_t$  vanishes asymptotically in the SVAR representation,  $\mathbf{f}_t$  is approximately equivalent to the vector of structural disturbances  $\mathbf{u}_t$ . However, in the reduced-form formulation the effect of  $\boldsymbol{\eta}_t$  remains important and it has to be estimated alongside other parameter. This term has a natural interpretation too. From a statistical perspective, idiosyncratic disturbances

$\boldsymbol{\eta}_t$  are considered to arise due to measurement error, outliers and other irregularities in each individual time series. From an economic point of view, the idiosyncratic disturbance term corresponds to expectations errors, heterogeneous information sets, myopia and other forms of irrational behavior (Gorodnichenko, 2005).

The result above suggest that impulse response functions (IRFs) are able to reflect the propagation of economically interpretable structural shocks, even in the presence of heterogeneous and noisy time series. IRFs trace the dynamic effect of a structural shock  $f_{jt}$  on the vector of observables  $\mathbf{y}_t^* = [\mathbf{m}_t', \mathbf{y}_t']'$  over time. These effects can be obtained by considering the vector moving average (VMA) representation of the system in (6), which takes the form:

$$\mathbf{y}_t^* = \boldsymbol{\mu} + \sum_{h=0}^{\infty} \boldsymbol{\Psi}_h (\boldsymbol{\Gamma}^* \mathbf{f}_{t-h} + \mathbf{W}^{*1/2} \boldsymbol{\eta}_{t-h}), \quad (10)$$

where  $\boldsymbol{\mu}$  is a constant mean vector,  $\boldsymbol{\Psi}_h$  are the VMA coefficients implied by the VAR companion form, and  $\boldsymbol{\Gamma}^*$  stacks the proxy coefficients  $\boldsymbol{\Gamma}$  and the macroeconomic loadings  $\boldsymbol{\Lambda}$ .

The impulse response at horizon  $h$  to the  $j$ -th structural factor  $f_{jt}$  is:

$$\text{IRF}_h^{(j)} = \frac{\partial \mathbf{y}_{t+h}^*}{\partial f_{jt}} = \boldsymbol{\Psi}_h \boldsymbol{\Gamma}_{\cdot j}^*, \quad (11)$$

where  $\boldsymbol{\Gamma}_{\cdot j}^*$  is the  $j$ -th column of  $\boldsymbol{\Gamma}^*$ . In particular, the impact IRF (i.e., the contemporaneous effect) is given by:

$$\text{IRF}_0^{(j)} = \boldsymbol{\Psi}_0 \boldsymbol{\Gamma}_{\cdot j}^* = \boldsymbol{\Gamma}_{\cdot j}^*. \quad (12)$$

This result is immediate from the identity  $\boldsymbol{\Psi}_0 = \mathbf{I}$  and demonstrates that the loading vector  $\boldsymbol{\Gamma}_{\cdot j}^*$  defines the contemporaneous response of each observable variable to structural factor  $f_{jt}$ . Therefore, any parametric restrictions imposed on the sign or magnitude of elements of  $\boldsymbol{\Gamma}$  and  $\boldsymbol{\Lambda}$  (i.e.,  $\boldsymbol{\Gamma}^*$ ) directly translate into sign and shape restrictions on the impact responses of the observed series. This insight underpins the identification strategy employed in this paper. Specifically, the combination of  $\boldsymbol{\Gamma}$  and  $\boldsymbol{\Lambda}$  determines the direction and magnitude of the contemporaneous responses to the structural factors  $\mathbf{f}_t$ . In the context of an instrument-

augmented VAR, the matrix  $\mathbf{\Gamma}$  governs the impact of the factors on the instrument block  $\mathbf{m}_t$ , while  $\mathbf{\Lambda}$  governs the impact on the macroeconomic variables  $\mathbf{y}_t$ .

In this paper, I impose that  $\mathbf{\Gamma}$  is a diagonal matrix with positive entries, ensuring that each instrument loads positively on its corresponding factor while maintaining orthogonality across instruments. This choice is natural given the economic interpretation of the instruments and facilitates factor identification without requiring an arbitrary normalization such as  $\mathbf{\Gamma} = \mathbf{I}_r$ . While this restriction uniquely identifies a rotation of the factors, it does not guarantee that the resulting impulse responses of the macroeconomic variables are economically meaningful. Therefore, I impose additional economically motivated sign restrictions on selected rows of  $\mathbf{\Lambda}$ , corresponding to the expected effects of monetary policy shocks on key aggregates such as output and inflation. This two-layer identification strategy ensures that the factors are both statistically identified through the instrument block and economically interpretable through their impact on macroeconomic variables. These sign restrictions are trivial to impose. [Korobilis \(2022\)](#) shows that such restrictions translate into truncated normal posterior distributions that are trivial to sample from: identification is embedded directly into the conditional posteriors, ensuring that every draw automatically satisfies the imposed restrictions. This one-step identification procedure provides major computational advantages, allowing the method to scale efficiently to large VAR systems. An outline of the posterior sampler is given in [subsection 2.3](#) and full technical details of the sampler are provided in the online supplement.

## 2.2 Adding time-variation, stochastic volatility and fat-tailed errors

Time-varying parameter models have gained lots of prominence in macroeconomics, especially since the global financial crisis. They allow to accommodate structural breaks in the data in a flexible way. Compared to Markov switching and other regime-switching methods, time-varying parameters are able to unveil the full transmission mechanism of shocks. For that

reason, I propose to extend the VAR of the previous Section into the following time-varying specification

$$\begin{bmatrix} \mathbf{m}_t \\ \mathbf{y}_t \end{bmatrix} = \begin{bmatrix} \Phi^{mm} & \Phi^{my} \\ \Phi^{ym} & \Phi^{yy} \end{bmatrix} \begin{bmatrix} \mathbf{m}_{t-1} \\ \mathbf{y}_{t-1} \end{bmatrix} + \begin{bmatrix} \Gamma \\ \Lambda_t \end{bmatrix} \mathbf{f}_t + \begin{bmatrix} \mathbf{W}^{1/2} \\ \Sigma_t^{1/2} \end{bmatrix} \boldsymbol{\eta}_t. \quad (13)$$

where the vectorized forms of the matrix  $\Lambda_t$  as well as each diagonal element  $\log(\Sigma_{ii,t})$ ,  $i = 1, \dots, n$  follow random walk specifications of the form

$$vec(\Lambda_t) = vec(\Lambda_{t-1}) + \boldsymbol{\xi}_t, \quad (14)$$

$$\log(\Sigma_{ii,t}) = \log(\Sigma_{ii,t-1}) + \delta_{i,t}, \quad \forall i = 1, \dots, n, \quad (15)$$

where  $\boldsymbol{\xi}_t \sim N(\mathbf{0}, \mathbf{Q})$  with  $\mathbf{Q}$  an  $n \times n$  covariance matrix, and  $\delta_{i,t} \sim N(0, \omega_i^2)$  with  $\omega_i^2$  a scalar variance parameter, for  $i = 1, \dots, n$ .

This form of time-variation has been used consistently since the early macroeconomic work of the 1970s and it implies the prior belief that each time varying parameter at time  $t$  is centered at its time  $t - 1$  value. It is important to emphasize that the extent to which the responses of macroeconomic variables vary over time is not determined arbitrarily, but is governed by the prior placed on the innovation variance matrix  $\mathbf{Q}$ . In particular, I adopt a hierarchical horseshoe prior for each diagonal element of  $\mathbf{Q}$ , which ensures adaptive shrinkage (Carvalho et al., 2010). This prior is both automatic and tuning-free, and requires no manual calibration of shrinkage parameters. Each element of  $\mathbf{Q}$  is decomposed into global and local variance components, which jointly control the degree of time variation in the corresponding loading coefficient. The prior shrinks small innovations in  $\Lambda_t$  aggressively toward zero, effectively removing time variation where not supported by the data. At the same time, it allows large innovations to persist, thereby capturing structural breaks. This regularization mechanism ensures that the impulse responses of disaggregated prices and other unrestricted variables are shaped primarily by the data, rather than by arbitrary

assumptions about the nature or smoothness of time variation.

Time-varying parameters and stochastic volatility are important for capturing structural breaks. However, they may not effectively handle the large outliers associated with the COVID-19 period. For this reason, I decompose the idiosyncratic error  $\boldsymbol{\eta}_t$  into two components: one corresponding to the proxies,  $\boldsymbol{\eta}_t^m$ , and another corresponding to the macroeconomic variables,  $\boldsymbol{\eta}_t^y$ . I assume:

$$\boldsymbol{\eta}_t = \begin{bmatrix} \boldsymbol{\eta}_t^m \\ \boldsymbol{\eta}_t^y \end{bmatrix} \sim \begin{pmatrix} N(\mathbf{0}_{m \times 1}, \mathbf{I}_m) \\ t(\boldsymbol{\nu}, \mathbf{0}_{n \times 1}, \mathbf{I}_n) \end{pmatrix} \quad (16)$$

where  $t(\boldsymbol{\nu}, \mathbf{0}, \mathbf{I})$  denotes the standardized Student-t distribution with  $\boldsymbol{\nu}$  an  $n \times 1$  vector indicating the degrees of freedom for each macro variable. In this specification,  $\boldsymbol{\nu}$  is a parameter estimated from the data. Smaller values of  $\boldsymbol{\nu}$  allow for fatter tails and lower kurtosis, thus accommodating extreme outliers in the data.

Two key aspects of the specification above should be emphasized. First, the autoregressive parameters are not time-varying. Given the interest in high-dimensional VARs, the matrix  $\boldsymbol{\Phi} = \begin{bmatrix} \boldsymbol{\Phi}^{mm} & \boldsymbol{\Phi}^{my} \\ \boldsymbol{\Phi}^{ym} & \boldsymbol{\Phi}^{yy} \end{bmatrix}$  may contain tens of thousands of elements. Consequently, converting such a large number of parameters into time-varying values would lead to significant estimation error and bias.<sup>6</sup> Moreover, ensuring that each draw of the autoregressive coefficients corresponds to a stationary VAR would be computationally impractical. Empirical evidence from the macroeconomics literature provides an additional reason to avoid time variation in  $\boldsymbol{\Phi}$ . For example, [Sims and Zha \(2006\)](#) find minimal evidence of variation in  $\boldsymbol{\Phi}$  over time, despite finding clear evidence of time-varying volatilities in a VAR. Even when  $\boldsymbol{\Phi}$  is time-invariant, it can contain thousands of elements in high-dimensional settings. To regularize the estimation of this matrix, I follow [Korobilis \(2022\)](#) and impose a horseshoe prior.

---

<sup>6</sup>The limited degrees of freedom in the data and the infeasibility of running the MCMC chain for a sufficient number of iterations make accurate estimation difficult in very high-dimensional settings.

The second important aspect of the proposed time-varying extension is that the proxy equation parameters  $\mathbf{\Gamma}$  and  $\mathbf{W}$  should also remain constant, and the idiosyncratic term  $\boldsymbol{\eta}_t^m$  should remain Gaussian. In the context of an external instruments VAR model, [Mumtaz and Petrova \(2023\)](#) allow these parameters to be time-varying. However, in the proposed specification, the internal instruments load on  $\mathbf{f}_t$ , a vector of latent time-varying parameters that must be estimated from the data. If  $\mathbf{\Gamma}$  and  $\mathbf{W}$  were also time-varying, they would absorb much of the informative content of the proxies, making the latent factors “flat” and unidentifiable. The same issue would arise if  $\boldsymbol{\eta}_t^m$  followed a Student-t distribution: it would absorb many of the spikes in the surprise proxies, rendering the informational content of the factors economically irrelevant. In this model, a lack of factor identification translates into a failure to identify structural shocks, including the monetary policy shock of interest. Therefore, it is more appropriate to treat the relevance of the proxies as constant and the idiosyncratic shocks as Gaussian, while allowing the effects of identified shocks on macroeconomic variables to vary over time. Maintaining constant relevance of the proxies ensures that the model can more accurately and robustly identify the time-varying effects of monetary policy shocks on the macroeconomic variables of interest.

## 2.3 Bayesian Estimation

The details of the Bayesian estimation algorithm are provided in the Online Supplement. In this section, I briefly outline the main steps and emphasize the computational tractability of the proposed framework.

Estimation proceeds via a Gibbs sampler that iteratively samples from the complete conditional posterior distributions of each block of parameters. Despite incorporating several important extensions relative to [Korobilis \(2022\)](#) — namely, the inclusion of high-frequency instruments, time-varying loadings and volatilities, and fat-tailed idiosyncratic shocks — all conditional posterior distributions remain available in closed form. This feature ensures that estimation remains computationally efficient even in very high-dimensional settings.

The key steps of the Gibbs sampler are as follows:

1. **VAR coefficients  $\Phi$ :** The autoregressive parameters are sampled equation-by-equation under a hierarchical Horseshoe prior that induces shrinkage and regularizes estimation in large VARs.
2. **Factor loadings  $\Gamma$  and  $\Lambda$ :** The factor loading matrices for the instruments and macro variables are sampled, respecting any imposed sign restrictions to facilitate structural interpretation. If certain loadings are time-varying, allowing exploration of time-varying IRFs, then a Horseshoe prior is used to shrink the expanded parameter space.
3. **Latent factors  $f_t$ :** The structural shocks (latent factors) are drawn conditional on the current parameters and data, ensuring orthogonality and unit variance normalization.
4. **Idiosyncratic variances  $W$  and  $\Sigma$ :** The variances of the idiosyncratic disturbances are sampled using standard inverse-gamma posteriors.
5. **Time-varying loadings  $\Lambda_t$  and volatilities  $\Sigma_t$ :** When time variation is allowed, the loadings and log-volatilities are modeled as random walks, and their entire trajectories are sampled using a stacked regression representation that avoids sequential Kalman filtering (Korobilis, 2021).
6. **Degrees of freedom  $v$ :** When allowing for leptokurtic errors, the degrees of freedom parameters governing the Student-t distributions of idiosyncratic shocks can be sampled using a standard Metropolis-within-Gibbs step (Geweke, 1993).

By carefully designing the model structure and prior specification, the estimation algorithm avoids the need for computationally intensive steps. This makes the proposed framework highly scalable and well-suited for applications involving a large number of macroeconomic variables and predictors. In the empirical application, which involves a large VAR model with 30 variables (including the two instruments) and  $p = 2$  lags, it takes approximately two hours



to generate 100,000 posterior draws using a desktop computer equipped with an Intel Core i9 12900K processor and 32 GB of RAM.

### 3 Data

The US macroeconomic dataset used in this paper spans January 1995 to May 2024 at a monthly frequency. Data transformations are selected to induce stationarity and align with standard VAR practice: levels for rate variables and indexes (interest rates, uncertainty, sentiment, financial conditions); first log differences for equity indices, exchange rates, and oil prices; and annual growth rates for output, inflation, and consumer spending. The use of annual growth rates reflects the smoother dynamics of these series, which helps produce clearer IRFs given that the data are not pre-filtered for outliers.

The data fall into four broad categories: high-frequency monetary policy surprises, core aggregate macroeconomic variables, disaggregated CPI components, and predictor variables:

**High-Frequency Monetary Policy Surprises** I follow [Gürkaynak et al. \(2005\)](#) in constructing two orthogonal factors from intradaily Treasury futures around FOMC announcements.

- Target factor: captures unexpected changes in the current federal funds rate target.
- Path factor: measures surprises in the anticipated future path of policy (e.g. forward guidance).

Monthly aggregates of these high-frequency surprises are drawn from [Acosta et al. \(2024\)](#), who constructs them by measuring 30-minute windows around scheduled FOMC announcements. In months when no announcements are available, the values of the surprises are set to zero, denoting the absence of information. [Figure 1](#) plots these two series.

In the online supplement, I further assess robustness by re-estimating the VAR using the directly observed surprise series of [Jarociński \(2024\)](#), imposing the exact sign restriction




Figure 1: Plots of target and path factors for the monthly sample 1995M1 to 2024M5.

identification scheme of [Jarociński and Karadi \(2020\)](#); these data are publicly available on the author’s website.

**Core Aggregate Macroeconomic Variables** Our core VAR includes seven U.S. aggregates: real gross domestic product (GDP), personal consumption expenditures (PCE) inflation, the federal funds rate, 1- and 10-year Treasury yields, real M2 money and the Standard and Poor’s 500 (S&P500) stock market index. For robustness, in the online supplement I consider several other monthly proxies of output and price inflation, such as total industrial production (IP) and core consumer price index (CPI).

**Disaggregated CPI Components** To study heterogeneous effects across consumption categories, I incorporate eight CPI subindices: Apparel; Education; Food; Other Goods; Housing; Medical Care; Recreation; and Transportation. These series are transformed into annual percentage changes to ensure comparability with aggregate inflation measures.

**Other macro and predictor variables** In addition to the seven core macro aggregates, the large VAR is augmented with a broad set of other macro and predictor series.

[Table 1](#) details of all variables that are used in the remainder of this section. When a

variable is originally available at a higher frequency (daily or weekly), monthly values are constructed by averaging all high frequency values over the month. All variables are downloaded from their sources seasonally adjusted, where this is relevant. The data were collected on February 12, 2025. Column Tcode shows short codes for applying stationarity transformations to the time series data, prior to estimating the VAR. The codes are as follows, 1: level, 2: first differences, 4: log-level, 5: first log differences (month-on-month growth rate), 7: annual differences (annual growth rate).

Table 1: U.S. macro/financial variables, sources and definition

No	Mnemonic	Description	Tcode	Source
<u>PROXIES</u>				
1	Target	Target factor of <a href="#">Gürkaynak et al. (2005)</a>	1	<a href="#">Acosta et al. (2024)</a> <sup>1</sup>
2	Path	Path factor of <a href="#">Gürkaynak et al. (2005)</a>	1	<a href="#">Acosta et al. (2024)</a> <sup>1</sup>
<u>CORE MACRO VARIABLES WHOSE RESPONSE IS SIGN-RESTRICTED</u>				
3	RGDP	Monthly real GDP	7	S&P <sup>2</sup>
4	PCE	Personal consumption expenditure deflator	7	FRED <sup>3</sup>
5	SHADOW	Wu-Xia Shadow Federal Funds Rate	1	FRB <sup>13</sup>
6	GS1	Market Yield on U.S. Treasury Securities at 1-Year	1	FRED <sup>3</sup>
7	GS10	Market Yield on U.S. Treasury Securities at 10-Year	1	FRED <sup>3</sup>
8	M2REAL	Real M2 Money Stock	5	FRED <sup>3</sup>
9	SP500	S&P 500 price index	5	investing.com <sup>4</sup>
<u>DISAGGREGATE CPI VARIABLES USED IN THE LARGE-SCALE VAR</u>				
10	CPIA	CPI: Apparel	7	BLS <sup>5</sup>
11	CPIE	CPI: Education	7	BLS <sup>5</sup>
12	CPIF	CPI: Food	7	BLS <sup>5</sup>
13	CPIO	CPI: Other goods	7	BLS <sup>5</sup>
14	CPIH	CPI: Housing	7	BLS <sup>5</sup>
15	CPIM	CPI: Medical care	7	BLS <sup>5</sup>
16	CPIR	CPI: Recreation	7	BLS <sup>5</sup>
17	CPIT	CPI: Transportation	7	BLS <sup>5</sup>
<u>ADDITIONAL PREDICTOR VARIABLES USED IN THE LARGE-SCALE VAR</u>				
18	UNRATE	Unemployment Rate	1	FRED <sup>3</sup>
19	HOUST	New Privately-Owned Housing Units Started	4	FRED <sup>3</sup>
20	RCONS	Real Personal Consumption Expenditures	5	FRED <sup>3</sup>
21	TWEXAFEGSMTHx	Nominal Major Currencies U.S. Dollar Index	5	FRED <sup>3</sup>
22	MCOILBRENTU	Crude Oil Prices: Brent - Europe	5	FRED <sup>3</sup>
23	USEPU	Global economic policy uncertainty	1	EPU website <sup>6</sup>

24	GPR	Geopolitical risk index	1	<a href="#">Caldara and Iacoviello (2022)</a> <sup>7</sup>
25	EBP	Excess bond premium	1	FRB <sup>8</sup>
26	JLNF12	Financial uncertainty	1	<a href="#">Jurado et al. (2015)</a> <sup>9</sup>
27	UMSCENT	University of Michigan Consumer Sentiment	1	FRED <sup>3</sup>
28	VIXCLS	CBOE Volatility Index: VIX	1	FRED <sup>3</sup>
29	MPU	Monetary Policy Uncertainty	1	FRB <sup>10</sup>
30	GFC	Global Financial Cycle	1	<a href="#">Miranda-Agrippino and Rey (2020)</a> <sup>11</sup>

#### ALTERNATIVE MEASURES

1*/2*	MP1	Fed funds futures (surprises)	1	<a href="#">Jarociński (2024)</a> <sup>12</sup>
1*/2*	TFUT10	10-year U.S. Treasury Note futures (surprises)	1	<a href="#">Jarociński (2024)</a> <sup>12</sup>
1*/2*	SP500FUT	S&P futures (surprises)	1	<a href="#">Jarociński (2024)</a> <sup>12</sup>
3*	BBKM GDP	Brave-Butters-Kelley Real Gross Domestic Product	1	FRED <sup>3</sup>
3*	INDPRO	INDPRO	7	FRED <sup>3</sup>
4*	CPILFESL	CPI: All Items Less Food and Energy	7	FRED <sup>3</sup>
5*	FFR	Federal Funds Effective Rate	1	FRED <sup>3</sup>

<sup>1</sup> Author's webpage <https://www.acostamiguel.com/replication/MPshocksAcosta.xlsx>

<sup>2</sup> S&P Market Intelligence <https://www.spglobal.com/marketintelligence/en/mi/products/us-monthly-gdp-index.html>

<sup>3</sup> Federal Reserve Economic Data <https://fred.stlouisfed.org/>

<sup>4</sup> Investing.com <https://www.investing.com/indices/us-spx-500-historical-data>

<sup>5</sup> Bureau of Labor Statistics <https://www.bls.gov/cpi/data.htm>

<sup>6</sup> <https://www.policyuncertainty.com/global-monthly.html>

<sup>7</sup> Author's website <https://www.matteoiacoviello.com/gpr.htm>

<sup>8</sup> Federal Reserve Board <https://www.federalreserve.gov/econres/notes/feds-notes/ebp.csv.csv>

<sup>9</sup> Authors' website <https://www.sydneyludvigson.com/macro-and-financial-uncertainty-indexes>

<sup>10</sup> Federal Reserve Board <https://www.frbsf.org/research-and-insights/data-and-indicators/market-based-monetary-policy-uncertainty/>

<sup>11</sup> Author's website <https://silviamirandaagrippino.com/code-data>

<sup>12</sup> Author's website [https://github.com/marekjarocinski/tshocks\\_fomc\\_update](https://github.com/marekjarocinski/tshocks_fomc_update)

<sup>13</sup> Federal Reserve Bank of Atlanta <https://www.atlantafed.org/cqer/research/wu-xia-shadow-federal-funds-rate>

## 4 Empirics

The empirical section is structured in two parts. I first estimate a small VAR that includes only the external instruments (monetary policy surprises) and a core set of macroeconomic variables whose responses to policy are assumed to be stable over time and are identified using sign restrictions. Therefore, the model consists of variables 1-9 in [Table 1](#). The VAR features stochastic volatility and student-t idiosyncratic errors, as changes in volatility and outliers are present in all macro data. This small-scale VAR serves as a baseline model and

illustrates the effectiveness of the identification strategy that combines sign and instrument-based restrictions. The results show that conventional monetary policy shocks can be identified robustly across a wide range of modeling choices, including the definition of instruments (target and path factors versus observed surprises), the choice of macroeconomic indicators (such as different measures of output, prices, or interest rates), and the specification of the VAR itself (including lag length, instrument exogeneity assumptions, and orthogonality conditions).

In the final part of the section, I expand the VAR by incorporating a richer set of variables, including disaggregated consumer price indices and forward-looking predictor variables. This model corresponds to the first 30 variables in [Table 1](#). The primary objective of this extension is to uncover how different components of the CPI responded to monetary policy after 2021, without imposing strong structural assumptions on the behavior of these disaggregated series. To this end, I retain the sign restrictions on the core macroeconomic variables to ensure identification of structural shocks in line with theory, while leaving the responses of the additional variables unrestricted and free to vary over time. These unrestricted impulse responses are entirely shaped by the data and are regularized through the use of shrinkage. This design allows the model to isolate structural shocks using theory-consistent restrictions on aggregate variables, and to trace out their effects on disaggregated prices in a flexible and data-driven way.

## 4.1 Monetary policy: small benchmark model

The first VAR is a small-scale benchmark model containing the two high-frequency monetary policy surprise factors (the “target” and “path” factors) alongside a core set of U.S. macroeconomic variables. The core endogenous variables include seven key aggregates observed at a monthly frequency – real economic activity (output), inflation, the policy rate<sup>7</sup> and short-term interest rates, a long-term interest rate, a monetary aggregate, and stock prices. I set the VAR lag order of this nine-variable VAR to  $p = 6$  (months) in the baseline

---

<sup>7</sup>In the benchmark specification I use the [\(Wu and Xia, 2016\)](#) shadow rate.

and I allow for stochastic volatility and student- $t$  errors to mitigate the influence of outliers and changes in volatility. However, for the sake of solidifying identification, these seven core macroeconomic variables have a time-invariant response to monetary policy shocks. In terms of the VAR formulations already introduced, this assumption corresponds to the model in equation (13) with the restriction  $\mathbf{\Lambda}_t = \mathbf{\Lambda}$ . This small VAR serves to illustrate conventional monetary transmission in a controlled setting before moving to a higher-dimensional model.

**Identification strategy.** I identify two structural monetary policy shocks in this VAR: a *Target shock*, corresponding to an unanticipated change in the current policy rate (conventional policy shock), and a *Path shock*, corresponding to news about future policy. The proposed identification strategy combines high-frequency external instruments with sign and zero restrictions on impact responses of select variables. High-frequency surprises around FOMC announcements are powerful proxies for policy shocks because they are measured in a narrow window that filters out most endogenous market responses. However, using these surprises alone to identify shocks can be problematic if they contain central-bank information effects or other confounding influences.

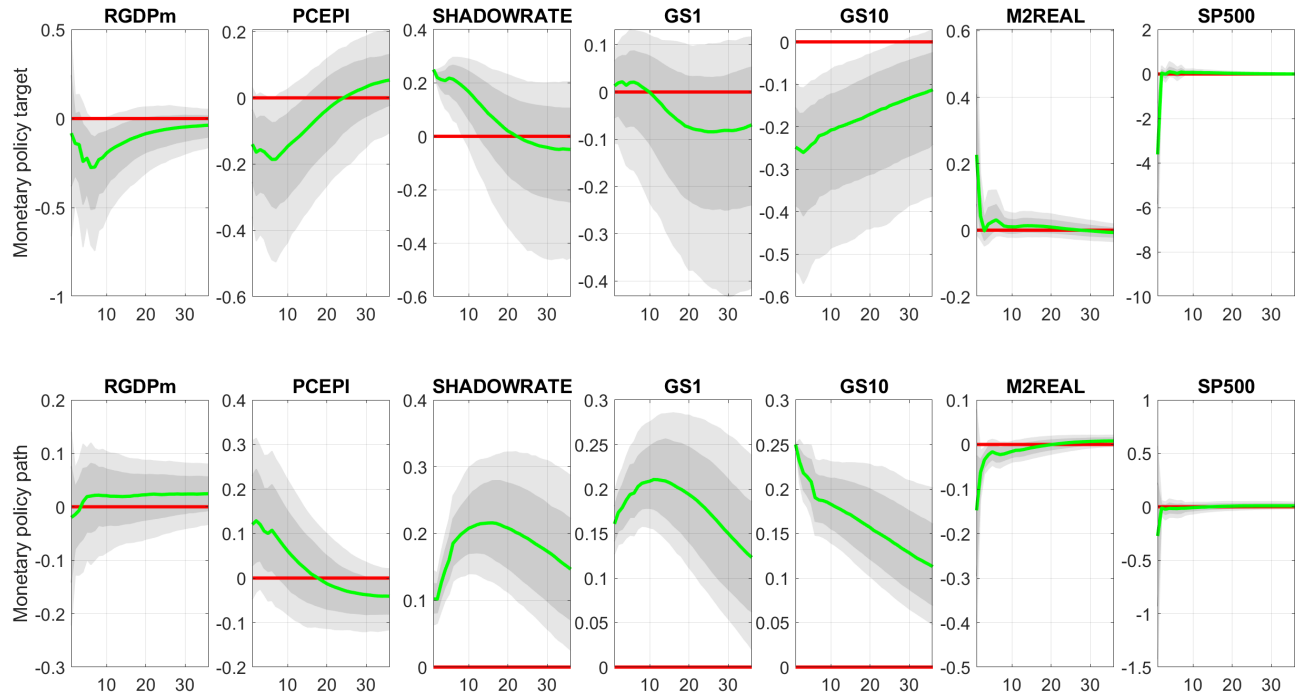


Figure 2: Responses of macro variables to Target (top row) and Path (bottom row) shocks in benchmark VAR without sign restrictions on macro variables. Solid green lines show the posterior median response, while dark grey and light grey bands represent the 68% and 90% posterior credible intervals, respectively; these summarize uncertainty across the posterior distribution of the structural impulse responses, reflecting both estimation error and parameter uncertainty.

Indeed, an initial attempt using only the two surprise factors to identify the shocks (imposing that the Target shock is the only driver of the target factor and the Path shock the only driver of the path factor, with no other restrictions) produced some counterintuitive results. As shown in Figure 2, a putative contractionary Target shock (tightening of the policy rate) did not yield a clearly negative effect on output or inflation – the posterior median responses of real GDP and PCE inflation were slightly negative but not significant – and it even led to a decline in the 10-year yield and a rise in real M2, contrary to typical theory predictions. Similarly, the identified Path shock had little effect on real activity and was associated with a slight rise in prices on impact. These anomalies echo concerns in the literature that high-frequency policy surprises may be contaminated by “information shocks” (when agents infer positive news about the economy from a rate hike) or may suffer from



Endogenous vars	Identified shocks		Residual shocks	
	Target shock	Path shock	Residual shock 1	Residual shock 2
<i>Target</i>	+	0	0	0
<i>Path</i>	0	+	0	0
-----	-----	-----	-----	-----
<i>realGDP</i>	−	<i>NA</i>	<i>NA</i>	<i>NA</i>
<i>PCEinf</i>	−	0	<i>NA</i>	<i>NA</i>
<i>FFR</i>	+	0	<i>NA</i>	<i>NA</i>
<i>GS1</i>	+	<i>NA</i>	<i>NA</i>	<i>NA</i>
<i>GS10</i>	0	+	<i>NA</i>	<i>NA</i>
<i>M2REAL</i>	−	<i>NA</i>	<i>NA</i>	<i>NA</i>
<i>SP500</i>	−	<i>NA</i>	<i>NA</i>	<i>NA</i>

Table 2: Sign restrictions imposed to core macro variables, which correspond to sign restrictions to the matrix  $\mathbf{\Lambda}$ . Matrix  $\mathbf{\Lambda}$  here is constant, not time-varying.

weak instrument problems, leading to implausible impulse responses. Recent studies have emphasized the need to disentangle pure policy shocks from central bank information components (e.g., [Jarociński, 2024](#)), and the findings here reinforce that message.

To pin down a conventional monetary tightening shock more robustly, I impose additional sign and zero restrictions on impact responses of the core macro variables, guided by economic theory. Table 2 summarizes the identification scheme. A Target shock is defined as an unanticipated rate hike that requires: (i) the target factor (current FFR surprise) to react positively, (ii) the path factor (expected future rate surprise) to be zero on impact, and (iii) core macro variables move in line with textbook predictions. With regards to the last point, output and the price level *fall* on impact signifying a contraction in demand and disinflationary pressure, the policy rate and short-term interest rates rise, real money supply falls due to tightening liquidity and credit conditions, and equity prices fall reflecting higher discount rates and weaker expected earnings. I leave the response of the 10-year yield unrestricted; prior intuition would suggest it likely rises or remains flat, as a monetary tightening can either raise long-term rates via expectations of future short rates or, in some cases, lower them if it strongly depresses future growth/inflation expectations. For the Path shock, I impose: (i) zero impact on the current policy rate (target factor) and a positive impact on

the path factor by definition, and (ii) no contemporaneous change in output and zero or muted impact on inflation.<sup>8</sup> A positive Path shock raises the 10-year yield on impact, since this shock should shift expectations of the rate trajectory. Other variables’ responses to the Path shock are left largely unrestricted. These hybrid identification restrictions – combining external instrument relevance with sign constraints on macroeconomic aggregates – ensure that the Target shock recovered corresponds to a standard contractionary monetary policy shock (stripped of information effects), while the Path shock is identified as a distinct structural disturbance capturing revisions to expected future policy.<sup>9</sup>

---

<sup>8</sup>I do not force inflation down for a path shock, allowing the data to determine if this shock acts like news of future policy.

<sup>9</sup>On top of these restrictions, in [Table 2](#) I specify two residual shocks that are not identified. This is due to the fact that in the proposed VAR model the number of shocks assumed not only affects identification, but also estimation accuracy. The number of shocks corresponds to the number of the estimated factors  $\mathbf{f}_t$ , which in turn affect how precisely-estimated is the “true” VAR covariance matrix  $\mathbf{\Omega}$ . Note that these additional residual factors need not be identified; the common component  $\mathbf{\Lambda}\mathbf{f}_t$  of these two factors is identified and this is sufficient for estimation using MCMC.




Figure 3: Responses of macro variables to Target (top row) and Path (bottom row) shocks in benchmark VAR with sign restrictions on macro variables that follow Table 2. Solid green lines show the posterior median response, while dark grey and light grey bands represent the 68% and 90% posterior credible intervals, respectively; these summarize uncertainty across the posterior distribution of the structural impulse responses, reflecting both estimation error and parameter uncertainty.

Figure 3 presents the impulse response functions (IRFs) of the core seven macro variables to a contractionary monetary policy shock under the preferred identification (combining the surprise instruments with sign restrictions). The responses are economically sensible and in line with conventional monetary transmission mechanisms and previous empirical evidence. Overall, the small VAR results demonstrate classic monetary transmission channels at work. A true contractionary policy shock (Target factor) leads to: higher short-term interest rates, a broad-based tightening of financial conditions (less money, lower equity values, likely higher credit spreads), a decline in aggregate demand (output falls) and a moderation in inflation. These outcomes accord with the broad consensus from decades of VAR studies and with standard New Keynesian theory (monetary tightening moves the economy along the aggregate demand curve, and the Phillips curve translates the induced slack into lower

inflation). Importantly, these effects are estimated on data that extend through the recent high-inflation episode of 2021–2022, suggesting that even in that atypical period, the fundamental channels remained operative. The Path shock, meanwhile, illustrates the more complex interplay of policy signaling and expectations. While it does not strongly contract the economy on impact, its existence highlights that not all Fed actions map neatly onto “surprise tightenings” – some reflect information updates.

In the subsequent analysis with a larger VAR this identification of two separate monetary shocks is maintained. Nevertheless, the focus there is on the conventional Target shock as the driver of systematic contractionary policy actions, while acknowledging the Path shock as an important but less clear-cut disturbance. Before moving to the large-scale model, I verify that the small VAR results are robust to various alternative specifications. In the Online Supplement I report detailed robustness checks, but I summarize them briefly here. First, increasing the lag order to  $p = 12$  or using tighter/slacker priors yields very similar IRFs, indicating that the baseline lag length (6) is sufficient to capture the dynamics. Second, I considered alternative identifying assumptions suggested by [Jarociński and Karadi \(2020\)](#): (a) restricting the lagged influence of the surprises on macro variables ( $\Phi_{ym} = \mathbf{0}$ , so that the high-frequency shocks only enter contemporaneously and not in lagged form), and (b) conversely restricting the feedback of macro shocks on the surprise series ( $\Phi_{mm} = \mathbf{0}$ ). Imposing these did not materially change the identified IRFs of interest – the output and inflation responses to a Target shock remained significantly negative and of similar magnitude. Third, replacing the [\(Wu and Xia, 2016\)](#) shadow rate with the effective Federal funds rate does not alter the qualitative pattern of responses of output and inflation, although the response of the 10-year rate becomes significantly negative. Finally, allowing the Target and Path shocks to be correlated (relaxing the orthogonality assumption) did not appreciably affect the impulse responses, though for our baseline I keep them orthogonal for a cleaner structural interpretation. All these checks bolster the confidence that the small VAR captures a genuine monetary tightening shock and its transmission in a theoretically consistent manner. In what

follows, I build on this foundation to examine a richer array of variables and potential time variation in responses.

## 4.2 Empirical results: large VAR with time-varying disaggregated CPI responses

I now turn to a large-scale Bayesian VAR that incorporates a much broader information set, including disaggregated price indices and financial market indicators. The primary goal of this extension is to study how a monetary policy shock propagates across different sectors and over time – in particular, how various components of the consumer price index (CPI) respond (heterogeneously) to policy tightening, and how financial conditions and expectations evolve alongside. By expanding the variable set, we also allow the possibility for the effects of monetary policy to change over the sample, since certain relationships may not be constant in the face of structural changes (e.g., the shift from a low-inflation regime to the high-inflation environment after 2021). The large VAR includes 30 monthly variables: the same two policy surprise factors and seven core macro variables from the small VAR, plus 13 additional macro-financial predictor series and the eight disaggregated CPI sub-indices. The benchmark case uses  $p = 2$  lags of all variables, as this choice yields more stable and interpretable impulse response functions. The Online Supplement shows that results are qualitatively unchanged when  $p = 6$ .

The identification of the monetary policy shocks in the large VAR follows the *same scheme as in the small VAR*: I impose the sign and zero restrictions on the impact responses of the core macro block (real activity down, inflation down for the Target shock; policy rate up, etc., as detailed earlier in Table 2), and similarly use the high-frequency target/path instruments. No additional sign restrictions are imposed on the new variables – instead, I “let the data speak” for how those variables respond. This design – identify shocks with a theoretically grounded core subset, but leave the rest unrestricted – is akin to a factor-augmented VAR or a large panel SVAR analysis. It ensures that the monetary policy shock that is recovered

in the large model is directly comparable to the one in the small model, while leveraging the larger information set to observe richer dynamics and potential time variation.




Figure 4: Responses of core macro variables to Target and Path shocks in the large VAR.  $\Lambda$  is constant for these variables, therefore, IRFs are fixed over the sample.

**Validating core responses in the large VAR.** I first verify that the introduction of many additional variables, and allowing for time-varying coefficients, does not distort the baseline monetary transmission mechanism. Figure 4 plots the impulse responses of the core seven macro variables to the Target and Path shocks as estimated in the large VAR. These core responses are restricted to be constant over time, hence the figure represents the common response for the entire sample. Reassuringly, the core IRFs are fairly similar to those from the small VAR (Figure 3). A contractionary Target shock in the large VAR still produces a clear downturn in real activity, a gradual decline in inflation, an uptick in short-term interest rates, and tightening of money and equity markets, all with magnitudes and timing comparable to the small-model results. In fact, the peak output and inflation effects are nearly identical, indicating that the identification is stable and the additional variables have not introduced

any ambiguity about the nature of the shock. This consistency is important: it implies that our shock remains a “conventional” monetary policy shock, and any new insights in the large VAR will come from observing the behavior of the additional variables rather than from redefining the shock itself. It also highlights that the Bayesian shrinkage and time-variation do not undermine the core inference; rather, the prior helps to integrate the extra information without losing the signal from the primary series. In summary, the large VAR’s core results pass a key sanity check giving confidence to proceed with analyzing the richer dynamics.

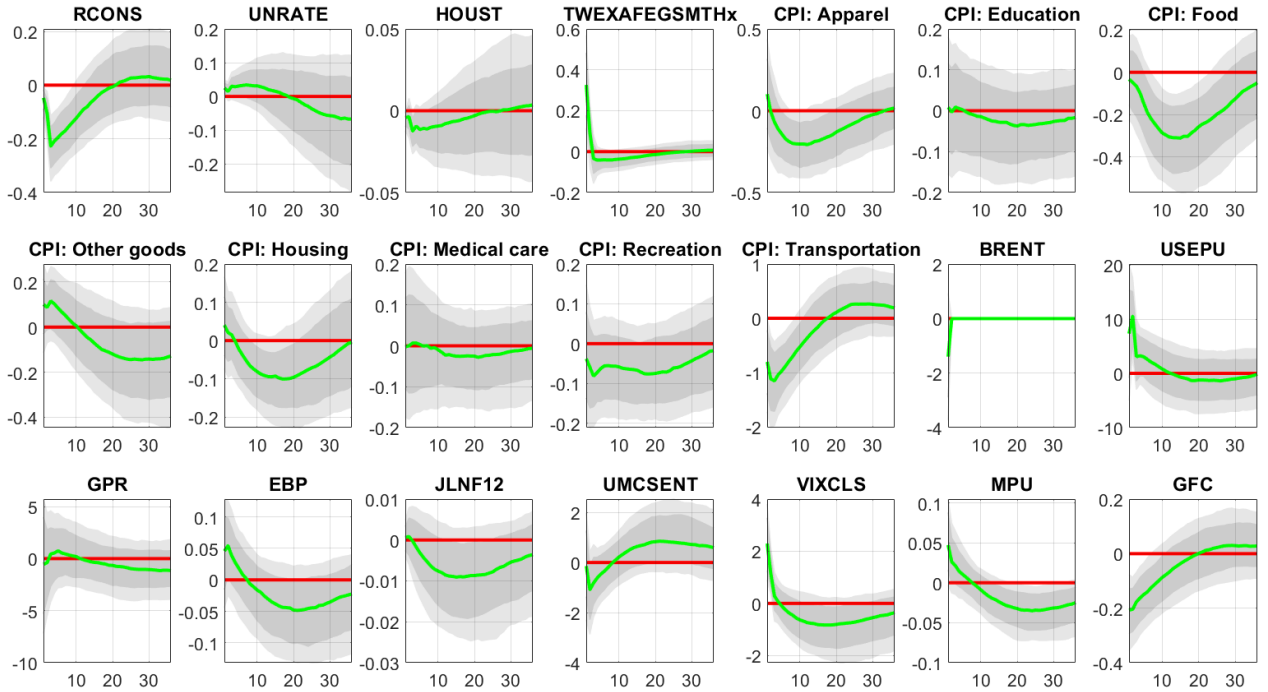


Figure 5: Responses of all other macro variables to Target shock in the large VAR.  $\Lambda$  is time-varying for these variables, therefore, plot gives responses at time  $T$ .






Figure 6: Responses of all other macro variables to Path shock in the large VAR.  $\Lambda$  is time-varying for these variables, therefore, plot gives responses at time  $T$ .

**Effects on broader financial conditions and expectations.** Expanding the VAR to include more financial indicators allows us to paint a more detailed picture of how monetary policy tightening percolates through financial markets and expectations. Figure 5 shows the impulse responses at the end of sample (May 2024) of all the additional variables to a Target shock, while Figure 6 does the same for a Path shock.<sup>10</sup>

I focus first on the Target shock (the conventional tightening). In short horizons the real variables have the correct sign: real consumption falls significantly, while unemployment rises and housing starts fall, although very marginally and insignificantly. A rate hike makes lenders and investors more risk-averse, leading them to demand higher premiums on risky debt, which in turn tightens credit supply to firms. The evidence in Figure 5 confirms that the monetary shock triggers a cascade where financial conditions deteriorate significantly. The Excess Bond Premium (EBP) captures the risk-related component of credit spreads

<sup>10</sup>Note that time-varying impact responses to these variables are analyzed later.

beyond expected defaults, and it typically rises after a contractionary monetary policy shock as investors become more risk-averse and demand higher compensation for holding corporate debt. Volatility (VIXCLS) and uncertainty measures (USEPU, MPU) also rise, reflecting increased trading volume and forward-looking uncertainty.

The responses to a Path shock (Figure 6, bottom panels) differ qualitatively, underscoring that this shock carries an informational element. Importantly, a Path shock does *not* cause credit spreads to widen – in fact, if anything, the estimates suggest a slight *narrowing* of corporate spreads immediately after a positive Path shock. At the same time, equity prices tend to increase on impact as the stock market often rallies in response to what appears to be a hawkish signal about the future. These financial reactions are consistent with a scenario in which the Fed’s guidance about future tightening is interpreted as a sign of confidence in economic strength. In other words, the Path shock behaves much like a positive “news” shock about the economy. This mirrors the [Jarociński and Karadi \(2020\)](#) result that an information shock (good news) leads to higher stock prices and lower credit spreads even as rates rise. In terms of other effects, a Path shock results in a non-significant increase in real consumption and housing starts, a significant currency depreciation, and decrease of economic policy uncertainty and stock market volatility. Unsurprisingly, monetary policy uncertainty does not increase, rather it rises significantly to this central-bank induced shock.

In summary, the large VAR confirms that the two identified shocks have distinct financial market and expectations footprints: the Target shock is an all-around contractionary disturbance with immediate tightening of financial conditions, while the Path shock is more of an expectations shock that can even be interpreted as “good news” at the moment of impact, with its contractionary bite coming later.

**Heterogeneous and time-varying inflation responses.** A central motivation for the large VAR is to examine how different components of inflation react to monetary policy shocks, and whether these reactions have changed in the recent high-inflation period.

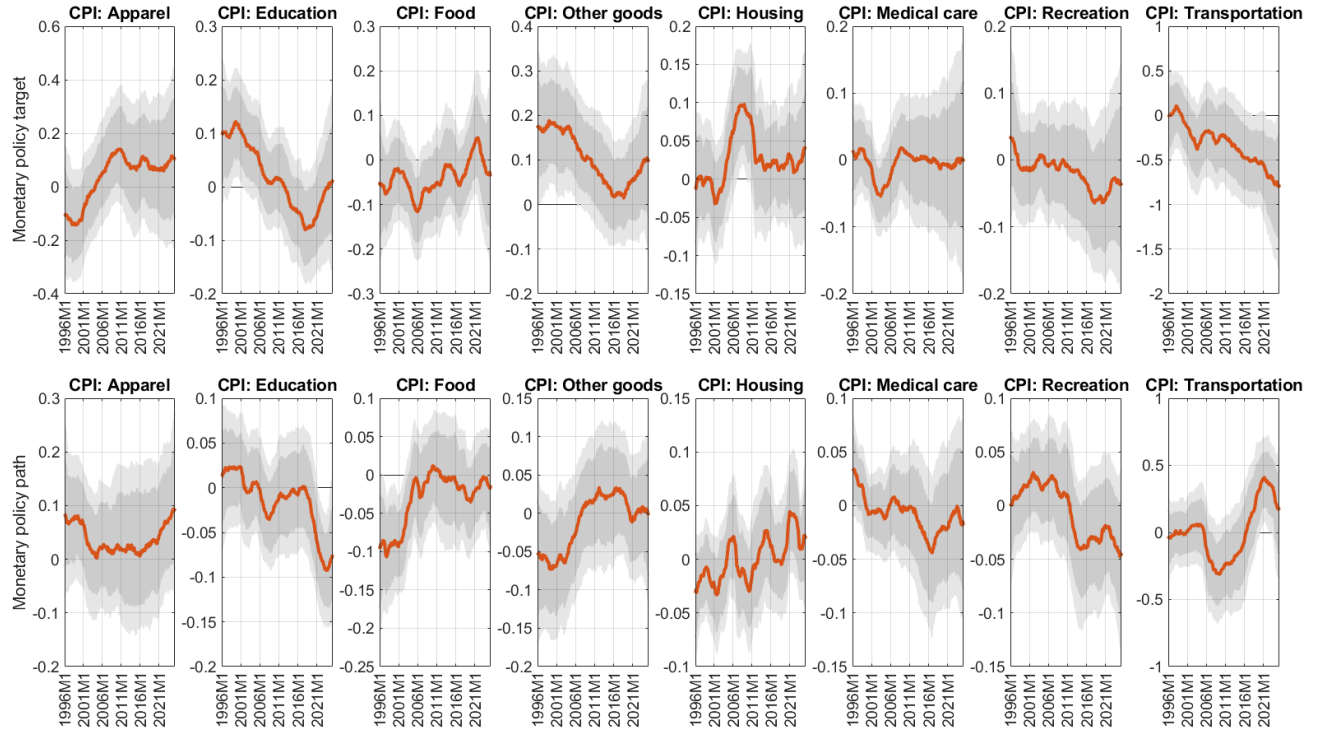


Figure 7: Impact responses of disaggregated CPI variables to Target and Path shocks.

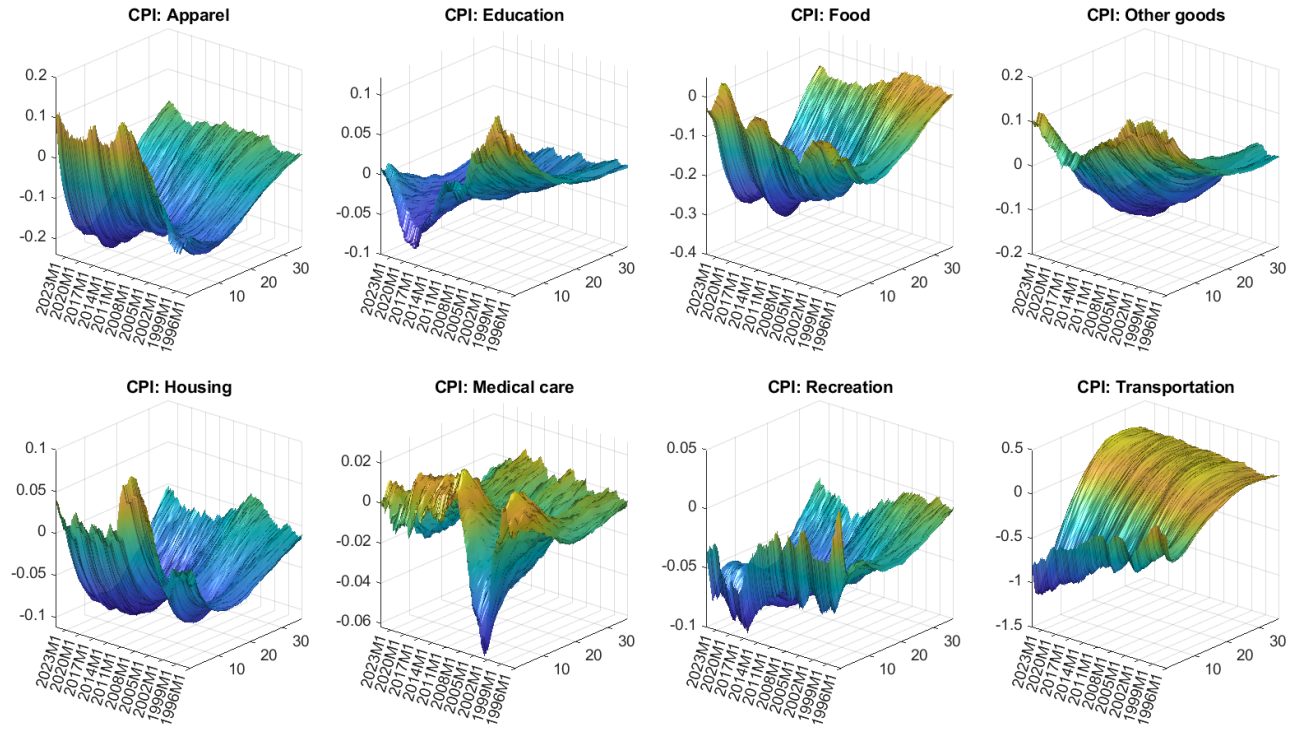


Figure 8: Responses over time and over horizons of disaggregated CPI variables to a Target shock.

Figure 7 presents the *impact* responses (at time  $t = 0$ ) of the eight disaggregated CPI categories to the Target and Path shocks, estimated at the end of our sample (i.e. for a shock occurring in May 2024). Figure 8 goes further, showing the evolution of the entire impulse response function of these CPI components to a Target shock over time (plotting the median responses at various horizons for shocks in different years).

Several noteworthy patterns emerge. First, the impact responses in Figure 7 reveal considerable heterogeneity across sectors in the instantaneous effect of a monetary tightening. A contractionary Target shock immediately lowers the rate of inflation in some categories, while others show virtually no reaction on impact. Rent inflation (CPI: Housing), which is a notoriously sticky and inertia-driven component of total inflation, shows essentially no movement on impact. Landlords do not reprice rental contracts in direct response to a rate hike; instead, the effect of monetary policy on rents operates with a long lag, via cooling of the housing market and slower rent growth over time. I indeed find in the three-dimensional IRFs that the response of rent inflation to a Target shock, while negligible on impact, becomes slightly negative a year or two after the shock – indicating that monetary tightening eventually tempers rent increases, but only gradually. Core goods prices (e.g., furnishings, electronics, vehicles), which are more interest-sensitive because they often involve credit financing (auto loans, for instance) and have more flexible pricing than services, show a notable immediate decline. For instance, the price index for vehicles and durable goods drops on impact after the Target shock, reflecting a sudden pullback in demand; higher loan rates discourage car and appliance purchases, leading retailers and manufacturers to offer discounts or slow price increases to stimulate sales. Meanwhile, food prices, which can be volatile but also heavily influenced by supply (weather, global commodity markets), show a very mild impact response – perhaps a slight dip, but within credible bands that include zero. Food and beverages may not immediately respond to domestic monetary conditions, except insofar as general demand changes slowly feed through to lower markups in food retail.

Turning to the Path shock impact responses: these are generally much smaller in absolute magnitude (often not significantly different from zero) for all CPI components. This again underscores that an expected future tightening alone does not immediately move prices – likely because neither aggregate demand nor costs are materially affected on impact. If anything, we can see a hint that a positive Path shock might cause a slight *uptick* in some prices like stock-sensitive goods or perhaps energy, consistent with the idea that it signals stronger future activity. However, these estimates are not very precise, and we should not ascribe too much economic significance to small positive bumps. The key takeaway is that unlike the Target shock, the Path shock does not uniformly push prices down at impact. Thus, from a policymaker’s perspective, an announced future hike (without current action) is not a substitute for a current hike in terms of immediately curbing inflation – its influence on prices will come later, mediated by expectations and subsequent activity changes.

Perhaps the most interesting insights come from the time variation in the Target shock responses (Figure 8). We find evidence that the responsiveness of certain inflation components to monetary policy shocks has evolved over the last few decades, especially in the post-2010 period and into the post-pandemic inflation surge. Notably, the decline in durable goods inflation after a monetary tightening has become more pronounced in the recent period. In the 1990s and early 2000s, a Target shock of similar size might have produced only a modest downturn in durable goods prices, partly because inflation was low and stable and firms were reluctant to cut prices. However, by 2022–2023, after the pandemic disruptions, our model infers that a contractionary shock yields a larger immediate reduction in durable goods inflation. One interpretation is that after the pandemic’s supply-driven price spikes (especially in vehicles and goods due to supply chain issues), there was pent-up disinflationary pressure once demand slowed. So when the Fed started tightening aggressively in 2022, it popped the “demand bubble” for goods, causing a faster normalization of goods prices (indeed, used car prices, for example, started falling in late 2022). Our time-varying IRFs capture this as a bigger impact of monetary shocks on goods inflation in the recent period.

Another dimension of heterogeneity is the peak timing of responses. Some components (energy, core goods) hit their maximum price decline within 3–6 months of the shock, whereas others (services) only show a significant change after 12–18 months. This staggered timing means that aggregate CPI inflation – a weighted sum of these – may exhibit a multi-phase reaction: an initial dip driven by energy and goods, potentially offset if energy prices rebound, and a later, more persistent slowing driven by services. This could explain why policymakers often see “headline” CPI inflation fall relatively quickly after tightening (as energy and food stabilize or drop), but “core” services inflation can remain high for a while, necessitating continued tight policy until it too comes down. Our large VAR results provide empirical backing for this narrative, which has been evident in the post-2021 episode: the Fed’s tightening in 2022 quickly lowered commodity-sensitive inflation by 2023, but core services inflation only started decelerating toward late 2023 and 2024.

Importantly, despite these differences in timing and magnitude, virtually all CPI components eventually move in the direction of disinflation in response to a contractionary Target shock (with the possible exception of certain idiosyncratic cases like medical care prices that might be governed by institutional factors). That is, monetary policy is a blunt tool that ultimately affects broad inflation – but the path it takes through various prices is uneven. This underscores the value of a large VAR approach: it lets us observe the rich cross-sectional detail behind the aggregate responses.

## 5 Conclusions

This paper develops a novel Bayesian inference framework for identifying monetary policy shocks in large-scale VARs, addressing longstanding challenges in structural identification. By combining high-frequency surprises from financial markets with economically motivated sign restrictions, the proposed method delivers robust and interpretable structural shocks, even in high-dimensional settings. A key innovation is the integration of this hybrid identification

strategy into a scalable Gibbs sampler that accommodates time-varying responses and fat-tailed disturbances, ensuring reliable inference in the presence of structural breaks and outliers.

Empirically, the paper offers new insights into the transmission of conventional monetary policy in the post-pandemic U.S. economy. By tracing the dynamic effects of interest rate surprises on disaggregated components of the consumer price index (CPI), the results highlight substantial heterogeneity in the speed and magnitude of inflation responses. While core goods prices exhibit swift and sizable declines following a rate hike, services (particularly housing) respond only gradually. These differences have become more pronounced during the 2022–24 inflationary episode, underscoring shifts in the monetary transmission mechanism and the importance of disaggregated analysis.

More broadly, the findings affirm that conventional monetary policy retains its effectiveness in shaping inflation dynamics, even in periods marked by elevated uncertainty and atypical price behavior. For policymakers, this underscores the importance of sector-specific diagnostics when assessing the stance and expected impact of policy. For researchers, the methodological framework provides a general-purpose tool for structural analysis in large systems, adaptable to a wide range of macroeconomic questions beyond monetary policy.



## References

- Acosta, M. (2023). The perceived causes of monetary policy surprises. Available at [https://www.acostamiguel.com/papers/acosta\\_mp.pdf](https://www.acostamiguel.com/papers/acosta_mp.pdf), accessed March 2024.
- Acosta, M., Brennan, C. M., and Jacobson, M. M. (2024). Constructing high-frequency monetary policy surprises from SOFR futures. *Economics Letters*, 242:111873.
- Arias, J. E., Caldara, D., and Rubio-Ramírez, J. F. (2019). The systematic component of monetary policy in SVARs: An agnostic identification procedure. *Journal of Monetary Economics*, 101:1–13.
- Aruoba, S. B. and Drechsel, T. (2024). The long and variable lags of monetary policy: Evidence from disaggregated price indices. *Journal of Monetary Economics*, 148:103635. Inflation in the COVID Era and Beyond.
- Bańbura, M., Giannone, D., and Reichlin, L. (2010). Large Bayesian vector auto regressions. *Journal of Applied Econometrics*, 25(1):71–92.
- Bauer, M. D. and Swanson, E. T. (2023). An alternative explanation for the “Fed information effect”. *American Economic Review*, 113(3):664–700.
- Bernanke, B. S. and Boivin, J. (2003). Monetary policy in a data-rich environment. *Journal of Monetary Economics*, 50(3):525–546. Swiss National Bank/Study Center Gerzensee Conference on Monetary Policy under Incomplete Information.
- Braun, R. and Brüggemann, R. (2023). Identification of SVAR models by combining sign restrictions with external instruments. *Journal of Business & Economic Statistics*, 41(4):1077–1089.
- Caldara, D. and Iacoviello, M. (2022). Measuring geopolitical risk. *American Economic Review*, 112(4):1194–1225.

- Carriero, A., Clark, T. E., Marcellino, M., and Mertens, E. (2024a). Addressing COVID-19 outliers in BVARs with stochastic volatility. *The Review of Economics and Statistics*, 106(5):1403–1417.
- Carriero, A., Marcellino, M., and Tornese, T. (2024b). Blended identification in structural VARs. *Journal of Monetary Economics*, page 103581.
- Carvalho, C. M., Polson, N. G., and Scott, J. G. (2010). The horseshoe estimator for sparse signals. *Biometrika*, 97(2):465–480.
- Casini, A. and McCloskey, A. (2024). Limitations of high-frequency event studies for identifying causal effects of monetary policy. *arXiv preprint arXiv:2406.15667*.
- Chan, J., Eisenstat, E., and Yu, X. (2022). Large bayesian VARs with factor stochastic volatility: Identification, order invariance and structural analysis. *arXiv preprint arXiv:2207.03988*. Accessed: 2025-02-09.
- Geweke, J. (1993). Bayesian treatment of the independent student-t linear model. *Journal of Applied Econometrics*, 8(S1):S19–S40.
- Giannone, D., Lenza, M., and Primiceri, G. E. (2015). Prior selection for vector autoregressions. *The Review of Economics and Statistics*, 97(2):436–451.
- Gorodnichenko, Y. (2005). Reduced-rank identification of structural shocks in VARs. SSRN working paper 590906, University of California, Berkeley.
- Gürkaynak, R. S., Sack, B., and Swanson, E. (2005). Do actions speak louder than words? The response of asset prices to monetary policy actions and statements. *International Journal of Central Banking*, 1(1).
- Jarociński, M. (2024). Estimating the Fed’s unconventional policy shocks. *Journal of Monetary Economics*, 144:103548.

- Jarociński, M. and Karadi, P. (2020). Deconstructing monetary policy surprises—The role of information shocks. *American Economic Journal: Macroeconomics*, 12(2):1–43.
- Jurado, K., Ludvigson, S. C., and Ng, S. (2015). Measuring uncertainty. *American Economic Review*, 105(3):1177–1216.
- Korobilis, D. (2021). High-dimensional macroeconomic forecasting using message passing algorithms. *Journal of Business & Economic Statistics*, 39(2):493–504.
- Korobilis, D. (2022). A new algorithm for structural restrictions in Bayesian vector autoregressions. *European Economic Review*, 148:104241.
- Lakdawala, A. (2019). Decomposing the effects of monetary policy using an external instruments SVAR. *Journal of Applied Econometrics*, 34(6):934–950.
- Lenza, M. and Primiceri, G. E. (2022). How to estimate a vector autoregression after March 2020. *Journal of Applied Econometrics*, 37(4):688–699.
- Mertens, K. and Ravn, M. O. (2013). The dynamic effects of personal and corporate income tax changes in the United States. *American Economic Review*, 103(4):1212–47.
- Miranda-Agrippino, S. and Rey, H. (2020). U.S. monetary policy and the global financial cycle. *The Review of Economic Studies*, 87(6):2754–2776.
- Mumtaz, H. and Petrova, K. (2023). Changing impact of shocks: A time-varying proxy SVAR approach. *Journal of Money, Credit and Banking*, 55(2-3):635–654.
- Nakamura, E. and Steinsson, J. (2018). High-frequency identification of monetary non-neutrality: The information effect. *The Quarterly Journal of Economics*, 133(3):1283–1330.
- Paccagnini, A. and Parla, G. (2021). Identifying high-frequency shocks in a Bayesian mixed-frequency VAR. *CAMA Working Paper Series*, 2021-85.
- Prüser, J. (2024). A large non-Gaussian structural VAR with application to monetary policy.

- Read, M. (2024). Set-identified structural vector autoregressions and the effects of a 100 basis point monetary policy shock. *The Review of Economics and Statistics*, page Online Early accepted version.
- Schlaak, T., Rieth, M., and Podstawski, M. (2023). Monetary policy, external instruments, and heteroskedasticity. *Quantitative Economics*, 14:161–200.
- Sims, C. A. and Zha, T. (2006). Were there regime switches in u.s. monetary policy? *American Economic Review*, 96(1):54–81.
- Swanson, E. T. (2021). Measuring the effects of Federal Reserve forward guidance and asset purchases on financial markets. *Journal of Monetary Economics*, 118:32–53.
- Wu, C. and Xia, F. D. (2016). Measuring the macroeconomic impact of monetary policy at the zero lower bound. *Journal of Money, Credit and Banking*, 48(2–3):253–291.

# Online Supplement to “Exploring Monetary Policy Shocks with Large-Scale Bayesian VARs”

Dimitris Korobilis<sup>1,2,3</sup>

<sup>1</sup>*University of Glasgow*

<sup>2</sup>*Center for Applied Macroeconomics and Commodity Prices (CAMP)*

<sup>3</sup>*Rimini Center for Economic Analysis (RCEA)*

May 9, 2025

## A Technical Appendix

### A.1 A Gibbs sampler for the constant parameter model

Recall the large monetary policy VAR with a factor structure in the residuals, as given in equation (6) of the main text. For convenience, define  $\mathbf{y}_t^* = [\mathbf{m}_t', \mathbf{y}_t']'$  as the  $N \times 1$  vector collecting the  $m$  high-frequency instruments and  $n$  macroeconomic variables (so  $N = m + n$ ). The VAR for  $\mathbf{y}_t^*$  can be written as

$$\mathbf{y}_t^* = \Phi \mathbf{x}_t + \Gamma^* \mathbf{f}_t + \mathbf{W}^{*1/2} \boldsymbol{\eta}_t,$$

where  $\mathbf{x}_t$  contains the intercept and lags (as in a standard VAR),  $\Phi$  is the  $N \times k$  matrix of VAR slope coefficients (with  $k$  the number of regressors per equation),  $\mathbf{f}_t$  is an  $r \times 1$  vector of latent factors (structural shocks) with  $r \ll N$ , and  $\boldsymbol{\eta}_t$  is an  $N \times 1$  idiosyncratic noise vector. We collect the factor loadings into the  $N \times r$  matrix  $\Gamma^* = \begin{pmatrix} \Gamma \\ \Lambda \end{pmatrix}$ , stacking the  $m \times r$

---

Correspondence: Professor of Econometrics, Adam Smith Business School, University of Glasgow, 2 Discovery Place, Glasgow, G11 6EY, United Kingdom; email: [dikorobilis@googlemail.com](mailto:dikorobilis@googlemail.com).

loadings  $\mathbf{\Gamma}$  for the instrument equations and the  $n \times r$  loadings  $\mathbf{\Lambda}$  for the macroeconomic equations. The idiosyncratic disturbances have covariance matrix  $\mathbf{W}^* = \text{diag}(\mathbf{W}, \mathbf{\Sigma})$ , with  $\mathbf{W} = \text{diag}(w_1, \dots, w_m)$  and  $\mathbf{\Sigma} = \text{diag}(\sigma_1^2, \dots, \sigma_n^2)$  for the instrument and macro equations respectively. For notational simplicity in what follows, let  $\sigma_i^2$  denote the  $i$ th diagonal element of  $\mathbf{W}^*$  (covering both  $w_i$  for  $i \leq m$  and  $\sigma_i^2$  for  $i > m$ ). By construction, we treat all variables in  $\mathbf{y}_t^*$  as endogenous in the VAR.

We impose a hierarchical shrinkage prior on the VAR coefficients and independent Gaussian (or truncated Gaussian) priors on the factor loadings. Specifically, for each equation  $i = 1, \dots, N$  of the VAR, we let  $\phi_i$  denote the  $k \times 1$  vector of regression coefficients (the  $i$ th row of  $\mathbf{\Phi}$ ), and  $\mathbf{\Gamma}_i^*$  denote the  $1 \times r$  vector of factor loadings (the  $i$ th row of  $\mathbf{\Gamma}^*$ ). The prior is given by:

$$\phi_i \mid \sigma_i^2, \tau_i^2, \mathbf{\Psi}_i^2 \sim \mathcal{N}(\mathbf{0}, \sigma_i^2 \tau_i^2 \mathbf{\Psi}_i^2), \quad \mathbf{\Psi}_i^2 = \text{diag}(\psi_{i,1}^2, \dots, \psi_{i,k}^2), \quad (\text{A.1})$$

$$\psi_{i,j} \sim \mathcal{C}^+(0, 1), \quad j = 1, \dots, k, \quad (\text{A.2})$$

$$\tau_i \sim \mathcal{C}^+(0, 1), \quad (\text{A.3})$$

$$\mathbf{\Gamma}_{ij}^* \sim \begin{cases} \mathcal{N}(0, \underline{h}_{ij}) I(\mathbf{\Gamma}_{ij}^* > 0), & \text{if } S_{ij} = +1, \\ \mathcal{N}(0, \underline{h}_{ij}) I(\mathbf{\Gamma}_{ij}^* < 0), & \text{if } S_{ij} = -1, \\ \delta_0(\mathbf{\Gamma}_{ij}^*), & \text{if } S_{ij} = 0, \\ \mathcal{N}(0, \underline{h}_{ij}), & \text{otherwise,} \end{cases} \quad j = 1, \dots, r, \quad (\text{A.4})$$

$$\mathbf{f}_t \sim \mathcal{N}_r(\mathbf{0}, \mathbf{I}_r), \quad (\text{A.5})$$

$$\sigma_i^2 \sim \text{Inv-Gamma}(\underline{\rho}_i, \underline{\kappa}_i). \quad (\text{A.6})$$

Equation (A.1) is the Horseshoe prior of [Carvalho et al. \(2010\)](#) on the VAR coefficients:  $\psi_{i,j}$  and  $\tau_i$  are local and global shrinkage parameters (with half Cauchy priors), so that  $\phi_i \mid \sigma_i^2, \tau_i, \psi_{i,j}$  is distributed as a multivariate normal with a diagonal covariance matrix  $\sigma_i^2 \tau_i^2 \mathbf{\Psi}_i^2$ . Equations (A.2)–(A.3) are the half Cauchy priors for these shrinkage hyperparameters. In

(A.4), we define  $S_{ij}$  as an optional sign indicator for each factor loading: if  $S_{ij} = +1$  ( $-1$ ), we impose a priori that  $\Gamma_{ij}^*$  must be positive (negative) by using a truncated normal prior restricted to  $(0, \infty)$  ( $(-\infty, 0)$ ); if  $S_{ij} = 0$ , we impose a zero restriction (point mass at 0); otherwise  $S_{ij}$  is empty and the prior is standard Gaussian with variance  $\underline{h}_{ij}$ . In practice we set  $\underline{h}_{ij} = 4$  as a relatively diffuse prior variance for all unrestricted loadings. Finally, (A.5)–(A.6) state that each latent factor is *a priori* standard normal and each idiosyncratic variance has an inverse gamma prior. We choose hyperparameters  $\underline{\rho}_i = 1$  and  $\underline{\kappa}_i = 0.01$  for all  $i$ , which is a weakly informative prior favoring small  $\sigma_i^2$  (this helps regularize the estimation of a high-dimensional covariance matrix).

Given these priors and the observed data  $\{\mathbf{y}_t^*\}_{t=1}^T$ , Bayesian inference is carried out by Gibbs sampling. The algorithm sequentially draws from the complete conditional posterior of each block of parameters in turn. We label this the **factor sign restrictions (FSR) sampler**, which proceeds as follows:

1. **Sample  $\Phi$  (VAR coefficients).** For each equation  $i = 1, \dots, N$ , draw  $\phi_i$  from its conditional normal posterior. Because the VAR can be written equation-by-equation as

$$y_{i,t}^* = \phi_i \mathbf{x}_t + \Gamma_i^* \mathbf{f}_t + \sigma_i \eta_{i,t},$$

and  $\mathbf{f}_t$  is currently treated as given, this is a standard linear regression for  $\phi_i$ . The prior for  $\phi_i$  in (A.1) is normal with covariance  $\sigma_i^2 \tau_i^2 \Psi_i^2$ . Let  $\tilde{y}_{i,t} = y_{i,t}^* - \Gamma_i^* \mathbf{f}_t$  denote the regression residual for equation  $i$  after removing the factor component. Then the conditional posterior is:

$$\phi_i \mid \Phi_{-i}, \Gamma^*, \mathbf{W}^*, \mathbf{f}_{1:T}, \mathbf{y}_{1:T}^* \sim \mathcal{N}_k(\bar{\mathbf{V}}_i (\sum_{t=1}^T \sigma_i^{-2} \mathbf{x}_t' \tilde{y}_{i,t}), \bar{\mathbf{V}}_i), \quad (\text{A.7})$$

where  $\bar{\mathbf{V}}_i^{-1} = \Sigma_{\phi_i}^{-1} + \sum_{t=1}^T \sigma_i^{-2} \mathbf{x}_t' \mathbf{x}_t$  and  $\Sigma_{\phi_i} = \sigma_i^2 \tau_i^2 \Psi_i^2$  is the prior covariance of  $\phi_i$ . Intuitively,  $\bar{\mathbf{V}}_i$  is the posterior covariance, given by the prior precision  $\Sigma_{\phi_i}^{-1}$  plus the

data precision  $\frac{1}{\sigma_i^2} \mathbf{X}'\mathbf{X}$ , and the posterior mean is a precision-weighted combination of the OLS estimate and prior mean 0. We sample  $\phi_i$  efficiently using the algorithm of [Bhattacharya et al. \(2016\)](#).

2. **Sample local shrinkage**  $\psi_{i,j}$ . For each  $i = 1, \dots, N$  and  $j = 1, \dots, k$ , update the Horseshoe local scale  $\psi_{i,j}$  via slice sampling ([Neal, 2003](#)). We introduce the auxiliary variable  $\eta_{ij} = 1/\psi_{i,j}^2$  and draw as follows:

- (a) Set  $\eta_{ij} = 1/\psi_{i,j}^2$  (using the most recent draw of  $\psi_{i,j}$ ).
- (b) Draw  $u \sim \mathcal{U}(0, 1/(1 + \eta_{ij}))$ .
- (c) Draw a new  $\eta_{ij}$  from its conditional density, which is proportional to  $\exp\{-\frac{\phi_{i,j}^2}{2\sigma_i^2}\eta_{ij}\}$  truncated to  $\eta_{ij} > \frac{u}{1-u}$ . Finally set  $\psi_{i,j} = 1/\sqrt{\eta_{ij}}$ .

(Step 2 exploits the scale-mixture representation of the half-Cauchy distribution; see [Neal, 2003](#) for details.)

3. **Sample global shrinkage**  $\tau_i$ . For each equation  $i = 1, \dots, N$ , update the Horseshoe global scale  $\tau_i$  via slice sampling:

- (a) Set  $\xi_i = 1/\tau_i^2$ .
- (b) Draw  $v \sim \mathcal{U}(0, 1/(1 + \xi_i))$ .
- (c) Draw  $\xi_i$  from its conditional density, which is proportional to  $\xi_i^{-(k+1)/2} \exp\{-\frac{\xi_i}{2\sigma_i^2} \sum_{j=1}^k (\phi_{i,j}/\psi_{i,j})^2\}$  for  $\xi_i > \frac{v}{1-v}$ . Then set  $\tau_i = 1/\sqrt{\xi_i}$ .

(In step (c) we sample  $\xi_i$  by inverting the incomplete gamma function as in [Bhattacharya et al., 2016](#).)

4. **Sample factor loadings**  $\mathbf{\Gamma}^*$ . Draw each element of the loading matrix from its conditional posterior, which is Gaussian or truncated Gaussian. Conditional on the



current values of  $\{\Phi, \mathbf{f}_{1:T}, \mathbf{W}^*\}$ , the latent factor model for equation  $i$  is

$$y_{i,t}^* - \phi_i \mathbf{x}_t = \Gamma_i^* \mathbf{f}_t + \sigma_i \eta_{i,t},$$

with prior  $\Gamma_i^* \sim \mathcal{N}_r(\mathbf{0}, \underline{\mathbf{H}}_i)$  where  $\underline{\mathbf{H}}_i = \text{diag}(\underline{h}_{i1}, \dots, \underline{h}_{ir})$  as in (A.4). Let  $\hat{y}_{i,t} \equiv \varepsilon_{i,t} = y_{i,t}^* - \phi_i \mathbf{x}_t$  be the  $i$ th equation's residual after removing the VAR component. Stacking over  $t = 1, \dots, T$ , the conditional likelihood for  $\Gamma_i^*$  is as in a standard Bayesian static factor model: the posterior for  $\Gamma_i^*$  is normal with covariance  $\overline{\mathbf{H}}_i = \left( \underline{\mathbf{H}}_i^{-1} + \sum_{t=1}^T \sigma_i^{-2} \mathbf{f}_t' \mathbf{f}_t \right)^{-1}$  and mean  $\overline{\mathbf{H}}_i \left( \sum_{t=1}^T \sigma_i^{-2} \mathbf{f}_t' \hat{y}_{i,t} \right)$ . Denote by  $\bar{\lambda}_{ij}$  and  $\bar{h}_{ij}$  the  $j$ th element of this posterior mean and variance, respectively, for loading  $\Gamma_{ij}^*$ . We sample each  $\Gamma_{ij}^*$  one-at-a-time from this multivariate normal, imposing any sign or zero restrictions via univariate truncation. In particular, conditioning on all other loadings  $\Gamma_{i,-(j)}^*$  in row  $i$ , the scalar  $\Gamma_{ij}^*$  has a univariate normal conditional posterior,

$$\Gamma_{ij}^* \mid \Gamma_{i,-(j)}^*, \Phi, \mathbf{W}^*, \mathbf{f}_{1:T}, \mathbf{y}_{1:T}^* \sim \mathcal{N}\left(\bar{\lambda}_{ij} - \bar{h}_{ij} \sum_{\ell \neq j} \bar{h}_{i\ell}^{-1} (\Gamma_{i\ell}^* - \bar{\lambda}_{i\ell}), \bar{h}_{ij}\right),$$

truncated to the interval  $(a_{ij}, b_{ij})$  as appropriate, where

$$(a_{ij}, b_{ij}) = \begin{cases} (-\infty, 0), & S_{ij} = -1, \\ (0, \infty), & S_{ij} = +1, \\ (0, 0), & S_{ij} = 0, \\ (-\infty, \infty), & \text{otherwise.} \end{cases}$$

Here  $(0, 0)$  denotes a degenerate distribution at zero (for  $S_{ij} = 0$ , we simply set  $\Gamma_{ij}^* = 0$ ). We sample each restricted element using the efficient univariate truncated normal generator of Botev (2017). Unrestricted loadings ( $S_{ij}$  empty) are drawn directly from the (untruncated) normal posterior. This element-by-element Gibbs step (see Geweke, 1996) ensures that any sign and zero restrictions on the rows of  $\Gamma^*$  are satisfied at each

draw.

5. **Sample latent factors  $\mathbf{f}_t$ .** Draw each factor vector  $\mathbf{f}_t$  for  $t = 1, \dots, T$  from its conditional multivariate normal posterior. Conditional on all other parameters, the model for  $\mathbf{f}_t$  is

$$\hat{\mathbf{y}}_t^* = \mathbf{\Gamma}^* \mathbf{f}_t + \mathbf{v}_t,$$

where  $\hat{\mathbf{y}}_t^* = \mathbf{y}_t^* - \mathbf{\Phi} \mathbf{x}_t$  is the  $N \times 1$  vector of current residuals (after removing the VAR fit), and  $\mathbf{v}_t = \mathbf{W}^{*1/2} \boldsymbol{\eta}_t$  has covariance  $\mathbf{W}^*$ . Combining the Gaussian likelihood with the prior  $\mathbf{f}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_r)$  yields:

$$\mathbf{f}_t \mid \mathbf{\Gamma}^*, \mathbf{W}^*, \mathbf{\Phi}, \mathbf{y}_{1:T}^* \sim \mathcal{N}_r(\overline{\mathbf{G}} \mathbf{\Gamma}^{*'} (\mathbf{W}^*)^{-1} \hat{\mathbf{y}}_t^*, \overline{\mathbf{G}}), \quad (\text{A.8})$$

with  $\overline{\mathbf{G}}^{-1} = \mathbf{I}_r + \mathbf{\Gamma}^{*'} (\mathbf{W}^*)^{-1} \mathbf{\Gamma}^*$ . This  $r \times r$  system is of low dimension, so sampling  $\mathbf{f}_t$  is fast. After obtaining draws  $\{\mathbf{f}_t\}_{t=1}^T$ , we post-process the  $T \times r$  matrix  $\mathbf{F} = (\mathbf{f}_1, \dots, \mathbf{f}_T)'$  to ensure the factors (interpreted as structural shocks) are orthonormal. In particular, we apply a Gram-Schmidt orthogonalization to the columns of  $\mathbf{F}$ , and normalize each column to have unit variance (and, if desired, multiply by  $-1$  to enforce any overall sign convention). This post-processing step resolves the sign and scale indeterminacy of the factors without loss of generality.

6. **Sample idiosyncratic variances  $\sigma_i^2$ .** Finally, update each variance  $\sigma_i^2$  (the  $i$ th diagonal element of  $\mathbf{W}^*$ ) from its inverse-gamma conditional posterior. Given current draws of  $\{\mathbf{\Phi}, \mathbf{\Gamma}^*, \mathbf{f}_{1:T}\}$ , the residuals  $\varepsilon_{i,t} = y_{i,t}^* - \phi_i \mathbf{x}_t - \mathbf{\Gamma}_i^* \mathbf{f}_t$  are observable. The conditional posterior for  $\sigma_i^2$  combines the prior (A.6) with the likelihood  $\prod_{t=1}^T \mathcal{N}(\varepsilon_{i,t} \mid 0, \sigma_i^2)$ . This yields

$$\sigma_i^2 \mid \mathbf{\Phi}, \mathbf{\Gamma}^*, \mathbf{f}_{1:T}, \mathbf{y}_{1:T}^* \sim \text{Inv-Gamma}\left(\underline{\rho}_i + \frac{T}{2}, \underline{\kappa}_i + \frac{1}{2} \sum_{t=1}^T \varepsilon_{i,t}^2\right). \quad (\text{A.9})$$

Thus each  $\sigma_i^2$  is drawn by conjugacy (closed-form inverse-gamma).

This Gibbs sampler iterates over steps 1–6 until convergence. Steps 1, 4, and 5 are standard conditional posteriors for linear Gaussian models (VAR and factor model); step 2–3 are specific to the Horseshoe prior; and step 4 implements any desired sign/zero restrictions by univariate truncation. Thinning of the MCMC draws (only saving every  $s$ th draw, for some  $s > 1$ ) is recommended to mitigate any residual autocorrelation in the Markov chain, since the latent factors and loadings enter the likelihood multiplicatively and can be slightly negatively correlated in the Gibbs updates. Convergence and mixing of the sampler can be monitored with standard MCMC diagnostics.

## A.2 Adding time-variation: sampling time-varying loadings and volatilities

We now extend the algorithm to handle time-varying factor loadings and stochastic volatility, as described in Section 4 of the main text. The time-varying parameter version of the model (see equations (13)–(14)) allows the macroeconomic impact matrix  $\mathbf{\Lambda}_t$  and idiosyncratic variances  $\mathbf{\Sigma}_t = \text{diag}(\sigma_{1,t}^2, \dots, \sigma_{n,t}^2)$  to evolve over time. In particular,  $\mathbf{\Lambda}_t$  and the log-volatilities follow independent random walks:

$$\text{vec}(\mathbf{\Lambda}_t) = \text{vec}(\mathbf{\Lambda}_{t-1}) + \boldsymbol{\xi}_t, \quad (\text{A.10})$$

$$\log(\sigma_{i,t}^2) = \log(\sigma_{i,t-1}^2) + \delta_{i,t}, \quad i = 1, \dots, n, \quad (\text{A.11})$$

with  $\boldsymbol{\xi}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$  and  $\delta_{i,t} \sim \mathcal{N}(0, \omega_i^2)$ . Here  $\mathbf{Q}$  is an  $(nr) \times (nr)$  covariance matrix governing the innovations in the  $n \times r$  loading matrix, and  $\omega_i^2$  is the variance of shocks to the log variance of series  $i$ . Equation (A.10) implies a prior belief that  $\mathbf{\Lambda}_t$  changes gradually, centered around its previous value  $\mathbf{\Lambda}_{t-1}$  (similarly for  $\log \sigma_{i,t}^2$  around  $\log \sigma_{i,t-1}^2$ ). This is the familiar state-space formulation for time-varying parameters, which is typically handled via Kalman filtering and smoothing. In our high-dimensional setting, however, it is more computationally convenient to use the *stacked regression* approach of Korobilis (2021). That is, we rewrite the entire

sample in a single regression equation where the time-varying loadings enter as additional static coefficients to be estimated.

To see this, note that we can integrate the state equations to express the loadings at time  $t$  as  $\mathbf{\Lambda}_t = \mathbf{\Lambda}_0 + \sum_{s=1}^t \boldsymbol{\xi}_s$ . For simplicity, assume the initial prior mean is  $\mathbf{\Lambda}_0 = \mathbf{0}$  (one can also treat  $\mathbf{\Lambda}_0$  as an additional parameter to estimate, but setting it to zero with large variance is practically equivalent). Then  $\mathbf{\Lambda}_t = \sum_{s=1}^t \Delta\mathbf{\Lambda}_s$ , where we define  $\Delta\mathbf{\Lambda}_t \equiv \boldsymbol{\xi}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$ . The observation equation for  $\mathbf{y}_t^*$  can be written as

$$\mathbf{y}_t^* = \Phi \mathbf{x}_t + [\Gamma, \mathbf{\Lambda}_t] \mathbf{f}_t + \mathbf{W}_t^{*1/2} \boldsymbol{\eta}_t,$$

where  $\mathbf{W}_t^* = \text{diag}(\mathbf{W}, \Sigma_t)$ . Plugging in  $\mathbf{\Lambda}_t = \sum_{s=1}^t \Delta\mathbf{\Lambda}_s$ , we obtain

$$\mathbf{y}_t^* = \Phi \mathbf{x}_t + \Gamma \mathbf{f}_t + \sum_{s=1}^t \Delta\mathbf{\Lambda}_s \mathbf{f}_t + \mathbf{W}_t^{*1/2} \boldsymbol{\eta}_t.$$

Now consider all  $T$  observations together. By stacking the above equation for  $t = 1$  to  $T$ , we can represent the model in the form

$$\underbrace{\begin{bmatrix} \mathbf{y}_1^* \\ \mathbf{y}_2^* \\ \vdots \\ \mathbf{y}_T^* \end{bmatrix}}_{\mathbf{Y}^*} = \underbrace{\begin{bmatrix} \mathbf{x}_1 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{x}_2 & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{x}_T \end{bmatrix}}_{\mathbf{X}^*} \underbrace{\begin{bmatrix} \Phi' \\ \Phi' \\ \vdots \\ \Phi' \end{bmatrix}}_{\Phi^*} + \underbrace{\begin{bmatrix} \mathbf{f}_1 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{f}_2 & \mathbf{f}_2 & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{f}_T & \mathbf{f}_T & \cdots & \mathbf{f}_T \end{bmatrix}}_{\mathbf{F}^*} \underbrace{\begin{bmatrix} \Gamma' \\ \Delta\mathbf{\Lambda}'_1 \\ \Delta\mathbf{\Lambda}'_2 \\ \vdots \\ \Delta\mathbf{\Lambda}'_T \end{bmatrix}}_{\Theta} + \underbrace{\begin{bmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \\ \vdots \\ \mathbf{v}_T \end{bmatrix}}_{\mathbf{v}}, \quad (\text{A.12})$$

where  $\mathbf{v}_t = \mathbf{W}_t^{*1/2} \boldsymbol{\eta}_t$  and  $\Theta$  collects all time-invariant and time-varying loading parameters. Equation (A.12) is simply a linear regression in which  $\Phi^*$  (which appears repeated for each  $t$ ) and  $\Theta$  are the unknown coefficients, and  $\mathbf{X}^*$  and  $\mathbf{F}^*$  are known matrices constructed from the data and latent factors. Importantly, *this stacked system is observationally equivalent to the state-space formulation (A.10)–(A.11), but now all time-varying loadings enter linearly.*

We can therefore apply the same conjugate Bayesian steps developed for the constant-loading model to jointly sample  $\mathbf{\Gamma}$  and  $\{\Delta\mathbf{\Lambda}_t\}_{t=1}^T$  as if they were “static” regression coefficients.

In practice, we augment the Gibbs sampler of the previous subsection with an expanded Step 4 that draws the entire trajectory of factor loadings. More specifically, we replace step 4 by:

**4'. Sample  $\mathbf{\Gamma}$  and  $\Delta\mathbf{\Lambda}_{1:T}$ .** Conditional on  $\{\Phi, \mathbf{f}_{1:T}, \mathbf{W}_{1:T}^*\}$ , we sample the initial loading matrix  $\mathbf{\Gamma}$  (for instruments) and all loading increments  $\Delta\mathbf{\Lambda}_t$  (for macro variables,  $t = 1, \dots, T$ ) in one block. We assign each  $\Delta\mathbf{\Lambda}_t$  the prior  $\mathcal{N}(\mathbf{0}, \mathbf{Q})$  (as in (A.10)), and  $\mathbf{\Gamma}$  retains the same Gaussian/truncated Gaussian prior as before (see (A.4)). Given the linear stacked model (A.12), the joint conditional posterior of  $[\mathbf{\Gamma}, \Delta\mathbf{\Lambda}_1, \dots, \Delta\mathbf{\Lambda}_T]$  is Gaussian (subject to truncation for any restricted elements). Its mean and variance can be derived exactly as in the constant-loading case, treating  $\mathbf{F}^*$  as the “regressor matrix” and  $\mathbf{Q}$  as the prior covariance for the coefficients  $\Delta\mathbf{\Lambda}_t$ . Rather than write out the enormous formula, we note that the computation can be performed efficiently by exploiting the block structure in (A.12). In particular, for each equation  $i$  (each row of the loading matrices), the conditional posterior of  $[\Gamma_{i,\cdot}, \Delta\Lambda_{i,1}, \dots, \Delta\Lambda_{i,T}]$  is a multivariate normal with covariance

$$\overline{\mathbf{H}}_i = \left( \underline{\mathbf{H}}_i^{-1} + \sum_{t=1}^T \sigma_{i,t}^{-2} \mathbf{f}_t' \mathbf{f}_t \right)^{-1},$$

and mean  $\overline{\mathbf{H}}_i \left( \sum_{t=1}^T \sigma_{i,t}^{-2} \mathbf{f}_t' \hat{y}_{i,t} \right)$ , where  $\hat{y}_{i,t} = y_{i,t}^* - \phi_i \mathbf{x}_t$  is the current regression residual (and  $\underline{\mathbf{H}}_i$  is the prior covariance of  $[\Gamma_{i,\cdot}, \Delta\Lambda_{i,1}, \dots, \Delta\Lambda_{i,T}]$ , block-diagonal with entries  $\underline{h}_{ij}$  for  $\Gamma_{ij}$  and the corresponding elements of  $\mathbf{Q}$  for  $\Delta\Lambda_{i,t,j}$ ). This has the same form as in the static case, except that it pools information from all time periods  $1 \dots T$  when computing the sum of squares  $\sum_t \mathbf{f}_t' \mathbf{f}_t$  and cross-product  $\sum_t \mathbf{f}_t' \hat{y}_{i,t}$ . We then sample each element of  $\mathbf{\Gamma}$  and  $\Delta\mathbf{\Lambda}_t$  sequentially from this joint normal, using univariate truncation for any sign-restricted entries (exactly as described in step 4 above). Note that if a particular factor

loading is required to remain constant over time (for example, to enforce identification via a key instrument or aggregate—see discussion in Section 4.1 of the main text), one can simply impose  $\Delta\Lambda_{i,t,j} = 0$  for all  $t$  on that element (a zero restriction on every increment). In this case, the loading for series  $i$  on factor  $j$  stays equal to its initial value  $\Gamma_{ij}^*$  in all periods.

Steps 1, 2, 3, 5, and 6 of the Gibbs sampler are modified only in minor ways for the time-varying specification. Step 1 (sampling  $\Phi$ ) and step 5 (sampling  $\mathbf{f}_t$ ) are implemented exactly as in the constant case, except that  $\Lambda_t$  (the current loading matrix for macros) is used when computing the residuals  $\tilde{y}_{i,t}$  and  $\hat{y}_{i,t}$  at time  $t$ . Step 2 and 3 (Horseshoe updates for  $\psi_{i,j}$  and  $\tau_i$ ) are unchanged. Step 6 now involves sampling the entire path of each stochastic volatility  $\sigma_{i,t}^2$ . Given draws of all other parameters, the observation equation for series  $i$  can be written as  $\varepsilon_{i,t} = \sigma_{i,t}\eta_{i,t}$ , where  $\varepsilon_{i,t} = y_{i,t}^* - \phi_i\mathbf{x}_t - \Lambda_{i,t}\mathbf{f}_t$  is the residual and  $\eta_{i,t} \sim N(0, 1)$ . The log-volatility state equation is (A.11). We sample  $\{\log(\sigma_{i,t}^2)\}_{t=1}^T$  using simulation smoothing methods for non-linear state-space models. In particular, we employ the auxiliary mixture sampler of Kim et al. (1998) to approximate the  $\chi_1^2$  likelihood of  $\eta_{i,t}^2$  as a mixture of normals, and then draw the log-variance path via forward-filtering backward-sampling (FFBS). This yields draws from

$$\log(\sigma_{i,1:T}^2) \mid \{\varepsilon_{i,t}\}_{t=1}^T, \omega_i^2$$

for each  $i = 1, \dots, n$ . (In our implementation we fix the innovation variances  $\omega_i^2$  at modest values, e.g. 0.02, to restrict volatility to evolve slowly; one could instead place hyper-priors on the  $\omega_i^2$  and sample them in an outer Gibbs loop if desired.) The end result of step 6' is a draw of each time-varying idiosyncratic variance  $\sigma_{i,t}^2$  for  $t = 1, \dots, T$ .

Putting everything together, the extended Gibbs sampler alternates between: (i) sampling the VAR coefficients  $\Phi$  (step 1), Horseshoe parameters (steps 2–3), latent factors  $\mathbf{f}_t$  (step 5), and then (ii) sampling the time-varying loading matrices (step 4') and stochastic volatilities (step 6'). By using the transformed-model approach in (A.12), we are able to

draw the entire trajectory of  $\mathbf{\Lambda}_t$  in one block using simple regression formulas. This approach is computationally efficient even in large dimensions, and it automatically incorporates shrinkage on the loading drift: for example, one can treat the prior covariance  $\mathbf{Q}$  in (A.10) as unknown and assign it a Horseshoe prior just as we did for  $\mathbf{\Phi}$  (shrinking the elements of  $\Delta\mathbf{\Lambda}_t$  toward zero, which effectively shrinks  $\mathbf{\Lambda}_t$  toward constant trajectories). In our implementation we follow this idea by factorizing  $\mathbf{Q} = \text{diag}(q_1, \dots, q_{nr})$  (assuming independent random walk increments for each loading) and placing a hierarchical Horseshoe prior on each  $q_\ell$  (akin to (A.1)–(A.3)). The conditional posterior draw of  $\Delta\mathbf{\Lambda}_t$  in step 4' is then obtained by substituting the updated  $\mathbf{Q}$  into  $\underline{\mathbf{H}}_i$ . We refer the reader to Korobilis (2021) for further technical details on this shrinkage implementation. Conceptually, the key point is that—apart from the additional state-space sampling for  $\log \sigma_{i,t}^2$ —our algorithm for the time-varying model uses the *same* conditional posteriors derived for the constant model, simply applied to an augmented parameter vector. This allows the large-scale VAR with time-varying impact matrix and stochastic volatility to be estimated with very little extra computational cost, while flexibly incorporating structural identification via sign and zero restrictions as described above.

## B Additional empirical results

### B.1 Robustness in the small VAR

Section B of the Online Supplement assembles a battery of sensitivity exercises that all point to the same conclusion: the contractionary Target shock that we identify remains economically interpretable and quantitatively stable across a wide range of modelling and identification choices. Below I organise the evidence into four thematic groups, mirroring the order of the figures and tables in the supplement.

#### Alternative sign restriction sets

- [Figure B1](#) and [Table B2](#) replicate the nine variable VAR but impose sign restrictions on the *Path* shock that comply with the empirical evidence in [Miranda-Agrippino and Ricco \(2021\)](#).
- Despite this change, the impulse responses to the *Target* shock are virtually identical to the baseline: short term rates rise, output and inflation fall, and monetary/financial aggregates contract.
- [Figure B10](#) shows impulse responses under the benchmark sign restrictions, where we allow *Target* and *Path* factors to load on each other. The new set of sign restrictions is given in [Table B4](#). Note that the shocks (columns of  $\mathbf{f}_t$ ) remain uncorrelated, as these are structural.

*Take away:* the behaviour of the *Target* shock is orthogonal to the specific sign pattern chosen for the *Path* shock.

**Data definition changes for core macro variables** [Figure B2–B5](#) replace one baseline series at a time in the small VAR. The full set of changes is summarised in the following table:

Table B1: Variable replacements used in robustness exercises

Figure	Variable replaced	Robustness series
<a href="#">Figure B2</a>	Real activity	Brave–Butters–Kelly GDP nowcast
<a href="#">Figure B3</a>	Real activity	Industrial production (IP)
<a href="#">Figure B4</a>	Inflation	Core CPI
<a href="#">Figure B5</a>	Policy rate	Effective federal funds rate

Across all four swaps, the qualitative and quantitative responses of key output and price series to the *Target* shock barely move.

*Take away:* results are not an artefact of a particular GDP, inflation or policy rate proxy.

## Model structure tweaks in the small VAR



- [Figure B6](#) keeps only the *Target* and *Path* shocks; [Figure B7](#) adds demand and supply shocks whose sign patterns are given in [Table B3](#).
- [Figure B8](#) imposes  $\Phi_{ym} = 0$  (surprises affect macro variables only contemporaneously); [Figure B9](#) sets  $\Phi_{mm} = \Phi_{my} = \Phi_{ym} = 0$ , fully decoupling the surprise block from the macro block in lags.

In every case the *Target* shock IRFs remain well-behaved, with only minuscule variations in magnitude.

*Take away:* neither the number of other structural shocks nor the degree of feedback allowed between surprises and macro series undermines identification.

### External instrument and lag length robustness

- [Figure B11](#) re-estimates the VAR with the high-frequency surprises of [Jarociński and Karadi \(2020\)](#). The second shock becomes a *central bank information* shock (sign patterns in [Table B5](#)). The *Target* shock IRFs remain virtually unchanged.
- [Figure B12](#) repeats the baseline with  $p = 12$  monthly lags. Dynamics are less smoother but the peak effects, timing and credible sets overlap those in the six lag benchmark.

*Take away:* results are insensitive to the specific high-frequency surprise measure and to doubling the lag length.

### Overall message

**Every perturbation implemented, whether in identification, measurement, lag structure or model size, leaves the qualitative transmission of the *Target* monetary policy shock intact.** Output and prices fall, policy and short-term rates rise, money and equity prices contract, and these patterns are reasonably stable in magnitude and timing.

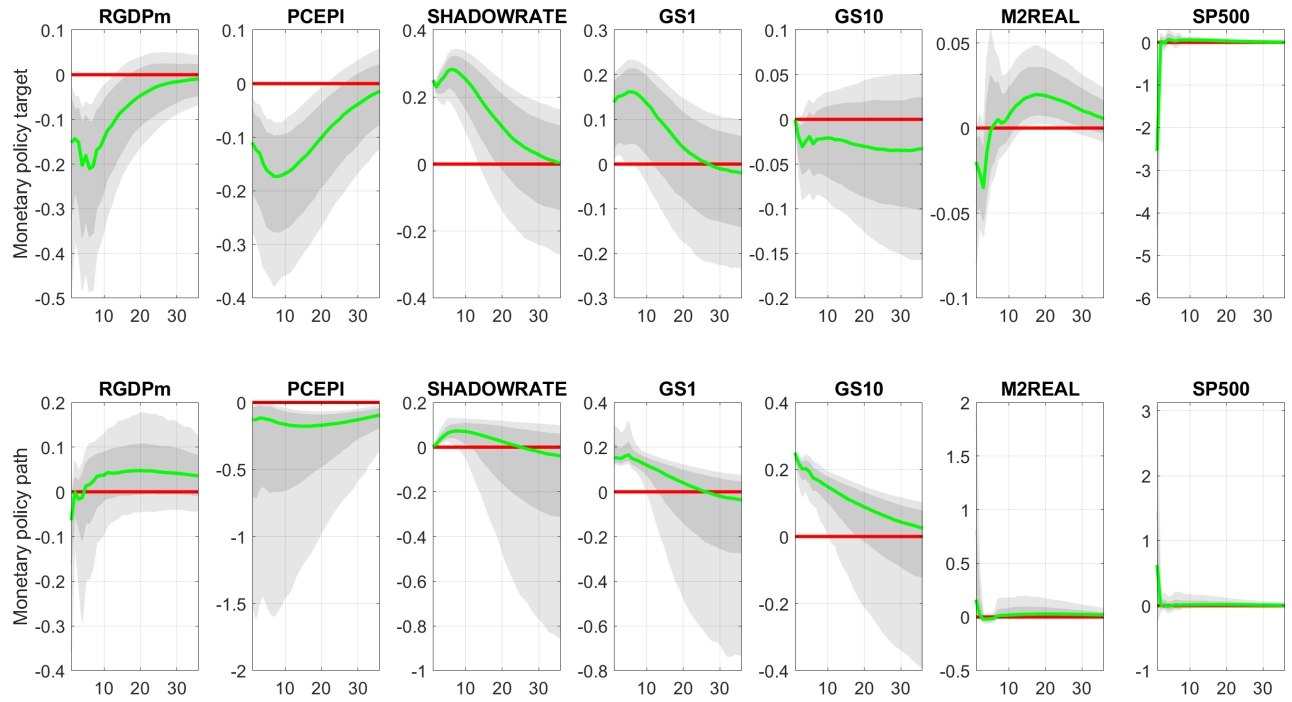


Figure B1: This is the small VAR with two instruments and seven core macro variables and sign restrictions of [subsection 4.1](#), but with a different set of sign restrictions on the Path shock. These restrictions comply with the empirical results of [Miranda-Agrippino and Ricco \(2021\)](#) and are given in [Table B2](#)

Endogenous vars	Identified shocks	
	Target shock	Path shock
<i>Target</i>	+	0
<i>Path</i>	0	+
-----	-----	-----
<i>realGDP</i>	−	−1
<i>PCEinf</i>	−	−1
<i>FFR</i>	+	0
<i>GS1</i>	+	<i>NA</i>
<i>GS10</i>	0	+
<i>M2REAL</i>	−	<i>NA</i>
<i>SP500</i>	−	<i>NA</i>

Table B2: Sign restrictions imposed to core macro variables, which correspond to sign restrictions to the matrix  $\mathbf{\Lambda}$ . Matrix  $\mathbf{\Lambda}$  here is constant, not time-varying.

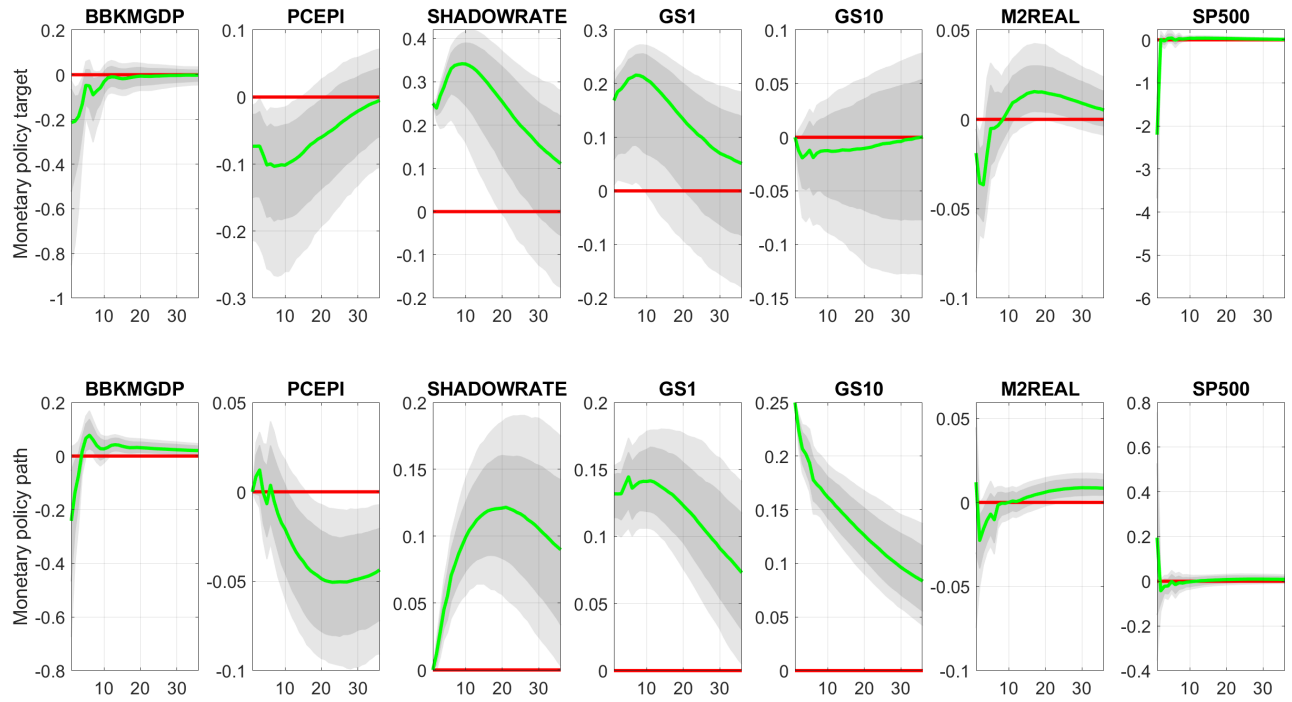


Figure B2: This is the small VAR with two instruments and seven core macro variables and sign restrictions of [subsection 4.1](#) but with real GDP replaced with the monthly BBKM GDP index.

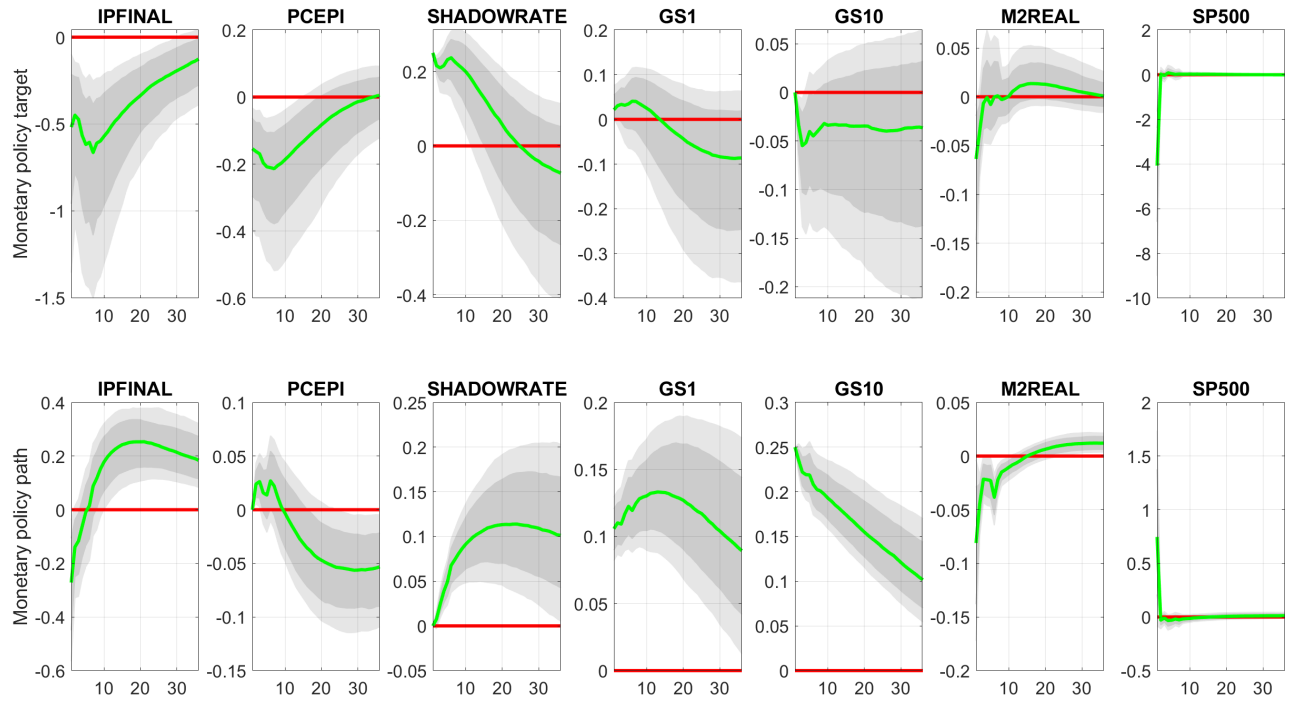


Figure B3: This is the small VAR with two instruments and seven core macro variables and sign restrictions of [subsection 4.1](#) but with real GDP replaced with total industrial production (IPFINAL).

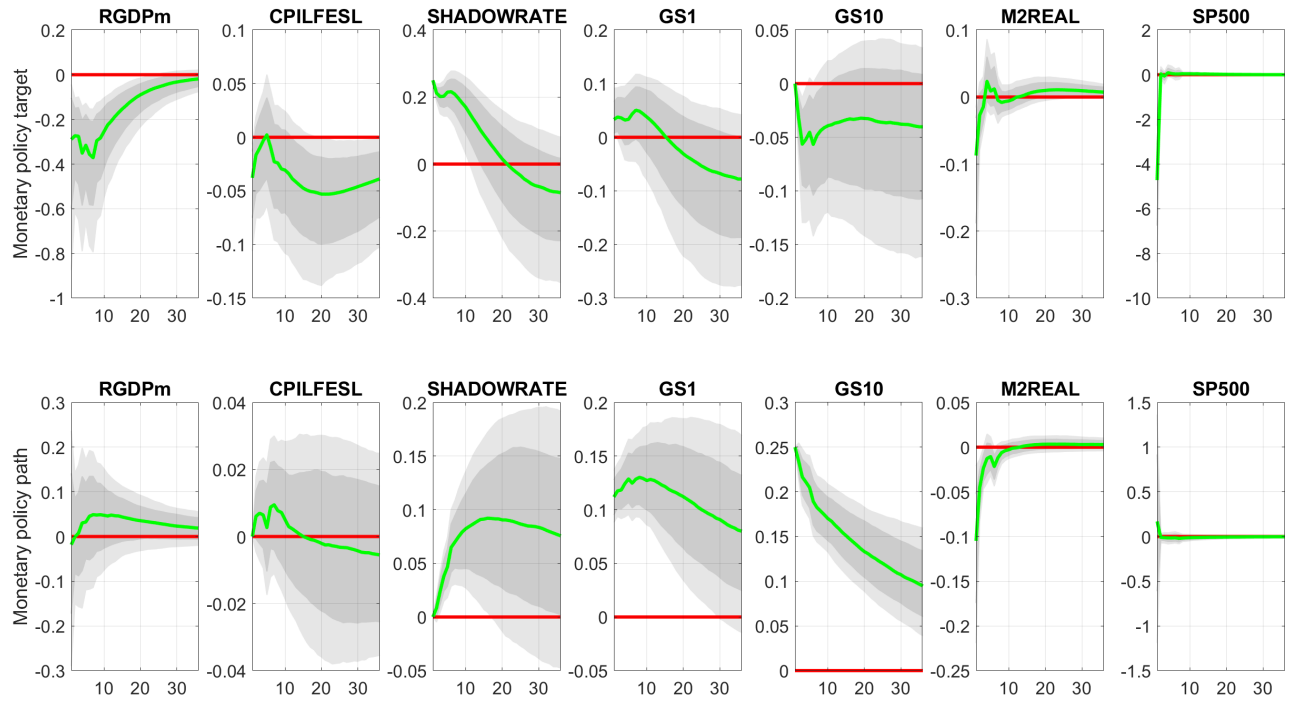


Figure B4: This is the small VAR with two instruments and seven core macro variables and sign restrictions of [subsection 4.1](#) but with PCE inflation replaced with core CPI (CPILFESL)

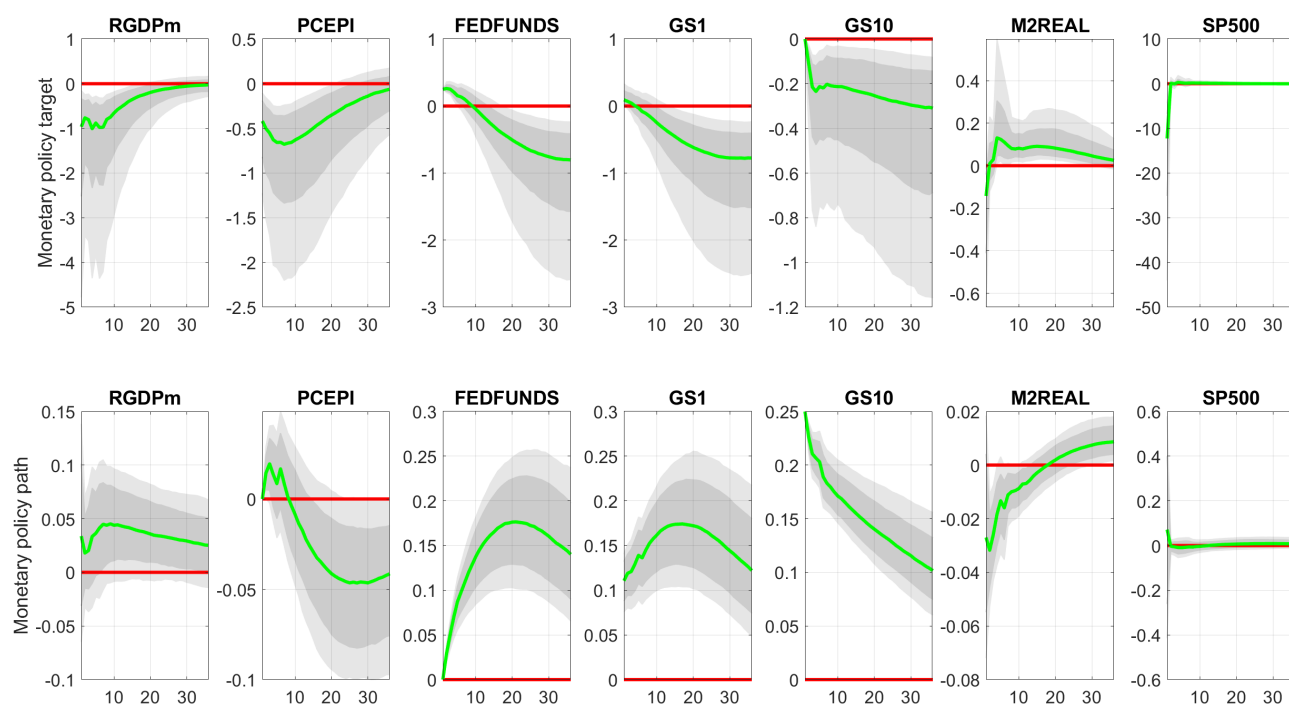


Figure B5: This is the small VAR with two instruments and seven core macro variables and sign restrictions of [subsection 4.1](#) but with the shadow rate replaced by the effective Fed funds rate (FEDFUNDS)

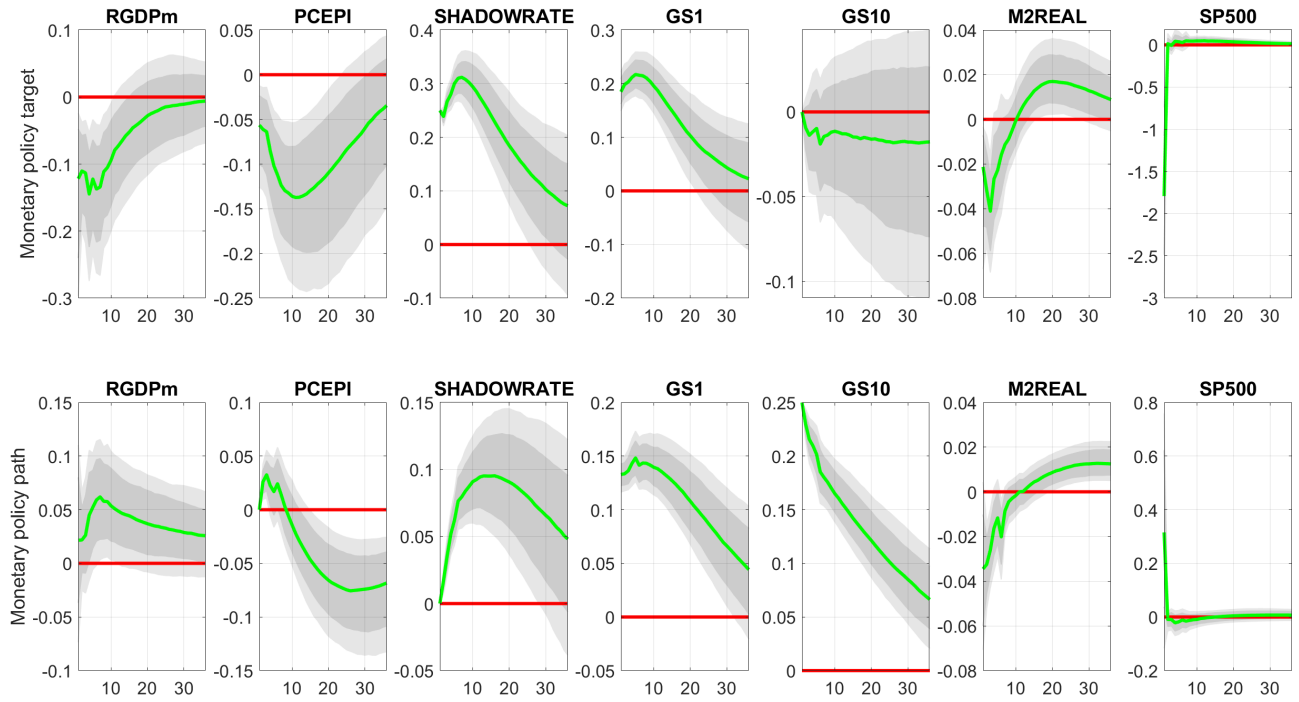


Figure B6: This is the small VAR with two instruments and seven core macro variables and sign restrictions of [subsection 4.1](#) but with only two shocks estimated (Target and Path) without any additional residual shocks.



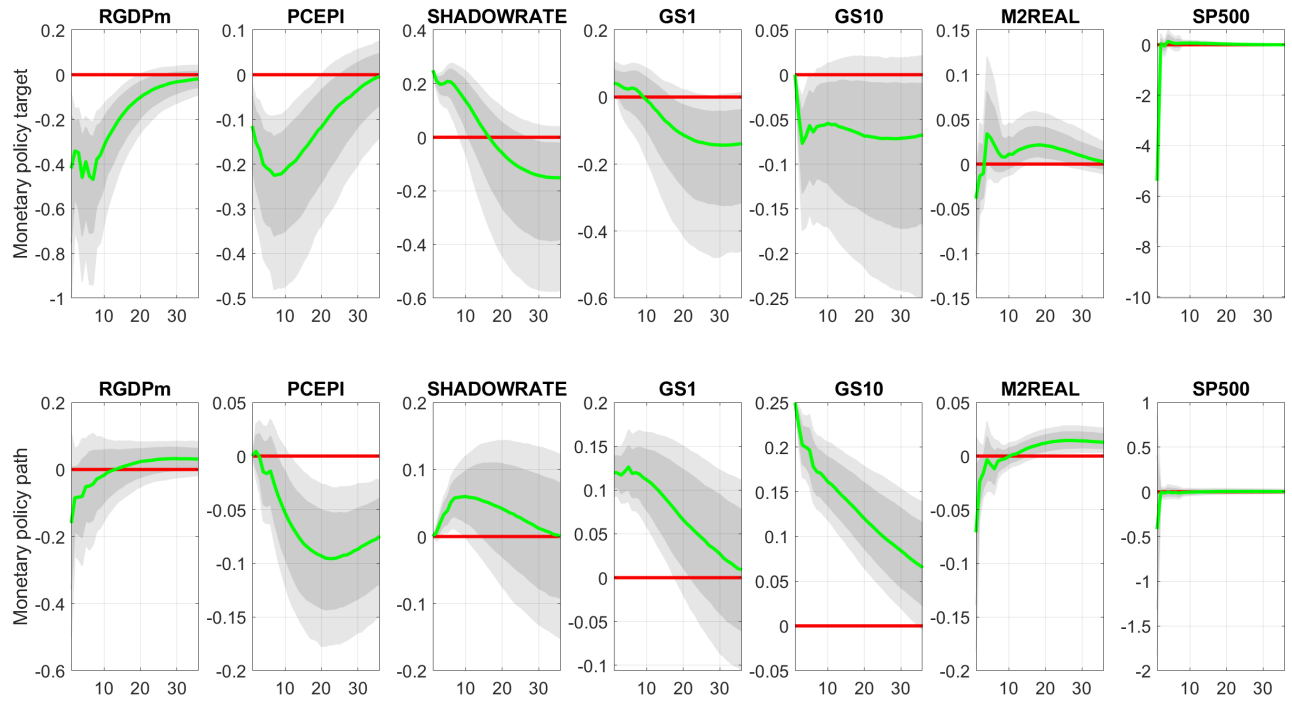


Figure B7: This is the small VAR with two instruments and seven core macro variables and sign restrictions of [subsection 4.1](#) with the Target and Path shocks augmented with a demand and supply shocks. The full set of sign restrictions used to identify all four shocks is given in [Table B3](#).

Endogenous vars	Identified shocks			
	Target shock	Path shock	Demand shock	Supply shock
<i>Target</i>	+	0	0	0
<i>Path</i>	0	+	0	0
-----	-----	-----	-----	-----
<i>realGDP</i>	−	−1	1	−1
<i>PCEinf</i>	−	−1	1	1
<i>FFR</i>	+	0	1	<i>NA</i>
<i>GS1</i>	+	<i>NA</i>	1	<i>NA</i>
<i>GS10</i>	0	+	<i>NA</i>	<i>NA</i>
<i>M2REAL</i>	−	<i>NA</i>	<i>NA</i>	<i>NA</i>
<i>SP500</i>	−	<i>NA</i>	1	−1

Table B3: Sign restrictions imposed to core macro variables, which correspond to sign restrictions to the matrix  $\mathbf{\Lambda}$ . This scheme identifies two monetary shocks and supply and demand shocks using sign restrictions.

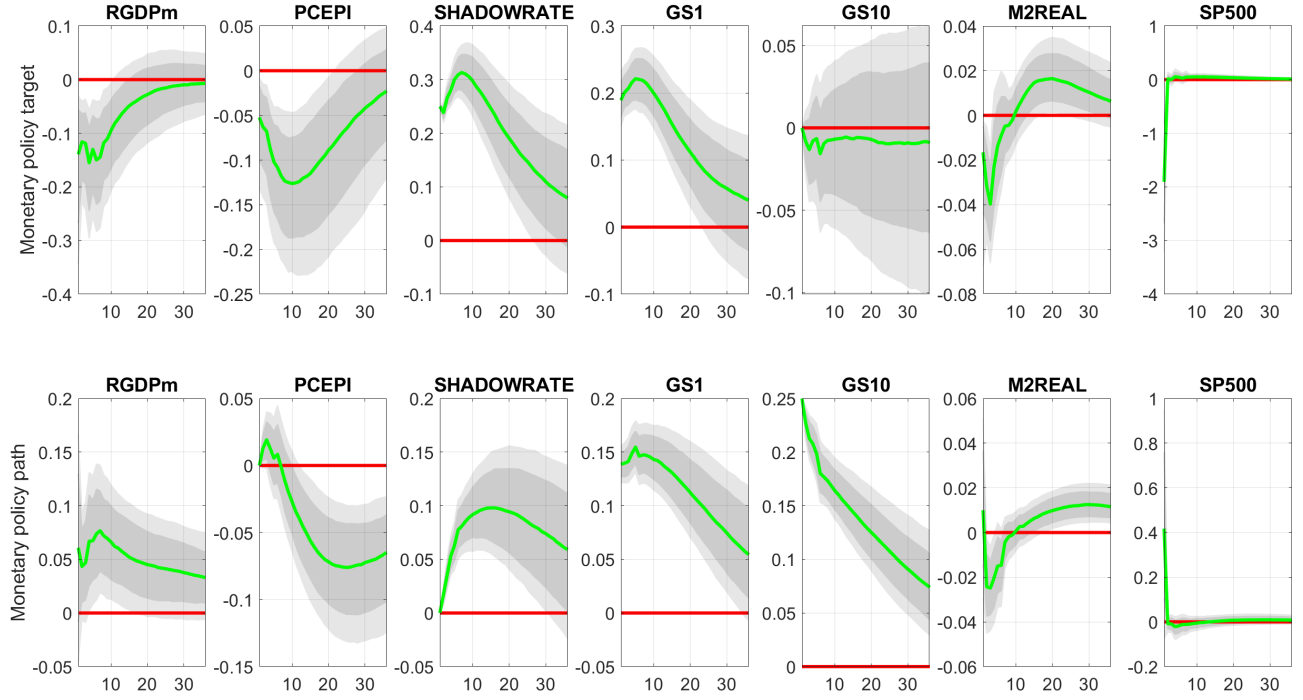


Figure B8: This is the small VAR with two instruments and seven core macro variables and sign restrictions of [subsection 4.1](#) with the additional parametric restriction that  $\Phi_{ym} = \mathbf{0}$ .

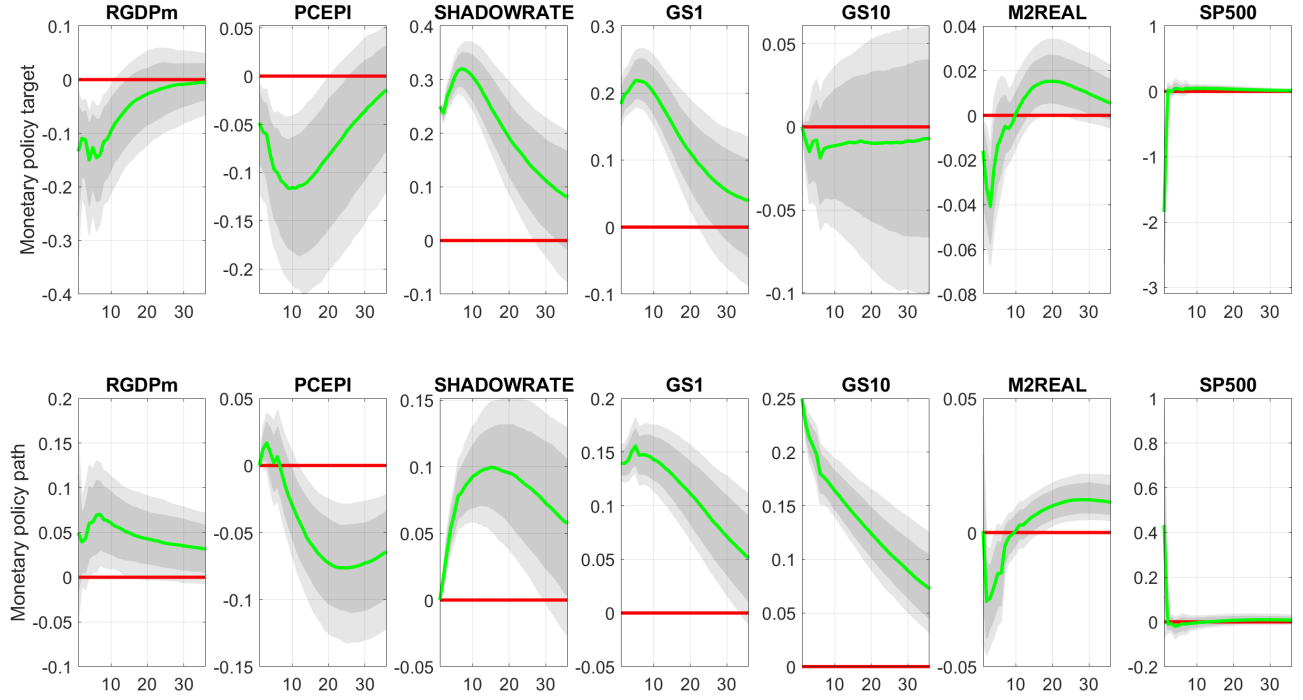


Figure B9: This is the small VAR with two instruments and seven core macro variables and sign restrictions of [subsection 4.1](#) with the additional parametric restriction that  $\Phi_{mm} = \Phi_{my} = \Phi_{ym} = \mathbf{0}$ .

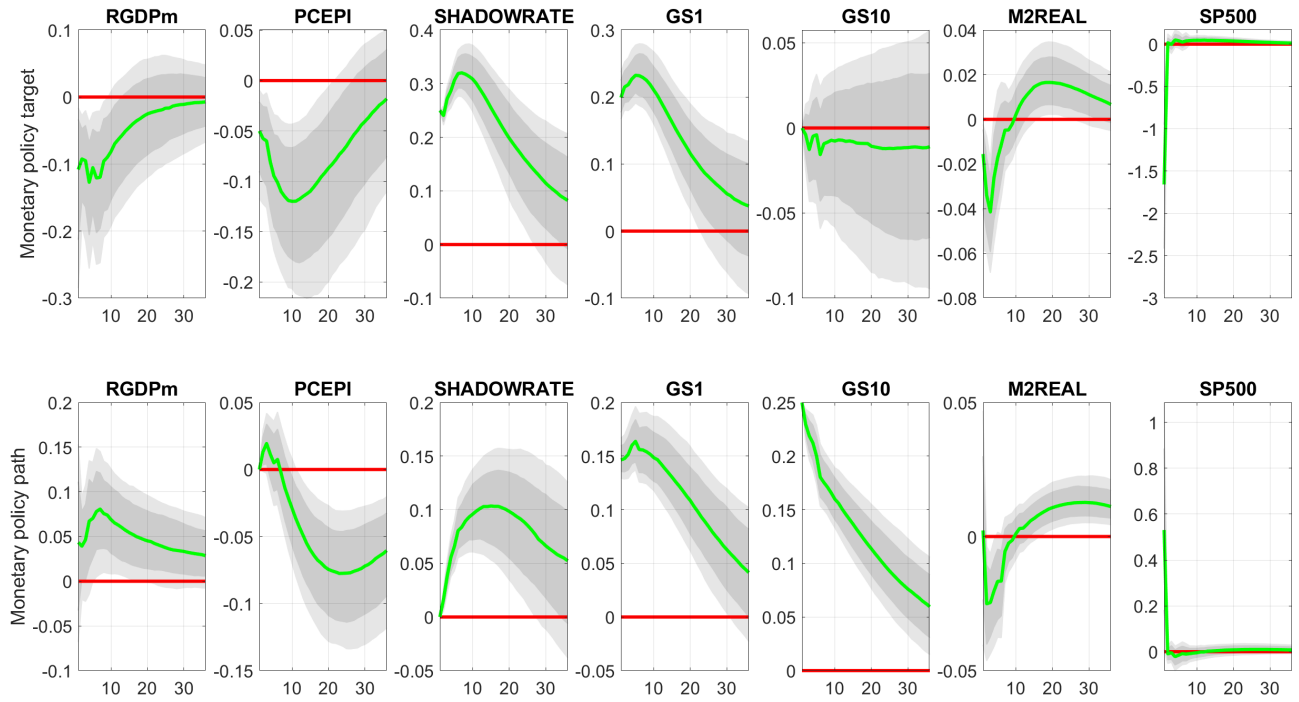


Figure B10: This is the small VAR with two instruments and seven core macro variables and sign restrictions of [subsection 4.1](#) but allowing the target factor to load on the path shock and vice versa for the path factor. This scheme is shown in [Table B4](#).

Endogenous vars	Identified shocks		Residual shocks	
	Target shock	Path shock	Residual shock 1	Residual shock 2
<i>Target</i>	+	<i>NA</i>	0	0
<i>Path</i>	<i>NA</i>	+	0	0
-----	-----	-----	-----	-----
<i>realGDP</i>	−	<i>NA</i>	<i>NA</i>	<i>NA</i>
<i>PCEinf</i>	−	0	<i>NA</i>	<i>NA</i>
<i>FFR</i>	+	0	<i>NA</i>	<i>NA</i>
<i>GS1</i>	+	<i>NA</i>	<i>NA</i>	<i>NA</i>
<i>GS10</i>	0	+	<i>NA</i>	<i>NA</i>
<i>M2REAL</i>	−	<i>NA</i>	<i>NA</i>	<i>NA</i>
<i>SP500</i>	−	<i>NA</i>	<i>NA</i>	<i>NA</i>

Table B4: Sign restrictions imposed to core macro variables, which correspond to sign restrictions to the matrix  $\mathbf{\Lambda}$ . Matrix  $\mathbf{\Lambda}$  here is constant, not time-varying.

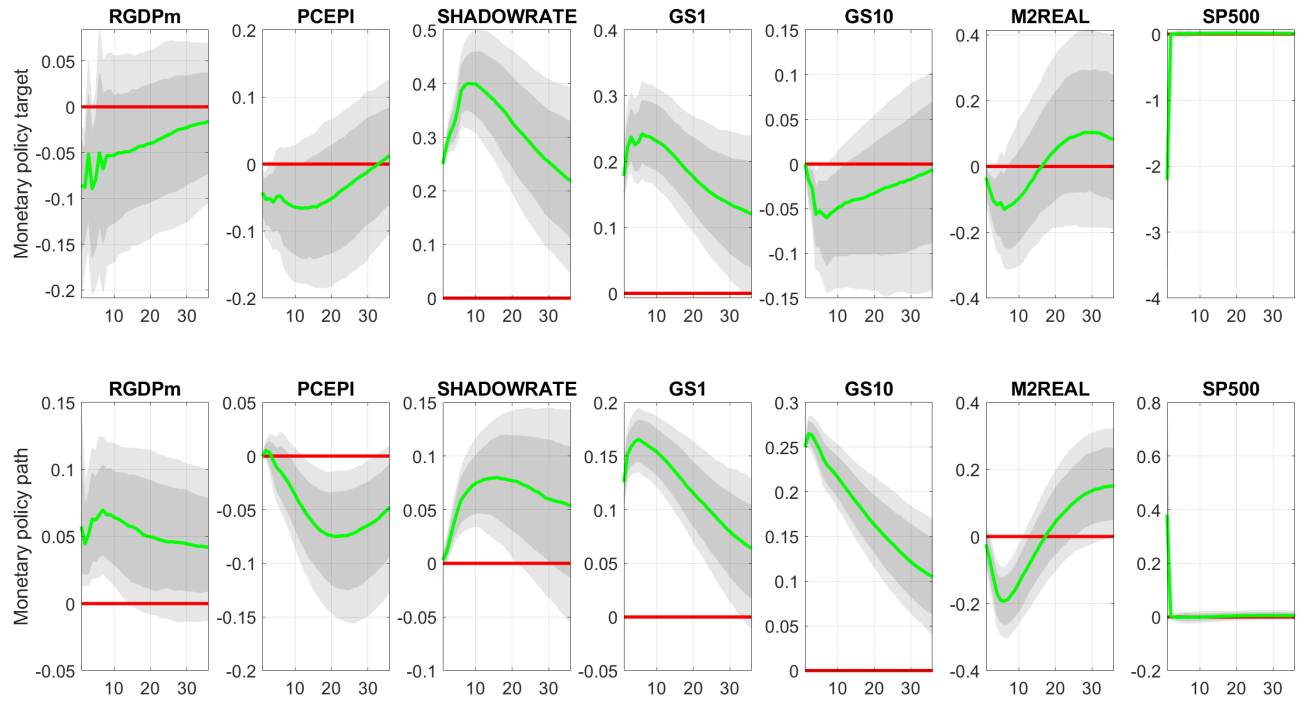


Figure B11: This is the small VAR with two instruments and seven core macro variables and sign restrictions of [subsection 4.1](#), but estimated using the surprise series in [Jarociński and Karadi \(2020\)](#). Now the second shock is interpreted as a “central bank information” shock, and I impose sign restrictions that are consistent with the findings of these authors; see [Table B5](#).

Endogenous vars	Identified shocks		Residual shocks	
	MP shock	CB info shock	Residual shock 1	Residual shock 2
<i>MP1</i>	+	0	0	0
<i>TFUT10</i>	+	+	0	0
<i>SP500</i>	−	+	0	0
-----	-----	-----	-----	-----
<i>realGDP</i>	−	+	<i>NA</i>	<i>NA</i>
<i>PCEinf</i>	−	0	<i>NA</i>	<i>NA</i>
<i>FFR</i>	+	0	<i>NA</i>	<i>NA</i>
<i>GS1</i>	+	+	<i>NA</i>	<i>NA</i>
<i>GS10</i>	0	+	<i>NA</i>	<i>NA</i>
<i>M2REAL</i>	−	<i>NA</i>	<i>NA</i>	<i>NA</i>
<i>SP500</i>	−	+	<i>NA</i>	<i>NA</i>

Table B5: Sign restrictions imposed to core macro variables, which correspond to sign restrictions to the matrix  $\mathbf{\Lambda}$ . Matrix  $\mathbf{\Lambda}$  here is constant, not time-varying.



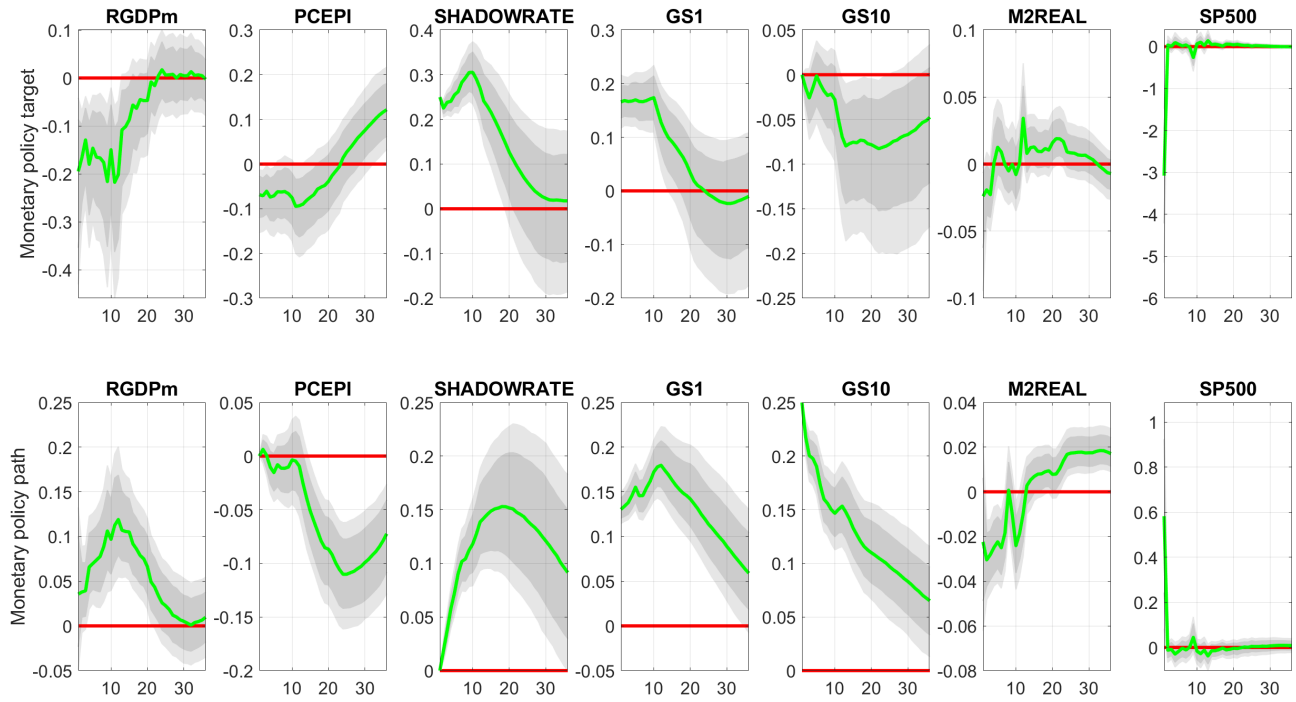


Figure B12: This is the small VAR with two instruments and seven core macro variables and sign restrictions of [subsection 4.1](#), but estimated with  $p = 12$  lags.

## B.2 Estimation of the large VAR with p=6 lags

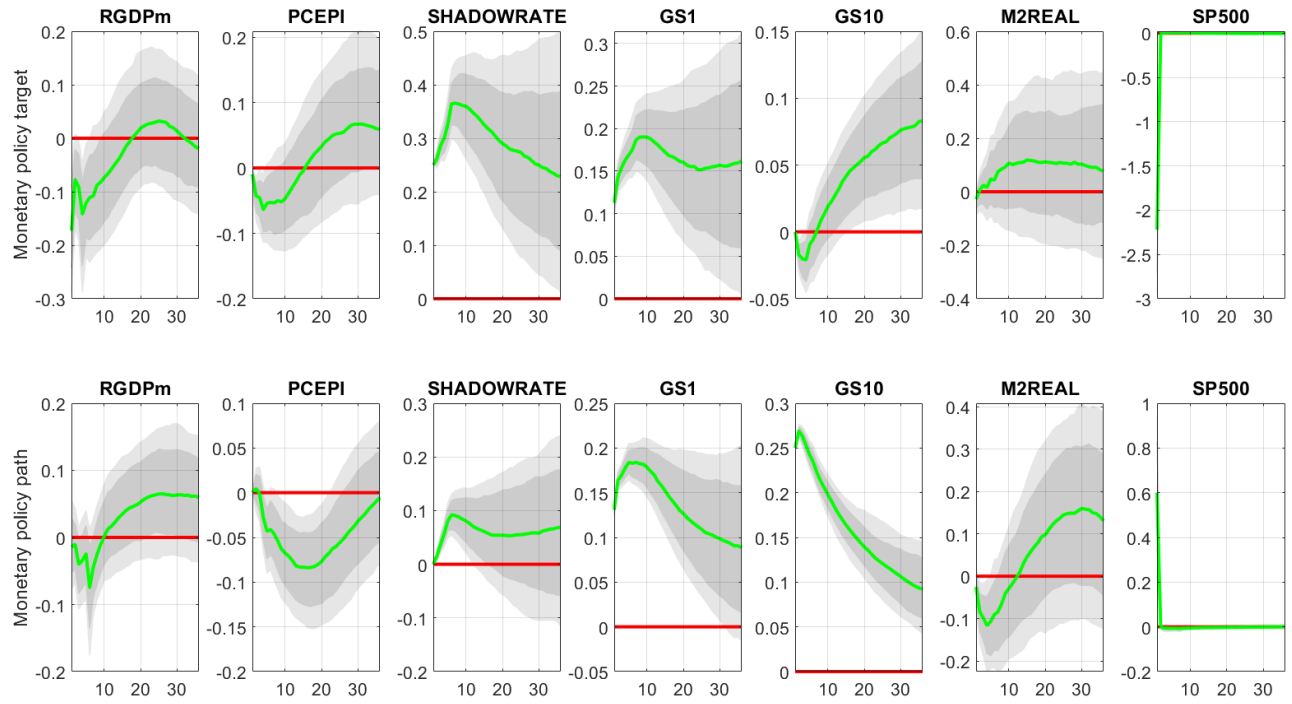


Figure B13: Responses of core macro variables to Target and Path shocks in the large VAR.  $\Lambda$  is constant for these variables, therefore, IRFs are fixed over the sample.

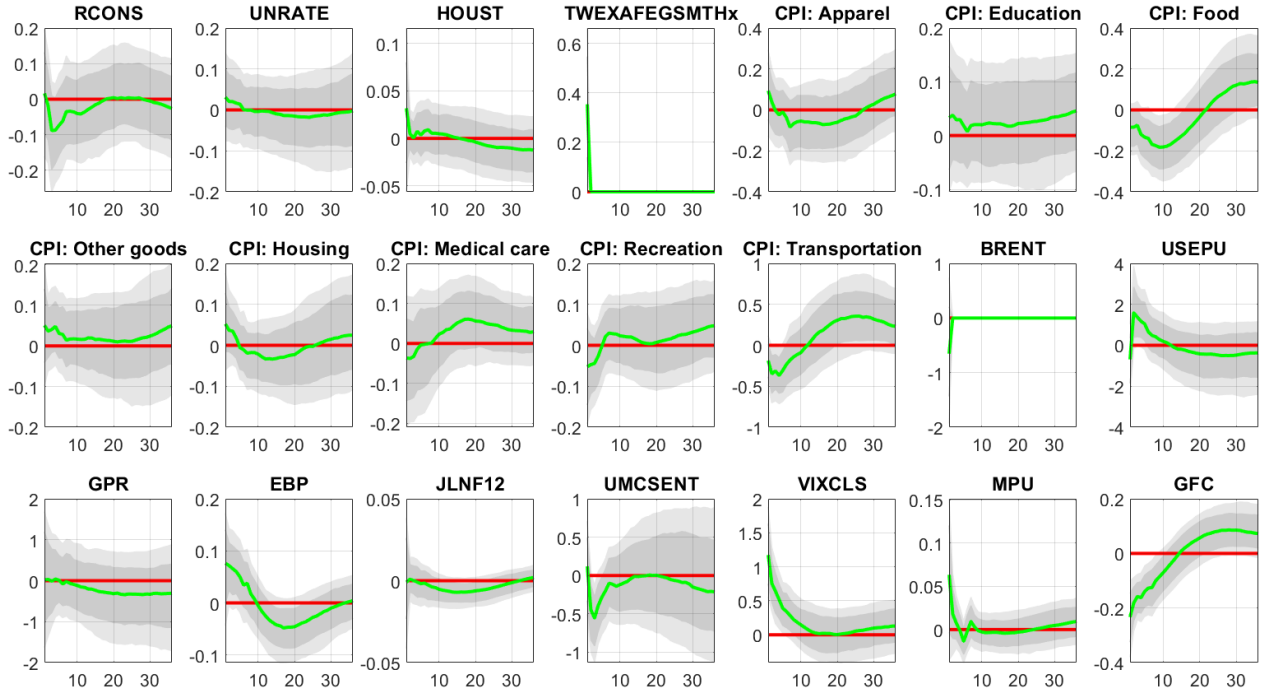


Figure B14: Responses of all other macro variables to Target shock in the large VAR.  $\Lambda$  is time-varying for these variables, therefore, plot gives responses at time  $T$ .

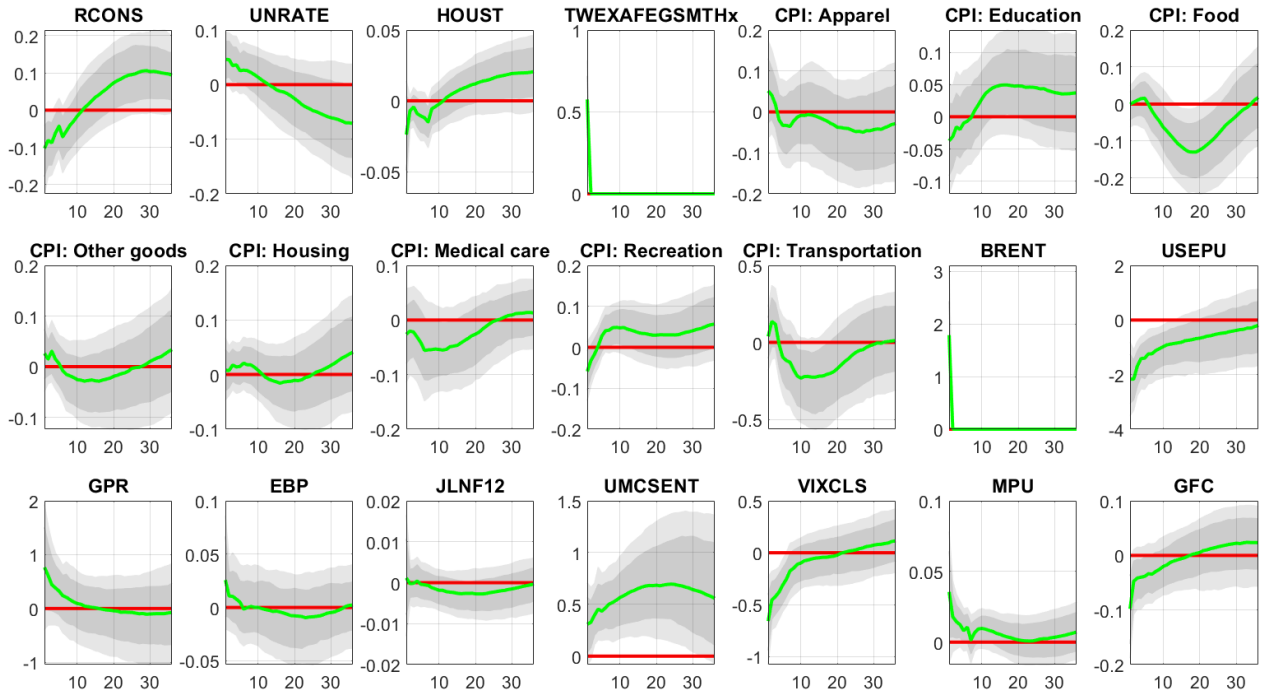


Figure B15: Responses of all other macro variables to Path shock in the large VAR.  $\Lambda$  is time-varying for these variables, therefore, plot gives responses at time  $T$ .

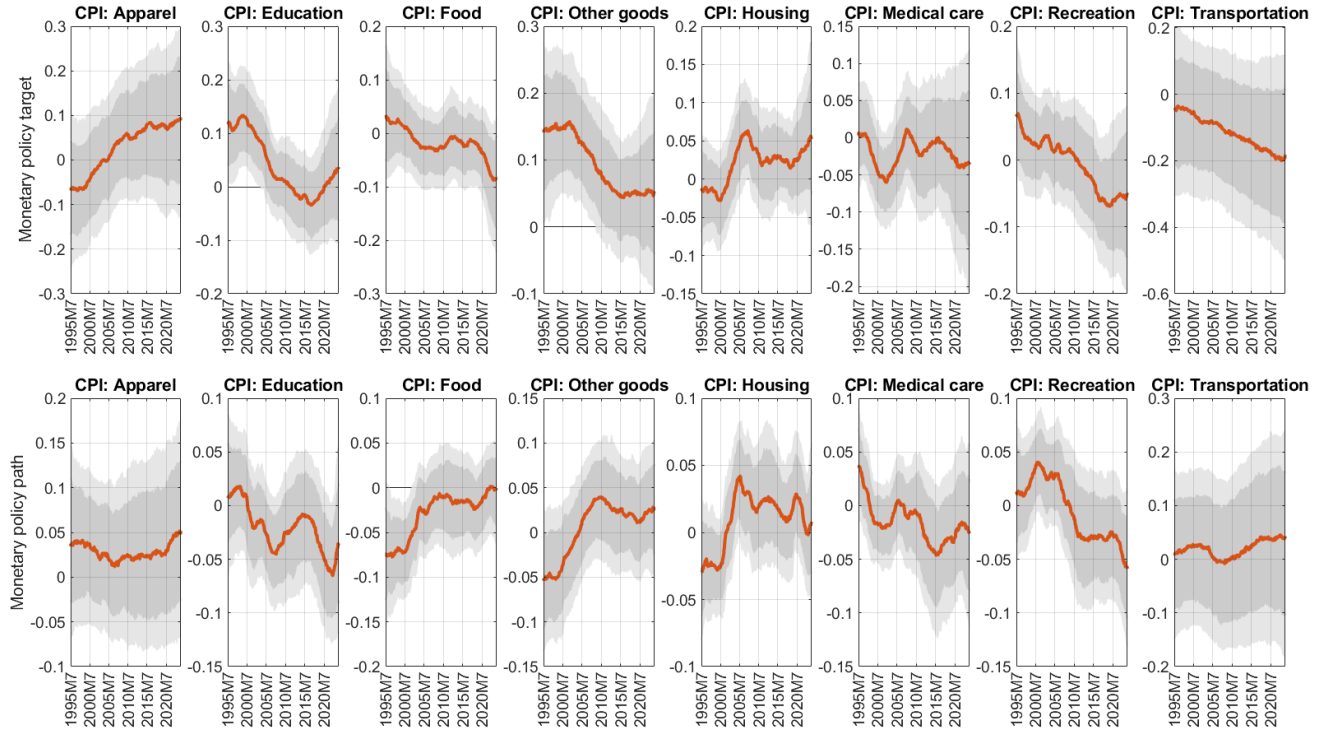


Figure B16: Impact responses of disaggregated CPI variables to Target and Path shocks.

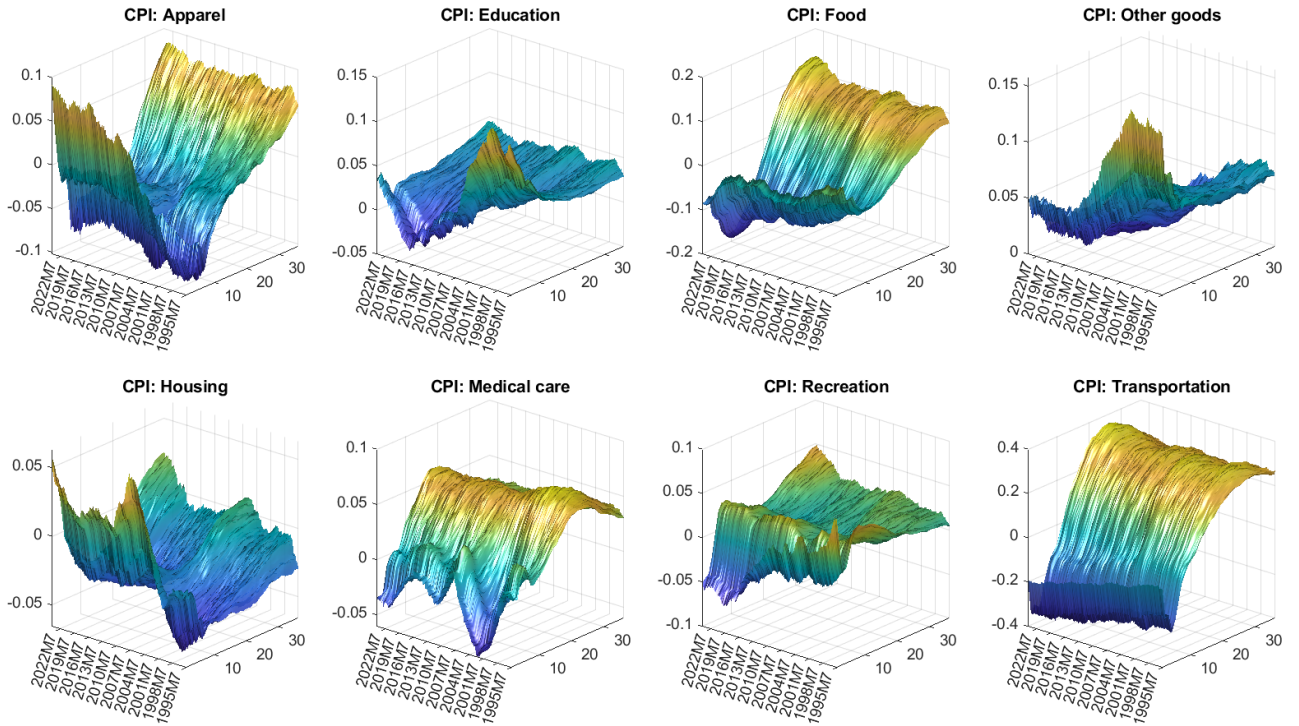


Figure B17: Responses over time and over horizons of disaggregated CPI variables to a Target shock.

## References

- Bhattacharya, A., Chakraborty, A., and Mallick, B. K. (2016). Fast sampling with Gaussian scale mixture priors in high-dimensional regression. *Biometrika*, 103(4):985–991.
- Botev, Z. I. (2017). The normal law under linear restrictions: Simulation and estimation via minimax tilting. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(1):125–148.
- Carvalho, C. M., Polson, N. G., and Scott, J. G. (2010). The horseshoe estimator for sparse signals. *Biometrika*, 97(2):465–480.
- Geweke, J. F. (1996). Bayesian inference for linear models subject to linear inequality constraints. In Lee, J. C., Johnson, W. O., and Zellner, A., editors, *Modelling and Prediction Honoring Seymour Geisser*, pages 248–263, New York, NY. Springer New York.
- Jarociński, M. and Karadi, P. (2020). Deconstructing monetary policy surprises—The role of information shocks. *American Economic Journal: Macroeconomics*, 12(2):1–43.
- Kim, S., Shephard, N., and Chib, S. (1998). Stochastic volatility: Likelihood inference and comparison with ARCH models. *The Review of Economic Studies*, 65(3):361–393.
- Korobilis, D. (2021). High-dimensional macroeconomic forecasting using message passing algorithms. *Journal of Business & Economic Statistics*, 39(2):493–504.
- Miranda-Agrippino, S. and Ricco, G. (2021). The transmission of monetary policy shocks. *American Economic Journal: Macroeconomics*, 13(3):74–107.
- Neal, R. M. (2003). Slice sampling. *The Annals of Statistics*, 31(3):705–767.