Restrictions on Conditional Correlations in SVARs*

Myunghyun Kim† Sunho Lee‡

May 24, 2026

Abstract

This paper introduces the *correlation restriction* as an identifying restriction in Bayesian structural vector autoregressions (SVARs). It is based on the key features of demand and supply shocks: The conditional correlation between output and prices on demand (supply) shocks is positive (negative). The correlation restriction ensures that the structural parameters generate a positive (negative) conditional correlation between the two variables on demand (supply) shocks. We use the restriction to identify monetary policy shocks (demand shocks) and oil supply shocks (supply shocks) and show that it is very useful in identifying the two shocks. In the case of monetary policy shocks, the correlation restriction narrows the credible sets for the impulse response functions (IRFs) and enables the price puzzle to disappear. Moreover, the correlation restriction helps to sharply identify oil supply shocks and to modify the shape of the price IRF, which makes the IRF consistent with economic theory.

JEL Classification: C32, E31, E52, Q43

Keywords: correlation restriction; SVARs; conditional correlations; demand shocks; supply shocks; monetary policy shocks; oil supply shocks; sign restrictions

* We would like to thank Dongho Choo, Xiyu Jiao, Kyu Ho Kang, Minchul Shin, Inwhan So, and the participants in the seminars at Sungkyunkwan University and Bank of Korea for their helpful comments. The findings, interpretations and conclusions expressed in this paper are entirely those of the authors and should not be attributed to the Bank of Korea. Kim acknowledges that this work was supported by the Ministry of Education of the Republic of Korea and the National Research Foundation of Korea (NRF-2022S1A3A2A01088457). All remaining errors are our own.

† Department of Economics, Sungkyunkwan University, 25-2, Sungkyunkwan-ro, Jongno-gu, Seoul, 03063, Korea (e-mail: mhkim7812@gmail.com).

‡ Office of Economic Modeling & Policy Analysis, Bank of Korea, 39, Namdaemun-ro, Jung-gu, Seoul, 04531, Korea (e-mail: econ.preference@gmail.com).

1 Introduction

The key features of demand and supply shocks are that a positive demand shock, such as an expansionary monetary policy shock, increases output and prices, whereas a supply shock, such as an oil supply shock reducing world oil production, decreases output and increases prices. In other words, output and prices are positively correlated with respect to demand shocks, while the conditional correlation between output and prices on supply shocks is negative. Although demand and supply shocks are easily distinguishable, and their key features are well understood in economic theory, these features are not typically imposed directly as identifying restrictions in Bayesian structural vector autoregressions (SVARs). Therefore, this paper incorporates the key features of demand and supply shocks into Bayesian SVARs through the *correlation restriction* and demonstrates its effectiveness in precisely identifying monetary policy shocks (demand shocks) and oil supply shocks (supply shocks).

The correlation restriction ensures that the structural parameters produce a positive (negative) conditional correlation between output and prices on demand (supply) shocks. Specifically, a correlation restriction for demand shocks rules out the structural parameters that generate negative conditional correlations between output and prices on demand shocks. In contrast, a correlation restriction for supply shocks rules out the structural parameters that produce positive conditional correlations between the two variables on supply shocks. The correlation restriction is highly flexible, and thus it can be easily applied to other restrictions. Despite its flexibility, our study focuses on combining the correlation restriction and sign restrictions on the structural parameters or the impulse response functions (IRFs) in SVARs.[1]

Since Uhlig (2005), it has become popular to use minimal, uncontroversial sign restrictions on the IRFs or the structural parameters when identifying structural shocks (e.g., monetary policy shocks and oil supply shocks), as they are generally thought to be weaker than classic restrictions, such as the recursive restrictions. Moreover, according to Rubio-Ramírez, Waggoner, and Zha (2010), since the structural parameters are set-identified under sign restrictions, they bring about robust conclusions across the set of SVARs satisfying the restrictions. However, sign restrictions often lead to overly agnostic results. In other words, they can result in very different implications for the IRFs from those predicted by economic theory. For example, Kilian and Murphy (2012) show that the values of the price elasticity of oil supply to demand shocks obtained from global oil market SVARs with only sign restrictions on the IRFs are disputable. Arias, Caldara, and Rubio-Ramírez (2019) also show that, in the SVAR with sign restrictions on the IRFs of Uhlig (2005), the posterior estimates

[1] It would be interesting to apply our correlation restriction to SVARs with other restrictions, such as recursive and long-run restrictions. However, that is beyond the scope of our paper, so we leave it for future research.

of the contemporaneous coefficients on the monetary policy equation are likely to violate the sign implied in the monetary rules embedded in dynamic stochastic general equilibrium (DSGE) models. Therefore, for sign-restricted SVARs, it is important to impose additional uncontroversial restrictions to reduce the set of the structural parameters that satisfy the sign restrictions. This is necessary in order to reach economically meaningful conclusions.

Our proposed correlation restriction helps identify structural shocks more precisely in SVARs with sign restrictions. Specifically, by adding an uncontroversial restriction on the conditional correlation between output and prices on demand or supply shocks based on economic theory, the correlation restriction narrows the set of the structural parameters that satisfy the sign restrictions, which enables the SVARs to lead to meaningful economic conclusions.

This paper begins by describing a basic SVAR framework in which the correlation restriction and sign restrictions are imposed. For the correlation restriction, simulation is used to compute the conditional correlations between output and prices on demand or supply shocks. Specifically, we simulate the SVAR model with respect to only demand (supply) shocks (with shutting down all other structural shocks) under each draw of the structural parameters to obtain the simulated series of output and prices. Then, we use the two simulated series to compute the simulated conditional correlation between output and price on demand (supply) shocks.

Next, we describe the functions that can characterize the correlation restriction and the sign restrictions on the IRFs and the structural parameters. These functions show that the correlation restriction depends solely on the structural parameters, like the sign restrictions. Therefore, we can use the Bayesian methods developed in Rubio-Ramírez, Waggoner, and Zha (2010) and Arias, Rubio-Ramírez, and Waggoner (2018) for the correlation restriction without further modification. We also explain the Bayesian inference and the algorithm for sampling the structural parameters under both the sign restrictions and the correlation restriction.

Moreover, one might think that imposing sign restrictions on the IRFs of output and prices simultaneously would lead to the same results as the correlation restriction. Therefore, we compare the correlation restriction with the sign restrictions. We demonstrate that this is not the case, and that these two restrictions differ significantly. The key point is that the sign restrictions are imposed at each horizon, whereas the correlation restriction is imposed on the overall conditional correlation. For example, to obtain the same results as the correlation restriction for demand shocks (i.e., a positive conditional correlation between output and prices on demand shocks), one may impose the sign restrictions that a positive demand shock increases output and prices over a sufficiently long period. However, the sign restrictions require pre-specifying when output and prices increase together, and in practice these periods are usually the first few periods. In contrast, the correlation restriction does not require pre-specifying when the two variables increase (or decrease) together. Under the correlation restriction, the conditional correlation can be temporarily negative, even during the first few periods, since the correlation restriction only constrains the structural parameters to ensure that the overall conditional correlation is positive. Furthermore, under the sign restrictions, the signs of the IRFs of output and prices are restricted (in this example, their signs must be positive). However, under the correlation restriction, the signs of the two IRFs are unrestricted. Their signs can either be positive or negative (even their signs can differ temporarily), as long as the overall conditional correlation between output and prices is positive.

Since monetary policy shocks are well-known demand shocks, we apply a correlation restriction for demand shocks to an SVAR for monetary policy shocks. As is well understood from economic theory, a contractionary monetary policy shock raises the nominal and real interest rates, which depresses aggregate demand. Accordingly, output and prices go down. In other words, conditional on monetary policy shocks, output and prices are positively correlated. Therefore, we impose a correlation restriction that the correlation between output and prices conditional on monetary policy shocks must be positive in an SVAR model for monetary policy shocks. For the SVAR model, we consider that of Arias, Caldara, and Rubio-Ramírez (2019), in which sign restrictions on the structural parameters are imposed and prices fall in response to a contractionary monetary policy shock (i.e., the price puzzle).[2] We show that the correlation restriction improves the identification of monetary policy shocks. More importantly, it enables the price puzzle to disappear (i.e., prices fall in response to a contractionary monetary policy shock).[3]

A good example of supply shocks is oil supply shocks, as the name of the shocks implies. Accordingly, we use a correlation restriction to identify oil supply shocks. As is well understood, an oil supply shock, which decreases world oil production, leads to an increase in the real oil price. Since oil is used in the production of almost all goods, the increased real oil price results in a rise in production costs, which depresses aggregate supply. Thus, output declines, but prices go up. That is, output and prices are negatively correlated with respect to oil supply shocks. For this application, the SVAR model for oil supply shocks by Baumeister and Peersman (2013), in which sign restrictions on the IRFs are used, is con-

[2] In fact, Arias, Caldara, and Rubio-Ramírez (2019) impose zero restrictions, as well as the sign restrictions. For simplicity, we consider their SVAR with only the sign restrictions. However, in Online Appendix B, we also consider their SVAR with both the sign and zero restrictions, as in their original SVAR. Note that even when both the sign and zero restrictions are imposed in their SVAR, the price puzzle is alive.

[3] Some might think that the correlation restriction eliminates the price puzzle by definition. This is not the case. The correlation restriction does not rule out the structural parameters that lead to an increase in output and prices in response to a contractionary monetary policy shock. Furthermore, as previously mentioned, under the correlation restriction, output and prices can react in opposite directions, even during the first few periods after the shock.

sidered.[4] Specifically, we add a correlation restriction—the conditional correlation between output and prices on oil supply shocks must be negative—to the SVAR model for oil supply shocks. Estimation results clearly show that the correlation restriction helps oil supply shocks be sharply identified. Furthermore, unlike the SVAR without the correlation restriction, in which prices do not rise significantly in response to an oil supply shock, the restriction allows prices to increase significantly in response to the shock, which is consistent with economic theory.

One related issue raised by Baumeister and Hamilton (2024) is that, in set-identified SVARs, conventional Bayesian posterior summaries may be inappropriate. In particular, a given reduced form can generate multiple admissible structural paths through different rotations, and these paths have the same likelihood. The conventional practice based on a uniform prior over the rotation matrix can place some of these paths inside a reported credible set and others outside it, even though the likelihood does not distinguish among them. Ranking such paths therefore requires either an informative prior over the rotation matrix or a robust object that does not condition on a particular prior over rotations. We follow Giacomini and Kitagawa (2021), which provides the latter approach, to examine the robustness of the identifying content of the correlation restriction. Specifically, we report the set of posterior means for the IRFs. This set gives the lower and upper bounds on posterior mean responses attainable by rotations that satisfy the identifying restrictions. We then compare the set under the sign restrictions and under the sign and correlation restrictions to assess whether the correlation restriction adds identifying content beyond the sign restrictions. The results show that the correlation restriction rules out a sizable part of the rotations admitted by the sign restrictions and narrows the set of posterior means by removing responses that are less consistent with the economic interpretations of monetary policy and oil supply shocks.

This paper is related to three strands of literature. The first one is studies on SVARs with sign restrictions on the IRFs or the structural parameters. Since Faust (1998), Canova and De Nicoló (2002), and Uhlig (2005), sign restrictions on the IRFs or the structural parameters in SVARs have been frequently used in many contexts of empirical studies, including not only monetary policy shocks and oil shocks, as in our study, but also government spending shocks, technology shocks, and financial shocks.[5]

The second strand is studies on the effects of monetary policy shocks using SVARs.

[4] Different from Baumeister and Peersman (2013), we do not allow the parameters to be time-varying because the focus of this paper is not the time-varying effects of oil supply shocks.

[5] For example, Faust (1998), Canova and De Nicoló (2002), Uhlig (2005), Amir-Ahmadi and Uhlig (2015), and Arias, Caldara, and Rubio-Ramírez (2019) for monetary policy shocks, Pappa (2009) and Mountford and Uhlig (2009) for government spending shocks, Dedola and Neri (2007) and Peersman and Straub (2009) for technology shocks, Kilian and Murphy (2012), Baumeister and Peersman (2013), Baumeister and Hamilton (2019), and Kim (2020) for oil shocks, and Hristov, Hülsewig, and Wollmershäuser (2012) and Gambetti and Musso (2017) for financial shocks.

For instance, Christiano, Eichenbaum, and Evans (1996) use classic zero restrictions (i.e., the recursive restrictions). Since their seminal work, various methods to identify monetary policy shocks, such as factor-augmented VARs (Bernanke, Boivin, and Eliasz, 2005), sign restrictions (Faust, 1998; Canova and De Nicoló, 2002; Uhlig, 2005; and Arias, Caldara, and Rubio-Ramírez, 2019), and high-frequency identification and proxy SVARs (Gertler and Karadi, 2015; and Jarociński and Karadi, 2020), etc., have been used in SVARs.

The last strand is studies on the effects of oil shocks using SVARs. Kilian (2009) estimates an SVAR with zero restrictions to disentangle oil supply, aggregate demand, and oil-specific demand shocks. Since his seminal work, researchers have used a variety of approaches to identify oil shocks within the SVAR framework, which includes sign restrictions (Kilian and Murphy, 2012; Lippi and Nobili, 2012; Peersman and Van Robays, 2012; Baumeister and Peersman, 2013; Baumeister and Hamilton, 2019; and Kim, 2020), institutional features of OPEC and high-frequency data (Känzig, 2021), and narrative information (Antolín-Díaz and Rubio-Ramírez, 2018; and Caldara, Cavallo, and Iacoviello, 2019).

Among the studies reviewed above, Canova and De Nicoló (2002) is particularly close to our paper. Unlike most sign-restriction studies that impose signs directly on IRFs or structural parameters, Canova and De Nicoló (2002) use signs of conditional cross-correlations to assign structural interpretations to orthogonal innovations. Our paper is closely related to theirs because both papers use comovement patterns conditional on structural shocks as identifying information. The key distinction is whether the randomness of the rotation matrix is treated as part of inference. Canova and De Nicoló (2002) start from orthogonal but economically unlabeled VAR innovations, rotate them over a finite grid, and choose the rotation under which as many rotated innovations as possible display the theory-implied conditional cross-correlation signs. Once this rotation is selected, however, uncertainty over admissible rotations is not propagated through the inference. By contrast, we draw the rotation matrix from the Haar measure and impose the correlation restriction by truncating the prior support of the structural parameters, thereby incorporating rotation uncertainty without discretizing the rotation space.

A second difference is modularity. In Canova and De Nicoló (2002), the conditional cross-correlation restrictions are part of a grid-search and rotation-selection procedure, so it is less straightforward to combine them with other identifying restrictions such as sign restrictions on IRFs or structural parameters, zero restrictions, or narrative sign restrictions. In our framework, the correlation restriction is instead written as a nonlinear inequality restriction evaluated on each structural draw, so it can be imposed jointly with other identifying restrictions within the same posterior-sampling algorithm. This allows us to bring correlation restrictions into Bayesian SVAR inference as a modular identifying restriction.

Taken together, the above literature has developed a variety of identifying strategies for

SVARs. This paper contributes to this literature by introducing the *correlation restriction* into Bayesian SVAR inference as an additional economically plausible source of structural identification. The restriction is based on the key features of demand and supply shocks: The correlation between output and prices conditional on demand (supply) shocks must be positive (negative). Importantly, the correlation restriction does not replace existing identifying restrictions. Rather, it can be imposed jointly with sign, zero, or narrative sign restrictions and provides an additional source of identifying information through a theory-consistent restriction on conditional comovement. We show that this additional restriction is useful in identifying monetary policy shocks and oil supply shocks by narrowing the admissible set of structural parameters and sharpening posterior inference.

The paper proceeds as follows. Section 2 presents the basic SVAR framework for the correlation restriction as well as sign restrictions on the IRFs or the structural parameters, and defines the simulated conditional correlation between output and prices on demand or supply shocks. Section 3 describes the functions that can characterize the correlation restriction and the sign restrictions on the IRFs or the structural parameters, and compares the correlation restriction with the sign restrictions. Section 4 explains the Bayesian inference for the SVAR with the correlation and sign restrictions, and describes the algorithm for sampling the structural parameters under the sign restrictions and the correlation restriction for demand shocks. Section 5 applies the correlation restriction to identify monetary policy shocks and oil supply shocks and reports the robust Bayesian inference results. Section 6 concludes the paper.

2 SVAR framework

Consider the general form of an SVAR:

$$\mathbf{Y}_t'\mathbf{A} = \sum_{j=1}^{p} \mathbf{Y}_{t-j}'\mathbf{C}_j + \boldsymbol{\alpha} + \boldsymbol{\varepsilon}_t', \tag{1}$$

where $\mathbf{Y}_t$ is an $n \times 1$ vector of endogenous variables, $\mathbf{A}$ and $\mathbf{C}_j$ $(j = 1, \ldots, p)$ are $n \times n$ parameter matrices, $\boldsymbol{\alpha}$ is a $1 \times n$ vector of constants, p is the lag length, T is the sample size, and $\boldsymbol{\varepsilon}_t$ is an $n \times 1$ vector of structural shocks whose covariance matrix is $\mathbf{I}_n$ (the $n \times n$ identity matrix). Equation (1) can be written more compactly as

$$\mathbf{Y}_t'\mathbf{A} = \mathbf{X}_t'\boldsymbol{\Gamma} + \boldsymbol{\varepsilon}_t', \tag{2}$$

where $\boldsymbol{\Gamma}' = \left[\mathbf{C}_1', \cdots, \mathbf{C}_p', \boldsymbol{\alpha}'\right]$ is an $n \times m$ $(m = np + 1)$ matrix of parameters, and $\mathbf{X}_t' = \left[\mathbf{Y}_{t-1}', \cdots, \mathbf{Y}_{t-p}', 1\right]$ is a $1 \times m$ vector consisting of lagged endogenous variables and the constant term. Multiplying both sides of Equation (2) by $\mathbf{A}^{-1}$ gives the reduced-form representation:

$$\mathbf{Y}_t' = \mathbf{X}_t'\mathbf{B} + \mathbf{u}_t', \tag{3}$$

where $\mathbf{B} = \mathbf{\Gamma}\mathbf{A}^{-1}$ and $\mathbf{u}_t' = \boldsymbol{\varepsilon}_t'\mathbf{A}^{-1}$. Here, $\mathbf{u}_t$ denotes the vector of reduced-form shocks, and its covariance matrix is $\mathbb{E}_t\left[\mathbf{u}_t\mathbf{u}_t'\right] = \left(\mathbf{A}\mathbf{A}'\right)^{-1} = \mathbf{\Sigma}$.

2.1 Impulse response functions

Let the response of the ith variable to the kth structural shock at horizon h be the element in row i and column k of $\mathbf{IR}_h\left(\mathbf{\Psi}\right)$, where $\mathbf{\Psi} = \left(\mathbf{A}, \mathbf{\Gamma}\right)$ collects the structural parameters. Then, $\mathbf{IR}_h\left(\mathbf{\Psi}\right)$ is defined by

$$\mathbf{IR}_0\left(\mathbf{\Psi}\right) = \left(\mathbf{A}^{-1}\right)' \tag{4}$$

$$\text{and } \mathbf{IR}_h\left(\mathbf{\Psi}\right) = \sum_{j=1}^{\min(h,p)} \left(\mathbf{C}_j\mathbf{A}^{-1}\right)' \mathbf{IR}_{h-j}\left(\mathbf{\Psi}\right), \text{ for } h > 0. \tag{5}$$

Equations (4) and (5) clearly show that the IRFs depend only on the structural parameters. Moreover, from Equation (4), the contemporaneous response, $\mathbf{IR}_0$, of the ith variable to the kth structural shock is the (i,k)th element of $(\mathbf{A}^{-1})'$, one of the structural parameters. Hence, restricting the element in the (i,k)th element of $(\mathbf{A}^{-1})'$ to be positive (negative) is identical to restricting the initial IRF of the ith variable to the kth structural shock to be positive (negative).

2.2 Conditional correlations between output and prices

Let the vector of structural shocks $\boldsymbol{\varepsilon}_t$ contain demand shocks ε_t^d and supply shocks ε_t^s. Then, given the structural parameters $\mathbf{\Psi}$, we can obtain simulated endogenous variables generated only by demand shocks (or only by supply shocks) with all other structural shocks shut down (i.e., setting all other structural shocks to zero for all t). We denote by $\left\{\tilde{y}_{z,t}^d\right\}_{t=1}^{\tau}$ and $\left\{\tilde{p}_{z,t}^d\right\}_{t=1}^{\tau}$ the zth ($z = 1, 2, \cdots, N$) simulated output and prices of length τ in simulations where all structural shocks, except for demand shocks, are set to zero. Similarly, we denote by $\left\{\tilde{y}_{z,t}^s\right\}_{t=1}^{\tau}$ and $\left\{\tilde{p}_{z,t}^s\right\}_{t=1}^{\tau}$ the zth simulated output and prices of length τ in simulations where all structural shocks, other than supply shocks, are set to zero. Using these simulated output and price series, we compute the simulated correlation between output and prices conditional on demand shocks $\rho_{y,p}^d$ and the simulated correlation between output and prices conditional on supply shocks $\rho^s_{y,p}$. Specifically,

$$\begin{aligned} \rho^d_{y,p} &= \frac{1}{N}\sum_{z=1}^{N} \widetilde{\text{corr}}\left(\left\{\tilde{y}^d_{z,t}\right\}_{t=b+1}^{\tau}, \left\{\tilde{p}^d_{z,t}\right\}_{t=b+1}^{\tau}\right) \\ \text{and } \rho^s_{y,p} &= \frac{1}{N}\sum_{z=1}^{N} \widetilde{\text{corr}}\left(\left\{\tilde{y}^s_{z,t}\right\}_{t=b+1}^{\tau}, \left\{\tilde{p}^s_{z,t}\right\}_{t=b+1}^{\tau}\right), \end{aligned} \tag{6}$$

where $\widetilde{\text{corr}}(\cdot)$ denotes the simulated correlation computed from the zth simulated series, and b denotes the burn-in period.

The simulated series $\left\{\tilde{y}^d_{z,t}, \tilde{p}^d_{z,t}\right\}_{t=1}^{\tau}$ depend not only on $\boldsymbol{\Psi}$ but also on the zth simulated demand shock sequence $\left\{\varepsilon^d_{z,t}\right\}_{t=1}^{\tau}$ and the initial values of output and prices. Thus, in finite samples, $\rho^d_{y,p}$ is not determined solely by $\boldsymbol{\Psi}$. However, the burn-in period removes the effect of initial values. Moreover, as τ and N become large, simulation noise vanishes, so that $\rho^d_{y,p}$ converges to the theoretical conditional correlation between output and prices when all non-demand shocks are shut down. Section 3.3 shows that this theoretical conditional correlation depends only on $\boldsymbol{\Psi}$. Therefore, for sufficiently large τ and N, restrictions on $\rho^d_{y,p}$ can be expressed as functions of $\boldsymbol{\Psi}$. By similar logic, restrictions on $\rho^s_{y,p}$ can also be expressed as functions of $\boldsymbol{\Psi}$. Hereafter, assuming that τ and N are sufficiently large, we denote the conditional correlations in Equation (6) by $\rho^d_{y,p}(\boldsymbol{\Psi})$ and $\rho^s_{y,p}(\boldsymbol{\Psi})$ to indicate that they depend only on $\boldsymbol{\Psi}$.

3 Identification via sign and correlation restrictions

The structural parameters in Equation (1) are not identified, so additional restrictions are required for identification. This paper proposes a new approach that uses the conditional correlation between output and prices conditional on demand or supply shocks: the correlation restriction. Our approach builds on the existing SVAR literature that imposes sign restrictions on the IRFs and the structural parameters, including Canova and De Nicoló (2002), Uhlig (2005), Baumeister and Peersman (2013), and Arias, Caldara, and Rubio-Ramírez (2019), among others.

The correlation restriction imposes economically plausible comovement between output and prices in the SVAR by requiring the structural parameters to deliver the appropriate sign of the correlation between the two variables conditional on each type of shock. In particular, demand shocks are associated with positive comovement, so output and prices driven by demand shocks tend to move in the same direction. A contractionary monetary policy shock, for example, raises nominal and real interest rates, reduces consumption and investment, and depresses aggregate demand. Consequently, both output and prices fall.

Supply shocks, by contrast, are associated with negative comovement between output and prices. For instance, an oil supply shock (reducing world oil production) raises oil prices. Oil prices are included directly in the consumer price index (CPI), and higher oil prices also increase production costs. As a result, prices rise while output falls. Given these economic intuitions, researchers wish to restrict the structural parameters so that the correlation between output and prices is positive conditional on demand shocks and negative conditional on supply shocks.

3.1 Sign restrictions

Rubio-Ramírez, Waggoner, and Zha (2010) and Arias, Rubio-Ramírez, and Waggoner (2018) show that sign restrictions can be written as

$$\Upsilon(\mathbf{\Psi}) = \left(\mathbf{e}_{1,n}'\mathbf{F}(\mathbf{\Psi})'\mathbf{S}_1', \cdots, \mathbf{e}_{n,n}'\mathbf{F}(\mathbf{\Psi})'\mathbf{S}_n'\right)' > \mathbf{0}, \tag{7}$$

where $\mathbf{e}_{k,n}$ is the kth column of the $n \times n$ identity matrix.

The definitions of $\mathbf{F}(\mathbf{\Psi})$ and $\mathbf{S}_k$ depend on whether the sign restrictions are imposed on the IRFs or on the structural parameters. When imposing sign restrictions on the IRFs, $\mathbf{F}(\mathbf{\Psi})$ stacks the IRFs vertically at the relevant horizons, and $\mathbf{S}_k$ selects the horizons and variables to which the sign restrictions are applied. When imposing sign restrictions on the structural parameters, $\mathbf{F}(\mathbf{\Psi}) = \mathbf{\Psi}$, and $\mathbf{S}_k$ selects the entries of $\mathbf{\Psi}$ on which the sign restrictions are imposed.

3.2 Correlation restriction

Sign restrictions are inequality restrictions and are therefore often viewed as less stringent than exact zero restrictions. A drawback, however, is that sign restrictions alone may leave a large set of admissible rotations, and hence wide credible sets for the IRFs. The correlation restriction provides an additional, economically interpretable restriction that targets the comovement between output and prices. As a result, it can shrink the set of admissible rotations and sharpen inference on the dynamic responses. Thus, we identify demand and supply shocks by combining sign restrictions on the IRFs or the structural parameters with the correlation restriction.

Specifically, since, conditional on the demand shock ε_t^d, the correlation between output and prices is expected to be positive, we impose

$$\rho_{y,p}^d(\mathbf{\Psi}) > 0. \tag{8}$$

Conversely, conditional on the supply shock ε_t^s, we expect a negative correlation between output and prices. We therefore impose

$$\rho_{y,p}^{s}(\mathbf{\Psi}) < 0. \tag{9}$$

Note that applying the correlation restriction requires the SVAR model in Equation (1) to be covariance stationary so that the conditional correlations are well defined.

Discussion for the size of the conditional correlations Some might consider this restriction on the conditional correlation to be too weak because many researchers typically consider correlations within the range of -0.2 to 0.2 to be meaningless. Accordingly, they might wish to impose a stronger restriction on the conditional correlation than the baseline restrictions in Equations (8) and (9). In this case, we can impose $\rho_{y,p}^{d}(\mathbf{\Psi}) > 0.2$ or $\rho_{y,p}^{s}(\mathbf{\Psi}) < -0.2$. We refer to this as the strong correlation restriction.

The strong correlation restriction is also useful, given the simulated conditional correlations in Equation (6). If a structural parameter draw implies a conditional correlation close to zero, finite-sample simulation noise, rather than the structural parameter itself, can determine the sign of the simulated conditional correlation. Therefore, imposing a buffer away from zero helps prevent noise-driven sign flipping, and makes draw acceptance and posterior inference more stable and reproducible.

The effect of the strong correlation restriction need not be uniform and may vary across datasets. In our case, to examine how imposing the strong correlation restriction affects the estimation results, Online Appendix A reports results under the strong correlation restriction. These estimation results are nearly identical to those under the baseline correlation restrictions.

3.3 Comparison of sign restrictions with the correlation restriction

One might think that simply imposing sign restrictions, that a positive demand shock increases output and price during a reasonable period, can reach the same results as a correlation restriction that the conditional correlation between the two on demand shocks is positive. This is because the sign restrictions can produce a positive conditional correlation between the two at least during the period. Hence, in this subsection, we compare sign restrictions with the correlation restriction and clearly show that it is not possible for the two restrictions to lead to the same results, that is, they cannot be identical. In addition, we also show that the simulated conditional correlation between the two variables is determined only by the structural parameters.

For expositional clarity, we focus our discussion on the identification of demand shocks.

Specifically, let y_t and p_t denote output and the price level, respectively, and consider their Wold representations in our SVAR model when all structural shocks, except for demand shocks (ε_t^d), are shut down:

$$\begin{aligned} y_t &= \alpha_y + \sum_{h=0}^{\infty} \mathbf{IR}_h^{(y,d)}(\mathbf{\Psi})\, \varepsilon_{t-h}^d \\ \text{and } \; p_t &= \alpha_p + \sum_{h=0}^{\infty} \mathbf{IR}_h^{(p,d)}(\mathbf{\Psi})\, \varepsilon_{t-h}^d, \end{aligned} \tag{10}$$

where $\mathbf{IR}_h^{(y,d)}$ and $\mathbf{IR}_h^{(p,d)}$ are the entries of $\mathbf{IR}_h$ associated with the output and price IRFs to a demand shock, and α_y and α_p are constant terms.

To ensure a positive correlation between output and prices with respect to demand shocks, we impose two sign restrictions that a positive (expansionary) demand shock increases output and prices for a sufficiently long horizon:

$$\begin{aligned} &\mathbf{IR}_h^{(y,d)}(\mathbf{\Psi}) > 0 \\ \text{and } &\mathbf{IR}_h^{(p,d)}(\mathbf{\Psi}) > 0, \;\; \text{for } h = 0, \ldots, H-1, \end{aligned}$$

where H is the horizon over which the sign restrictions are imposed. These sign restrictions can be equivalently rewritten as

$$\mathbf{IR}_h^{(y,d)}(\mathbf{\Psi}) > 0 \tag{11}$$

$$\text{and } \mathbf{IR}_h^{(y,d)}(\mathbf{\Psi})\, \mathbf{IR}_h^{(p,d)}(\mathbf{\Psi}) > 0. \tag{12}$$

Condition (12), $\mathbf{IR}_h^{(y,d)}(\mathbf{\Psi})\, \mathbf{IR}_h^{(p,d)}(\mathbf{\Psi}) > 0$, induces a positive correlation between output and prices at least during H periods when a demand shock, ε_t^d, hits.

For identifying demand shocks, we then describe a correlation restriction that the conditional correlation between output and prices on demand shocks must be positive. Given the assumption that ε_t^d is an i.i.d. random variable with unit variance and the number of simulations is sufficiently large such that the simulated correlation converges to the theoretical correlation, using Equation (10), the correlation restriction can be written as

$$\widetilde{\text{corr}}(\tilde{y}_t^d, \tilde{p}_t^d) = \frac{\widetilde{\text{cov}}(\tilde{y}_t^d, \tilde{p}_t^d)}{\tilde{\sigma}\left(\tilde{y}_t^d\right) \tilde{\sigma}\left(\tilde{p}_t^d\right)} \approx \frac{\sum_{h=0}^{\infty} \mathbf{IR}_h^{(y,d)}(\mathbf{\Psi})\, \mathbf{IR}_h^{(p,d)}(\mathbf{\Psi})}{\sqrt{\sum_{h=0}^{\infty}\left(\mathbf{IR}_h^{(y,d)}(\mathbf{\Psi})\right)^2}\sqrt{\sum_{h=0}^{\infty}\left(\mathbf{IR}_h^{(p,d)}(\mathbf{\Psi})\right)^2}} > 0, \tag{13}$$

where $\widetilde{\text{cov}}(\cdot)$ and $\tilde{\sigma}(\cdot)^2$ are the simulated covariance and variance, respectively. Condition (13) clearly shows that the correlation restriction depends solely on the structural parameters.

For ease of comparison, we rewrite the correlation restriction (Condition 13) more compactly as follows:

$$\widetilde{\text{cov}}(\tilde{y}_t^d, \tilde{p}_t^d) \approx \sum_{h=0}^{\infty} \mathbf{IR}_h^{(y,d)}(\boldsymbol{\Psi})\, \mathbf{IR}_h^{(p,d)}(\boldsymbol{\Psi}) > 0. \tag{14}$$

The correlation restriction (Condition 14) requires the sum of the products of the output and price IRFs to be positive over the full propagation path. Unlike the two sign restrictions (Conditions 11 and 12), the correlation restriction (Condition 14) does not require the product to be positive at every horizon. Therefore, under the correlation restriction (Condition 14), the output and price IRFs can temporarily move in opposite directions, even at early horizons, although the overall comovement is positive.

In order for the two sign restrictions to generate a positive conditional correlation between output and prices on demand shocks, like the correlation restriction, we need to set a sufficiently large H in Conditions (11) and (12). This is because increasing H expands the number of horizons h that satisfy $\mathbf{IR}_h^{(y,d)}(\boldsymbol{\Psi})\, \mathbf{IR}_h^{(p,d)}(\boldsymbol{\Psi}) > 0$, which increases the likelihood of satisfying Condition (14). However, a too large H can result in significant distortions of the estimated IRFs. Conversely, a small H can mitigate such distortions, but with a small H, the sign restrictions may fail to satisfy Condition (14). This failure can lead to a negatively correlated output and prices when all structural shocks, other than the demand shock, are eliminated, an outcome that directly contradicts theoretical intuition.

Even when the sign restrictions yield a positive conditional correlation between output and prices on demand shocks by setting a sufficiently large H, the correlation restriction and the sign restrictions are not equivalent. The correlation restriction in Condition (14) requires only that output and prices comove on average after the shock, and it does not fix the signs of their IRFs. That is, the signs of the output and price IRFs can either be positive or negative (even their signs can differ temporarily), as long as the overall conditional correlation between output and prices is positive. By contrast, the sign restrictions in Conditions (11) and (12) require both output and prices to rise at prespecified horizons after a demand shock. In addition, under the correlation restriction, the timing of comovement is determined by the data, whereas it must be specified ex ante under the sign restrictions. Therefore, even when both approaches imply a positive conditional correlation between output and prices, the results under the sign restrictions are obtained in a far more restrictive environment than those under the correlation restriction.

In introducing positive comovement between output and prices, the correlation restriction has additional advantages over sign restrictions. Under the correlation restriction, the simulated conditional correlation is computed after discarding an initial burn-in period. Therefore, it mainly constrains long-run comovement between output and prices, rather than short-run dynamics after the shock. By contrast, the sign restrictions in Conditions (11) and (12) substantially constrain short-run dynamics over the H periods after the shock.

Taken together, these comparisons show that imposing correlation and sign restrictions jointly offers advantages over relying on sign restrictions alone. The correlation restriction automatically enforces the necessary condition that output and prices are positively correlated with respect to demand shocks. Therefore, even if the demand shocks identified by sign restrictions do not fully capture the nature of demand shocks as described by economic theory, the correlation restriction can compensate for this limitation. Furthermore, in practice, researchers usually impose minimal, uncontroversial sign restrictions. That is, it is rare to impose sign restrictions on both output and price IRFs to identify demand shocks, as in this subsection. And the correlation restriction is uncontroversial, since it has a strong theoretical basis. Therefore, from a practical perspective, it seems better to impose a sign restriction on one variable (output or prices) and a correlation restriction for demand or supply shocks, which would substantially help identify demand or supply shocks more precisely.

4 Bayesian inference

This section shows how to impose the correlation restriction in the SVAR in which sign restrictions are already imposed. Our approach extends the Bayesian methodology of Rubio-Ramírez, Waggoner, and Zha (2010), Arias, Rubio-Ramírez, and Waggoner (2018), and Antolín-Díaz and Rubio-Ramírez (2018). For brevity, we focus on the correlation restriction for demand shocks, $\rho^{d}_{y,p}(\mathbf{\Psi}) > 0$. The correlation restriction for supply shocks, $\rho^{s}_{y,p}(\mathbf{\Psi}) < 0$, can be handled analogously.

4.1 Prior distribution

We begin by specifying a prior for the reduced-form parameters $(\mathbf{B}, \mathbf{\Sigma})$. We assign a conjugate Normal-Inverse-Wishart prior, $\mathcal{NIW}(\nu, \mathbf{\Phi}, \mathbf{S}, \mathbf{\Omega})$. Specifically,

$$\begin{aligned} \text{vec}(\mathbf{B}) \mid \mathbf{\Sigma} &\sim \mathcal{N}\big(\text{vec}(\mathbf{S}),\, \mathbf{\Sigma} \otimes \mathbf{\Omega}\big) \\ \text{and } \mathbf{\Sigma} &\sim \mathcal{IW}(\nu, \mathbf{\Phi}). \end{aligned}$$

Our study uses an improper, noninformative prior with ν, $\mathbf{\Phi}$, $\mathbf{S}$, and $\mathbf{\Omega}^{-1}$ set to zero (or zero matrices, as appropriate), following Arias, Caldara, and Rubio-Ramírez (2019), to ensure that the estimation results are driven primarily by the data rather than the prior.

In a Bayesian SVAR, sign restrictions are implemented by truncating the prior distribution for the reduced-form parameters given above. The correlation restriction is implemented in the same way. However, as Equations (7) and (8) show, these two restrictions are stated in terms of the structural parameters, rather than the reduced-form parameters. Therefore, to truncate the prior for the reduced-form parameters using these two restrictions, we first need to derive a mapping between the structural and reduced-form parameters, and to rewrite the restrictions as functions of the reduced-form parameters. This makes it clear which region of the prior distribution for the reduced-form parameters satisfies the restrictions.

To derive the mapping, we introduce the set $\mathcal{O}(n)$ of $n \times n$ orthogonal matrices and an orthogonal rotation matrix $\mathbf{Q} \in \mathcal{O}(n)$ that links the structural shocks to the reduced-form shocks. Expressing the reduced-form shocks in Equation (3) as rotations of the structural shocks via $\mathbf{Q}$ yields

$$\mathbf{Y}_t' = \mathbf{X}_t'\mathbf{B} + \boldsymbol{\varepsilon}_t'\mathbf{Q}'h\left(\boldsymbol{\Sigma}\right), \tag{15}$$

where $h\left(\boldsymbol{\Sigma}\right)$ is an $n \times n$ matrix satisfying $h\left(\boldsymbol{\Sigma}\right)' h\left(\boldsymbol{\Sigma}\right) = \boldsymbol{\Sigma}$. A Cholesky decomposition is typically used for $h\left(\cdot\right)$.

Building on Equations (2) and (15), Arias, Rubio-Ramírez, and Waggoner (2018) and Antolín-Díaz and Rubio-Ramírez (2018) show that the mapping between the structural parameters and the reduced-form parameters is given by

$$f_h\left(\boldsymbol{\Psi}\right) = \Big(\underbrace{\boldsymbol{\Gamma}\mathbf{A}^{-1}}_{\mathbf{B}}, \underbrace{\left(\mathbf{A}\mathbf{A}'\right)^{-1}}_{\boldsymbol{\Sigma}}, \underbrace{h\big(\left(\mathbf{A}\mathbf{A}'\right)^{-1}\big)\mathbf{A}}_{\mathbf{Q}}\Big) \tag{16}$$

and that its inverse is given by

$$f_h^{-1}\left(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}\right) = \Big(\underbrace{h\big(\boldsymbol{\Sigma}\big)^{-1}\mathbf{Q}}_{\mathbf{A}}, \underbrace{\mathbf{B}h\big(\boldsymbol{\Sigma}\big)^{-1}\mathbf{Q}}_{\boldsymbol{\Gamma}}\Big). \tag{17}$$

Equation (17) implies that the two restrictions can be expressed in terms of the reduced-form parameters and $\mathbf{Q}$.

Specifically, the sign restrictions in Equation (7) can be written as

$$\Upsilon\left(f_h^{-1}\left(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}\right)\right) > 0.$$

Under a uniform prior on $\mathbf{Q}$ over $\mathcal{O}(n)$, the sign restrictions truncate the prior distribution to the subset of the parameter space that satisfies the above inequality. The resulting prior distribution is given by

$$\pi\left(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q} \mid \Upsilon\left(f_h^{-1}\left(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}\right)\right) > 0\right) \propto \mathcal{NIW}\left(\nu, \boldsymbol{\Phi}, \mathbf{S}, \boldsymbol{\Omega}\right) [\![\Upsilon\left(f_h^{-1}\left(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}\right)\right) > 0]\!],$$

where $[\![\cdot]\!]$ denotes an indicator function that equals 1 if the argument is true and 0 otherwise.

Similarly, the correlation restriction for demand shocks, Equation (8), can be expressed as

$$\rho^d_{y,p}\left(f_h^{-1}\left(\mathbf{B},\mathbf{\Sigma},\mathbf{Q}\right)\right) > 0,$$

and the associated truncated prior is given by

$$\pi\left(\mathbf{B},\mathbf{\Sigma},\mathbf{Q} \mid \rho^d_{y,p}\left(f_h^{-1}\left(\mathbf{B},\mathbf{\Sigma},\mathbf{Q}\right)\right) > 0\right) \propto \mathcal{NIW}\left(\nu,\mathbf{\Phi},\mathbf{S},\mathbf{\Omega}\right) [\![\rho^d_{y,p}\left(f_h^{-1}\left(\mathbf{B},\mathbf{\Sigma},\mathbf{Q}\right)\right) > 0]\!].$$

Moreover, when the correlation restriction is imposed jointly with the sign restrictions, the prior is truncated to

$$\mathcal{NIW}\left(\nu,\mathbf{\Phi},\mathbf{S},\mathbf{\Omega}\right) [\![\Upsilon\left(f_h^{-1}\left(\mathbf{B},\mathbf{\Sigma},\mathbf{Q}\right)\right) > 0]\!][\![\rho^d_{y,p}\left(f_h^{-1}\left(\mathbf{B},\mathbf{\Sigma},\mathbf{Q}\right)\right) > 0]\!].$$

4.2 Posterior distribution

We first derive the posterior distribution in the absence of sign and correlation restrictions. Since $\mathcal{NIW}$ is conjugate to the likelihood function of the SVAR, the posterior distribution also has a closed form $\mathcal{NIW}$ representation (Arias, Rubio-Ramírez, and Waggoner, 2018). In particular, the posterior is

$$\begin{aligned} &\mathcal{NIW}(\hat{\nu},\hat{\mathbf{\Phi}},\hat{\mathbf{S}},\hat{\mathbf{\Omega}}), \qquad (18)\\ \text{with } &\hat{\nu} = T+\nu,\\ &\hat{\mathbf{\Omega}} = \left(\mathbf{X}'\mathbf{X}+\mathbf{\Omega}^{-1}\right)^{-1},\\ &\hat{\mathbf{S}} = \hat{\mathbf{\Omega}}\left(\mathbf{X}'\mathbf{Y}+\mathbf{\Omega}^{-1}\mathbf{S}\right),\\ \text{and } &\hat{\mathbf{\Phi}} = \mathbf{Y}'\mathbf{Y}+\mathbf{\Phi}+\mathbf{S}'\mathbf{\Omega}^{-1}\mathbf{S}-\hat{\mathbf{S}}'\hat{\mathbf{\Omega}}^{-1}\hat{\mathbf{S}}, \end{aligned}$$

for $\mathbf{Y} = [\mathbf{Y}_1 \ \cdots \ \mathbf{Y}_T]'$ and $\mathbf{X} = [\mathbf{X}_1 \ \cdots \ \mathbf{X}_T]'$.

When restrictions are imposed via a truncated prior, posterior simulation remains straightforward, provided that the posterior distribution without the restrictions is available in closed form. In particular, we use an acceptance-rejection sampler that draws candidate values from this closed form posterior and retains only those that satisfy the sign and correlation restrictions. The following algorithm provides a concrete description of how our sampling procedure for the structural parameters operates when both the sign restrictions and the correlation restriction are imposed.

Algorithm 1: `Sampling the structural parameters under the sign restrictions and the correlation restriction for demand shocks`

(i) *Draw* $(\mathbf{B},\mathbf{\Sigma})$ *independently from the* $\mathcal{NIW}$ *posterior distribution in Equation (18).*

(ii) *Draw* $\mathbf{Q}$ *independently from the uniform distribution over* $\mathcal{O}(n)$.

(iii) *Based on the sampled* $(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})$, *check whether the sign restrictions in Equation (7) are satisfied. If the sign restrictions are satisfied (i.e.,* $[\![\Upsilon\left(f_h^{-1}(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})\right) > 0]\!] = 1$*), proceed to Step (iv) and otherwise return to Step (i).*

(iv) *Obtain* N *simulated series of output and prices with the length* τ *by simulating the SVAR model with respect only to demand shocks. Using these simulated series, compute* $\rho_{y,p}^d$ *in Equation (6). If the correlation restriction is satisfied (i.e.,* $[\![\rho_{y,p}^d\left(f_h^{-1}(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})\right) > 0]\!] = 1$*), proceed to Step (v) and otherwise return to Step (i).*

(v) *Using Equation (17), transform* $(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})$ *into* $\mathbf{\Psi}$. *If the required number of draws of the structural parameters* $\mathbf{\Psi}$ *has been achieved, terminate the algorithm and otherwise return to Step (i).*

5 Applications

In this section, we use the correlation restriction to identify monetary policy and oil supply shocks. Monetary policy shocks are a type of demand shock, and oil supply shocks are, as the name suggests, a type of supply shock. Therefore, in the former case, we impose a correlation restriction that the correlation between output and prices conditional on monetary policy shocks must be positive. For identifying oil supply shocks, we impose a correlation restriction that the correlation between output and prices conditional on oil supply shocks must be negative. Using the two correlation restrictions, we first identify monetary policy shocks and then oil supply shocks.

5.1 Monetary policy shocks

The most typical example of demand shocks is monetary policy shocks. As is well-known, a contractionary monetary policy shock raises nominal and real interest rates, which depresses aggregate demand. Accordingly, output and prices fall in response to the shock. In other words, conditional on monetary policy shocks, output and prices are positively correlated. Hence, in this subsection, we use this feature of monetary policy shocks to identify monetary policy shocks more precisely.

Specifically, we use the SVAR proposed by Arias, Caldara, and Rubio-Ramírez (2019), which does a relatively good job in identifying monetary policy shocks. We add the correlation restriction for demand shocks to the SVAR of Arias, Caldara, and Rubio-Ramírez (2019). Then, we show that when adding the correlation restriction, a contractionary monetary policy shock decreases both output and prices (i.e., the price puzzle disappears).

5.1.1 Data and others

The SVAR model for monetary policy shocks includes six monthly US variables, as in Arias, Caldara, and Rubio-Ramírez (2019): real GDP, the GDP deflator, a commodity price index, total reserves, nonborrowed reserves, and the federal funds rate. The monthly real GDP and GDP deflator, constructed by interpolation of the corresponding quarterly time series, are taken from Arias, Caldara, and Rubio-Ramírez (2019). The source of the commodity price index is Global Financial Data, and total reserves, nonborrowed reserves, and the federal funds rate are obtained from the Federal Reserve Economic Data (FRED). All variables are seasonally adjusted (with the exception of the commodity price index and the federal funds rate). When estimating, we use the first difference of all the variables (with the exception of the federal funds rate) following Arias, Caldara, and Rubio-Ramírez (2019).

The sample period is the same as that in Arias, Caldara, and Rubio-Ramírez (2019): from January 1965 to June 2007. This sample period ensures the exclusion of the effects of the global financial crisis and the period of unconventional monetary policy. The SVAR includes 12 lags and a constant term, as in Arias, Caldara, and Rubio-Ramírez (2019).

5.1.2 Restrictions

To identify monetary policy shocks, we now introduce sign restrictions on the structural parameters, which are the same as those in Arias, Caldara, and Rubio-Ramírez (2019), and a correlation restriction, which uses the main feature of monetary policy (demand) shocks.[6] Specifically, we impose the following sign restrictions on the structural parameters.

Sign Restriction 1: *The contemporaneous reaction of the federal funds rate, the monetary policy instrument, to output and prices is positive.*

`Sign Restriction` 1 means that the central bank increases the federal funds rate when output and prices increase. In particular, we can write the equation for the federal funds rate as:

$$FFR_t = \beta_g GDP_t + \beta_d DEF_t + \beta_c CP_t + \beta_{tr} TR_t + \beta_n NBR_t + \sigma \varepsilon_{FFR,t} + \text{Other Terms},$$

where FFR, GDP, DEF, CP, TR, NBR, and $\varepsilon_{FFR,t}$ denote the federal funds rate, real GDP, the GDP deflator, the commodity price index, total reserves, nonborrowed reserves, and monetary policy shocks, respectively. The constant term, lag variables, and other structural

[6] Note that, different from Arias, Caldara, and Rubio-Ramírez (2019), we do not impose zero restrictions. The results in which zero restrictions are imposed are shown in Online Appendix B.

Table 1: Conditional correlations between output and prices on monetary policy shocks

Restriction	`Sign Restriction` 1	`Sign Restriction` 1, `Correlation Restriction` 1
Posterior median	-0.14	0.46
68% credible interval	[-0.82, 0.56]	[0.14, 0.75]

Note: The conditional correlations for each draw of the structural parameters are computed from 200 simulated series of real GDP and the GDP deflator of 250 periods each (50 periods as burn-in).

shocks are included in Other Terms. `Sign Restriction` 1 implies that β_g and β_d should be positive, which is consistent with monetary rules in New Keynesian DSGE models.

We complement `Sign Restriction` 1 with a correlation restriction, which exploits the main feature of monetary policy (demand) shocks. Specifically, the correlation restriction is summarized:

Correlation Restriction 1: *The correlation between real GDP (output) and the GDP deflator (prices) conditional on monetary policy shocks must be positive.*

`Correlation Restriction` 1 imposes a condition implied by the main feature of monetary policy shocks. That is, output and prices should be positively correlated with respect to monetary policy shocks.

5.1.3 Results

Table 1 presents the conditional correlations between real GDP and the GDP deflator on monetary policy shocks, when only `Sign Restriction` 1 is used, and when both `Sign Restriction` 1 and `Correlation Restriction` 1 are used. For each draw of the structural parameters, we obtain 200 simulated series of real GDP and the GDP deflator of 250 periods each (50 periods as burn-in) with respect to monetary policy shocks only. We compute 200 correlations between real GDP and the GDP deflator conditional on monetary policy shocks for each draw of the structural parameters. We average the 200 correlations, and use the averaged correlation as the conditional correlation between the two variables on monetary policy shocks for each draw of the structural parameters. The posterior median of the conditional correlation between real GDP and the GDP deflator on monetary policy shocks is -0.14, with a 68 percent credible interval of [-0.82, 0.56], when only `Sign Restriction` 1 is used. In contrast, when both restrictions are used, it is 0.46, with a 68 percent credible interval of [0.14, 0.75]. `Correlation Restriction` 1 increases the conditional correlation significantly,

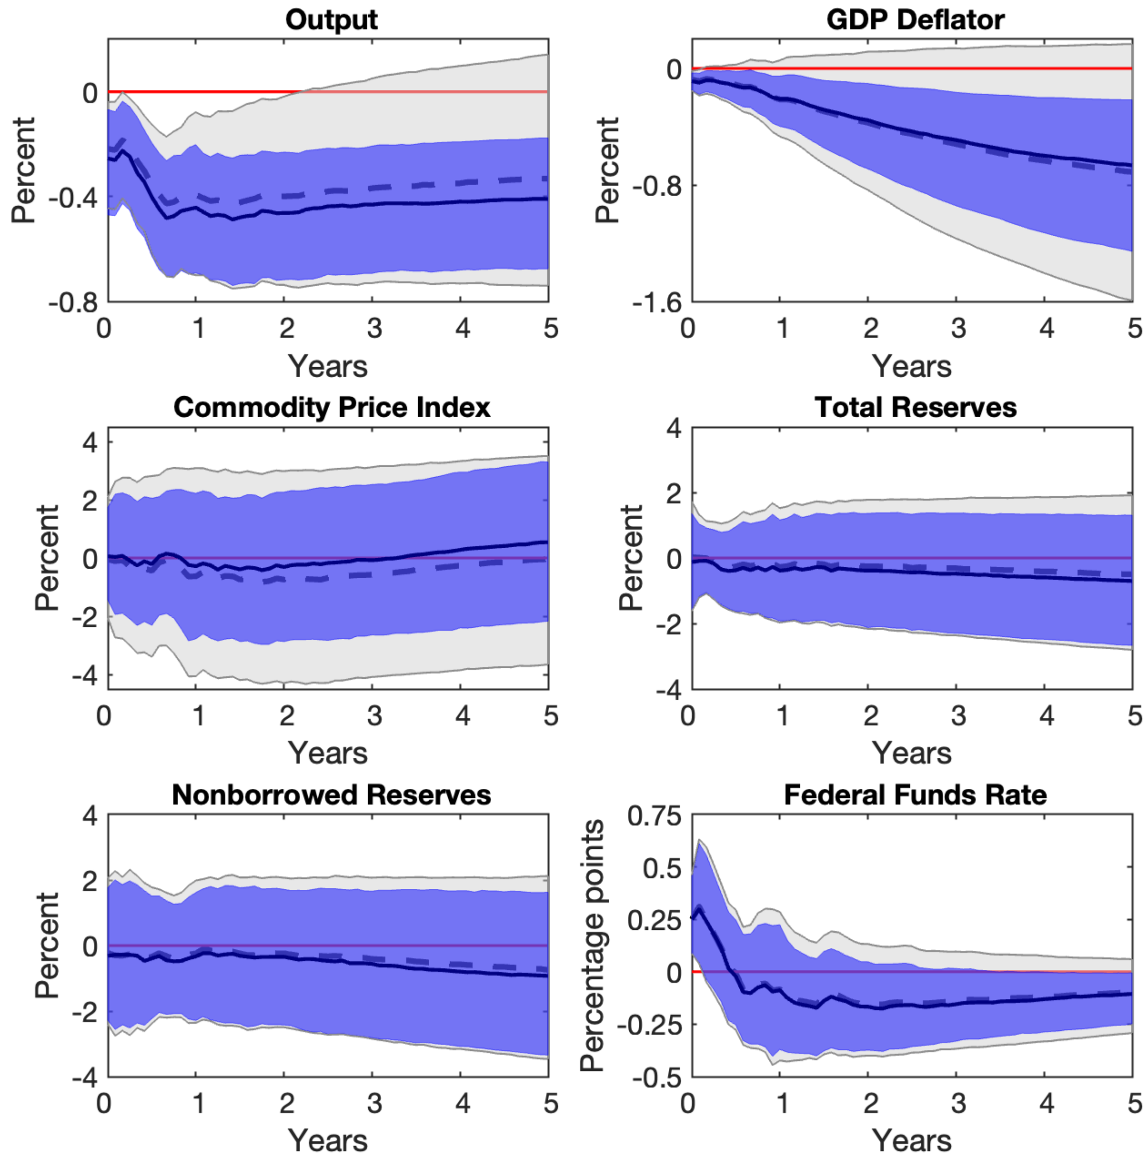

Figure 1: IRFs to a contractionary monetary policy shock

Notes: The lighter shaded area and dashed lines are the 68 percent point-wise credible sets and the median IRFs, respectively, using `Sign Restriction` 1. The darker shaded area and solid lines are the 68 percent point-wise credible sets and the median IRFs, respectively, using `Sign Restriction` 1 and `Correlation Restriction` 1. The IRFs of all variables, except for the federal funds rate, are accumulated to the level.

as the restriction implies.

Figure 1 displays IRFs of the six variables to a contractionary monetary policy shock with `Sign Restriction` 1 only and with both restrictions.[7] The lighter shaded area and dashed lines are the 68 percent point-wise credible sets and the median IRFs, respectively, using only `Sign Restriction` 1. The darker shaded area and solid lines show the equivalent quantities when both `Sign Restriction` 1 and `Correlation Restriction` 1 are used. In both cases,

[7] The IRFs are computed from 6,552 draws that satisfy `Sign Restriction` 1 and 2,892 draws that satisfy both `Sign Restriction` 1 and `Correlation Restriction` 1.

the monetary policy shock is normalized to raise the federal funds rate by 25 basis points in period 0.

Overall, the IRFs in both cases have similar shapes. Specifically, a contractionary monetary policy shock increases immediately the federal funds rate. The shock leads to an immediate and persistent decline in output. The responses of commodity prices and total and nonborrowed reserves to the shock are close to zero. However, the response of prices is notably different in the two cases. With `Sign Restriction` 1 only, the shock does not decrease prices significantly, which implies the price puzzle. In contrast, in response to the shock, prices fall with a very high posterior probability when both `Sign Restriction` 1 and `Correlation Restriction` 1 are imposed. That is, the correlation restriction effectively eliminates the price puzzle.

Furthermore, the credible sets for all the IRFs with both restrictions are narrower than those with `Sign Restriction` 1 only. In other words, the correlation restriction shrinks the credible sets, particularly for output and prices.

In a nutshell, the correlation restriction sharpens the identification of monetary policy shocks. Moreover, it makes the price puzzle disappear.

5.2 Oil supply shocks

A good example of supply shocks is oil supply shocks. An oil supply shock reduces world oil production, and hence the real oil price increases. Since oil is used in the production of almost every good, an increase in the real oil price raises production costs, which reduces aggregate supply. Accordingly, in response to the shock, output falls while prices rise. That is, output and prices are negatively correlated with respect to oil supply shocks. Therefore, we use this feature of oil supply shocks to identify oil supply shocks more precisely in this subsection.

5.2.1 Data and others

The SVAR model for oil supply shocks consists of four quarterly variables, as in Baumeister and Peersman (2013): US real GDP, US CPI, world oil production, and the real US refiners' acquisition cost of imported crude oil (the real oil price).[8] The real oil price is defined as the nominal oil price deflated by US CPI. World oil production and the nominal US refiners' acquisition cost of imported crude oil are obtained from the US Energy Information Adminis-

[8] Note that we do not include a variable for world real economic activity, such as the index of global real economic activity proposed by Kilian (2009), or the world industrial production index proposed by Baumeister and Hamilton (2019), because the focus of this subsection is oil supply shocks. Therefore, only oil supply shocks are explicitly identified, and there is no structural interpretation for the conglomerate of residual oil demand shocks.

tration. The source of the seasonally adjusted US real GDP and CPI is the FRED. Following Baumeister and Peersman (2013), we estimate the SVAR using the first difference of all the variables.

The sample period is from Q1 1974 to Q4 2011, which is the same as that in Baumeister and Peersman (2013). The SVAR includes 4 lags and a constant term, as in Baumeister and Peersman (2013).

5.2.2 Restrictions

Here, to identify oil supply shocks, we provide sign restrictions on the IRFs, which are the same as those in Baumeister and Peersman (2013), and a correlation restriction, which is based on the key feature of supply shocks. Specifically, we use the following sign restrictions on the IRFs.

Sign Restriction 2: *An oil supply shock reduces world oil production and increases the real oil price for four quarters.*

`Sign Restriction` 2 to identify oil supply shocks is common in many studies, including Baumeister and Peersman (2013) and Kim (2020).

In addition, we also impose a correlation restriction based on the primary feature of oil supply shocks (supply shocks). Specifically, the correlation restriction is as follows:

Correlation Restriction 2: *The correlation between real GDP (output) and CPI (prices) conditional on oil supply shocks must be negative.*

`Correlation Restriction` 2 imposes a condition implied by the primary feature of oil supply shocks, i.e., output and prices driven only by oil supply shocks should be negatively correlated.

5.2.3 Results

The conditional correlations between real GDP and CPI on oil supply shocks, when only `Sign Restriction` 2 is used, and when both `Sign Restriction` 2 and `Correlation Restriction` 2 are used, are provided in Table 2. As in the previous subsection, the correlations between real GDP and CPI conditional on oil supply shocks for each draw of the structural parameters are computed from 200 simulated series of real GDP and CPI of 250 periods each (50 periods as burn-in). The posterior median of the conditional correlation between real GDP and CPI on oil supply shocks is -0.81, with a 68 percent credible interval of [-0.99, 0.68], when only `Sign Restriction` 2 is imposed. In contrast, when both restrictions are imposed, it is -0.96,

Table 2: Conditional correlation between output and prices on oil supply shocks

Restriction	`Sign Restriction` 2	`Sign Restriction` 2, `Correlation Restriction` 2
Posterior median	-0.81	-0.96
68% credible interval	[-0.99, 0.68]	[-0.99, -0.74]

Note: The correlations for each draw of the structural parameters are computed from 200 simulated series of real GDP and CPI of 250 periods each (50 periods as burn-in).

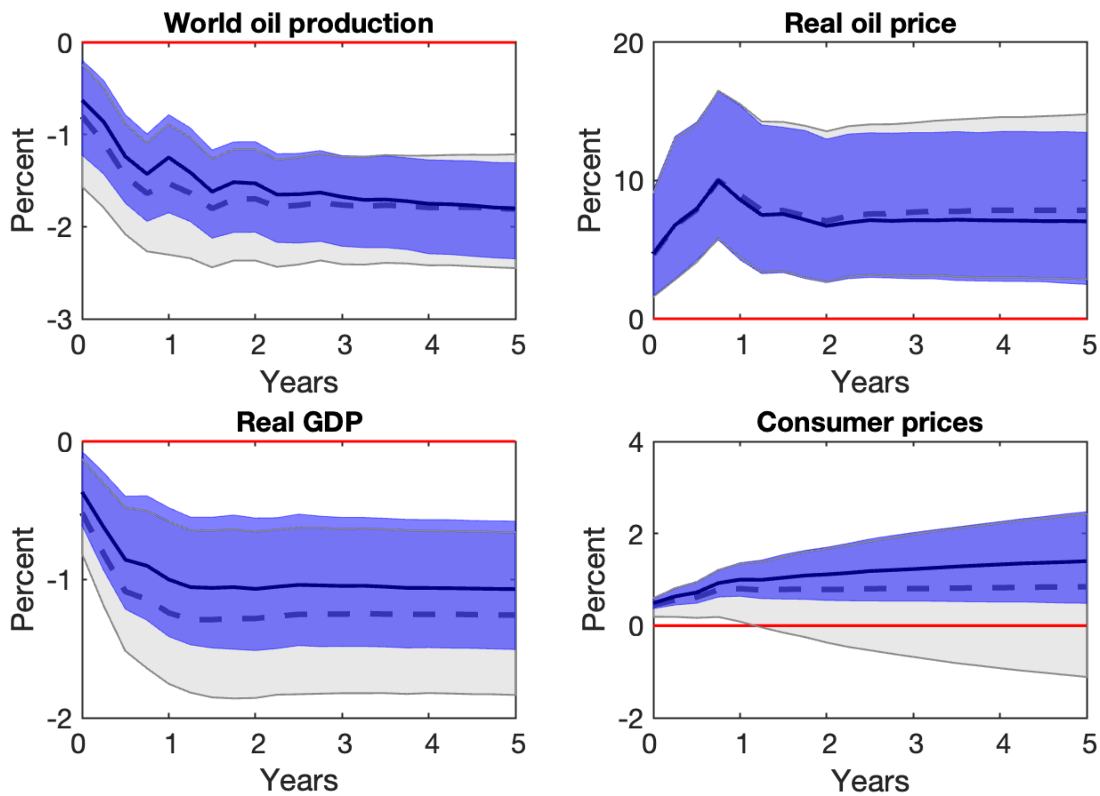

Figure 2: IRFs to an oil supply shock

Notes: The lighter shaded area and dashed lines are the 68 percent point-wise credible sets and the median IRFs, respectively, using `Sign Restriction` 2. The darker shaded area and solid lines are the 68 percent point-wise credible sets and the median IRFs, respectively, using `Sign Restriction` 2 and `Correlation Restriction` 2. The IRFs of all variables are accumulated to the level.

with a 68 percent credible interval of [-0.99, -0.74]. As expected, `Correlation Restriction` 2 leads to a fall in the conditional correlation.

The IRFs of the four variables to an oil supply shock with and without `Correlation Restriction` 2 are shown in Figure 2.[9] The lighter shaded area and dashed lines denote the 68 percent point-wise credible sets and the median IRFs with only `Sign Restriction`

[9] The IRFs are computed from 3,000 draws that satisfy `Sign Restriction` 2 and 1,871 draws that satisfy both `Sign Restriction` 2 and `Correlation Restriction` 2.

2, respectively. The darker shaded area and solid lines show the equivalent quantities with `Sign Restriction` 2 and `Correlation Restriction` 2. The oil supply shock is normalized to raise the real oil price by a maximum of 10 percent.

An oil supply shock decreases world oil production and increases the real oil price in both cases. The responses of the two variables are similar, although the 68 percent (point-wise) credible set for the IRF of world oil production with `Correlation Restriction` 2 is smaller than that without the restriction. However, the responses of output and prices are quite different in the two cases. Specifically, in response to the shock, output declines less in the case with `Correlation Restriction` 2 than in the case without the restriction, and the credible set for the IRF of output is tighter in the former case. The responses of prices to the shock in the two cases differ notably. When only `Sign Restriction` 2 is imposed, the shock does not increase prices significantly at most horizons. However, when both restrictions are imposed, the shock increases prices more with a high posterior probability. Moreover, the credible set for the IRF of prices is much narrower in the case with `Correlation Restriction` 2 than in the case without the restriction.

In short, the correlation restriction helps identify oil supply shocks more sharply. Moreover, it enables prices to increase significantly in response to an oil supply shock, which is in line with economic theory.

Finally, in Online Appendix C, we present the results when the narrative sign restrictions used in Antolín-Díaz and Rubio-Ramírez (2018), `Sign Restriction` 2, and `Correlation Restriction` 2 are imposed to identify oil supply shocks. The results show that the narrative sign restrictions do not make a notable difference.

5.3 Robust Bayesian inference

The IRF results in the applications, reported in Figures 1 and 2, use the conventional Bayesian summaries of posterior medians and 68 percent pointwise credible sets. These summaries are natural in our baseline analysis, but their interpretation is more delicate in a set-identified SVAR. The reason is that a given reduced form can generate multiple admissible structural paths through different rotations, and these paths have the same likelihood. The conventional practice based on a uniform prior over the rotation matrix can place some of these paths inside a reported credible set and others outside it, even though the likelihood does not distinguish among them. This point is emphasized by Baumeister and Hamilton (2024).

The issue matters for our applications because the SVARs also remain set-identified when the correlation restriction is imposed. One way to address this issue, emphasized by Baumeister and Hamilton (2024), is to specify and defend an explicit prior over the rotation matrix when ranking admissible structural paths. An alternative is the robust Bayesian approach of Giacomini and Kitagawa (2021). This approach avoids conditioning on a particular prior over rotations by summarizing the range of posterior mean IRFs attainable over admissible rotations. We implement it by computing the set of posterior means, obtained by taking the lower and upper bounds of the IRFs over the admissible rotations for each reduced-form posterior draw and then averaging these bounds over the reduced-form posterior distribution. Under this methodology, we compare the sign restrictions only with the sign and correlation restrictions.

The results are consistent with the main evidence: adding the correlation restriction rules out a sizable fraction of rotations that satisfy the sign restrictions and narrows the set of posterior means. The portions removed are responses that are less consistent with the shock interpretation, such as expansionary output or price responses to a contractionary monetary policy shock and negative consumer price responses to an oil supply shock. Online Appendix D provides the details.

6 Conclusion

This paper introduces the *correlation restriction* as an identifying restriction in Bayesian SVARs. This restriction is based on the key features of demand and supply shocks from economic theory. These features are that, conditional on demand shocks, output and prices are positively correlated, and that the conditional correlation between the two variables on supply shocks is negative. The correlation restriction constrains the structural parameters to ensure the conditional correlation between output and prices on demand (supply) shocks to be positive (negative). We also explain the Bayesian inference and the algorithm for sampling the structural parameters under both the sign restrictions and the correlation restriction.

The correlation restriction is applied to the SVARs for monetary policy shocks (demand shocks) and oil supply shocks (supply shocks). To identify structural monetary policy shocks, we impose a correlation restriction that the conditional correlation between output and prices on monetary policy shocks must be positive, whereas we impose a correlation restriction that the conditional correlation between the two variables on oil supply shocks must be negative. Estimation results show that the correlation restriction sharpens the identification of monetary policy shocks and enables the price puzzle to disappear. Furthermore, in the case of oil supply shocks, the correlation restriction helps shrink the uncertainty around the IRFs of all variables. In particular, the restriction allows prices to increase significantly in response to an oil supply shock, which is consistent with economic theory.

Finally, it should be noted that the approach in this paper has a clear advantage over the existing, often controversial, restrictions in SVARs. Specifically, demand and supply shocks can be easily distinguished. Furthermore, economic theory uncontroversially implies that the conditional correlation between output and prices on demand (supply) shocks is positive (negative). Therefore, many SVARs can easily make use of the correlation restriction that this paper introduces.

References

Amir-Ahmadi, P. and H. Uhlig (2015). Sign restrictions in Bayesian FAVARs with an application to monetary policy shocks. NBER Working Paper 21738.

Antolín-Díaz, J. and J. F. Rubio-Ramírez (2018). Narrative sign restrictions for SVARs. *American Economic Review 108 (10)*, 2802–2829.

Arias, J. E., D. Caldara, and J. F. Rubio-Ramírez (2019). The systematic component of monetary policy in SVARs: An agnostic identification procedure. *Journal of Monetary Economics 101*, 1–13.

Arias, J. E., J. F. Rubio-Ramírez, and D. F. Waggoner (2018). Inference based on structural vector autoregressions identified with sign and zero restrictions: Theory and applications. *Econometrica 86 (2)*, 685–720.

Baumeister, C. and J. D. Hamilton (2019). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. *American Economic Review 109 (5)*, 1873–1910.

Baumeister, C. and J. D. Hamilton (2024, June). ADVANCES IN USING VECTOR AUTOREGRESSIONS TO ESTIMATE STRUCTURAL MAGNITUDES. *Econometric Theory 40*(3), 472–510.

Baumeister, C. and G. Peersman (2013). Time-varying effects of oil supply shocks on the US economy. *American Economic Journal: Macroeconomics 5 (4)*, 1–28.

Bernanke, B. S., J. Boivin, and P. Eliasz (2005). Measuring the effects of monetary policy: A factor-augmented vector autoregressive (FAVAR) approach. *Quarterly Journal of Economics 120 (1)*, 387–422.

Caldara, D., M. Cavallo, and M. Iacoviello (2019). Oil price elasticities and oil price fluctuations. *Journal of Monetary Economics 103*, 1–20.

Canova, F. and G. De Nicoló (2002). Monetary disturbances matter for business fluctuations in the G-7. *Journal of Monetary Economics 49 (6)*, 1131–1159.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (1996). The effects of monetary policy shocks: Evidence from the flow of funds. *Review of Economics and Statistics 78 (1)*, 16–34.

Dedola, L. and S. Neri (2007). What does a technology shock do? A VAR analysis with model-based sign restrictions. *Journal of Monetary Economics 54 (2)*, 512–549.

Faust, J. (1998). The robustness of identified VAR conclusions about money. *Carnegie–Rochester Conference Series on Public Policy 49*, 207–244.

Gambetti, L. and A. Musso (2017). Loan supply shocks and business cycle. *Journal of Applied Econometrics 32 (4)*, 764–782.

Gertler, M. and P. Karadi (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics 7 (1)*, 44–76.

Giacomini, R. and T. Kitagawa (2021). Robust Bayesian Inference for Set-Identified Models. *Econometrica 89*(4), 1519–1556.

Hristov, N., O. Hülsewig, and T. T. Wollmershäuser (2012). Loan supply shocks during the financial crisis: Evidence for the Euro area. *Journal of International Money and Finance 31 (3)*, 569–592.

Jarociński, M. and P. Karadi (2020). Deconstructing monetary policy surprises—the role of information shocks. *American Economic Journal: Macroeconomics 12 (2)*, 1–43.

Känzig, D. R. (2021). The macroeconomic effects of oil supply news: Evidence from OPEC announcements. *American Economic Review 111 (4)*, 1092–1125.

Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review 99 (3)*, 1053–1069.

Kilian, L. and D. P. Murphy (2012). Why agnostic sign restrictions are not enough: Understanding the dynamics of oil market VAR models. *Journal of the European Economic Association 10 (5)*, 1166–1188.

Kim, M. (2020). How the financial market can dampen the effects of commodity price shocks. *European Economic Review 121*, 103340.

Lippi, F. and A. Nobili (2012). Oil and the macroeconomy: A quantitative structural analysis. *Journal of the European Economic Association 10 (5)*, 1059–1083.

Mountford, A. and H. Uhlig (2009). What are the effects of fiscal policy shocks? *Journal of Applied Econometrics 24 (6)*, 960–992.

Pappa, E. (2009). The effects of fiscal shocks on employment and the real wage. *International Economic Review 50 (1)*, 217–244.

Peersman, G. and R. Straub (2009). Technology shocks and robust sign restrictions in a euro area SVAR. *International Economic Review 50 (3)*, 727–750.

Peersman, G. and I. Van Robays (2012). Cross-country differences in the effects of oil shocks. *Energy Economics 34*, 1532–1547.

Rubio-Ramírez, J. F., D. F. Waggoner, and T. Zha (2010). Structural vector autoregressions: Theory of identification and algorithms for inference. *Review of Economic Studies 77 (2)*, 665–696.

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

Online Appendix for

"Restrictions on Conditional Correlations in SVARs"

Myunghyun Kim[*] Sunho Lee[†]

Not for Publication

This appendix contains:

A. Results under the strong correlation restriction

B. Adding zero restrictions to the SVAR for monetary policy shocks

C. Adding narrative sign restrictions to the SVAR for oil supply shocks

D. Results under robust Bayesian inference

[*] Department of Economics, Sungkyunkwan University, 25-2, Sungkyunkwan-ro, Jongno-gu, Seoul, 03063, Korea (e-mail: mhkim7812@gmail.com).

[†] Office of Economic Modeling & Policy Analysis, Bank of Korea, 39, Namdaemun-ro, Jung-gu, Seoul, 04531, Korea (e-mail: econ.preference@gmail.com).

A Results under the strong correlation restriction

Here, we consider a strong correlation restriction. Specifically, instead of Correlation Restriction 1, we impose the following correlation restriction for the SVAR for monetary policy shocks:

Correlation Restriction 1.1: *The correlation between real GDP (output) and the GDP deflator (prices) conditional on monetary policy shocks must be larger than 0.2.*

Figure A.1 shows the IRFs of the six variables to a contractionary monetary policy shock with Sign Restriction 1 and with both Sign Restriction 1 and Correlation Restriction 1.1.[1] The lighter shaded area and dashed lines are the 68 percent point-wise credible sets and the median IRFs, respectively, using only Sign Restriction 1. The darker shaded area and solid lines show the equivalent quantities when Sign Restriction 1 and Correlation Restriction 1.1 are used. In both cases, the monetary policy shock is normalized to raise the federal funds rate by 25 basis points in period 0. The correlations between real GDP and the GDP deflator conditional on monetary policy shocks for each draw of the structural parameters are computed from 200 simulated series of real GDP and the GDP deflator of 250 periods each (50 periods as burn-in). The posterior median of the correlation is -0.14, with a 68 percent credible interval of [-0.82, 0.56], when only Sign Restriction 1 is used. In contrast, when Sign Restriction 1 and Correlation Restriction 1.1 are used, it is 0.54, with a 68 percent credible interval of [0.32, 0.80].

The IRFs of all six variables are almost identical to those in Figure 1 in the main text. Therefore, the correlation restriction, regardless of whether it is weak or strong, sharpens the identification of monetary policy shocks and helps the price puzzle to disappear.

Then, we consider a strong correlation restriction in the SVAR for oil supply shocks. Instead of Correlation Restriction 2, we impose the following correlation restriction:

Correlation Restriction 2.1: *The correlation between real GDP (output) and CPI (prices) conditional on oil supply shocks must be smaller than -0.2.*

Figure A.2 presents the IRFs of the four variables to an oil supply shock with Sign Restriction 2 and with both Sign Restriction 2 and Correlation Restriction 2.1.[2] The lighter shaded area and dashed lines are the 68 percent point-wise credible sets and the median

[1] The IRFs are computed from 6,552 draws that satisfy Sign Restriction 1 and 2,287 draws that satisfy Sign Restriction 1 and Correlation Restriction 1.1.

[2] The IRFs are computed from 3,000 draws that satisfy Sign Restriction 2 and 1,811 draws that satisfy Sign Restriction 2 and Correlation Restriction 2.1.

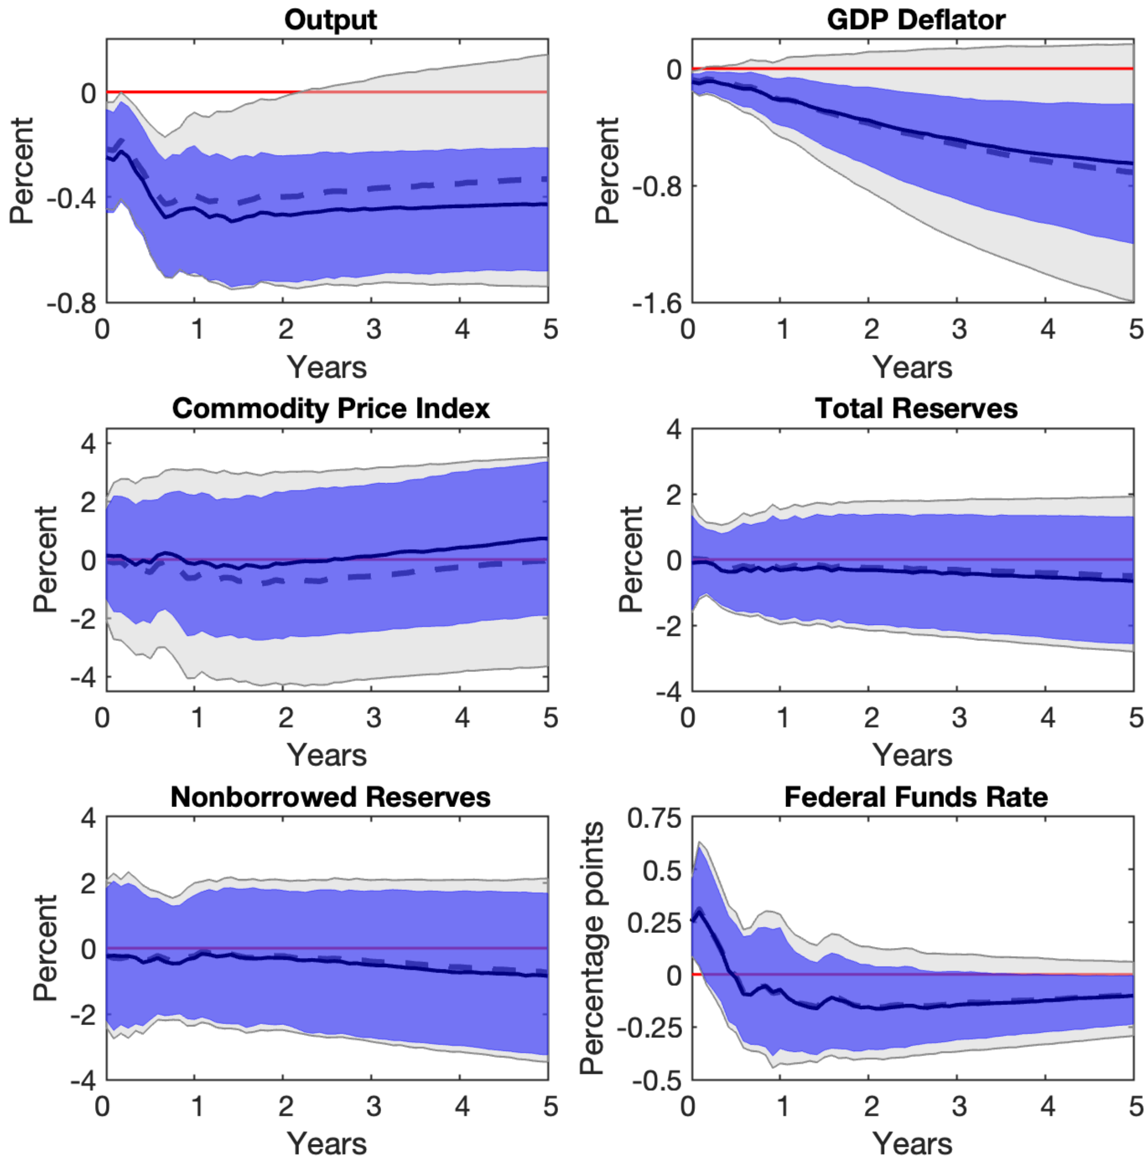

Figure A.1: IRFs to a contractionary monetary policy shock (`Correlation Restriction` 1.1)

Notes: The lighter shaded area and dashed lines are the 68 percent point-wise credible sets and the median IRFs, respectively, using `Sign Restriction` 1. The darker shaded area and solid lines are the 68 percent point-wise credible sets and the median IRFs, respectively, using `Sign Restriction` 1 and `Correlation Restriction` 1.1. The IRFs of all variables, except for the federal funds rate, are accumulated to the level.

IRFs with `Sign Restriction` 2. The darker shaded area and solid lines show the equivalent quantities with `Sign Restriction` 2 and `Correlation Restriction` 2.1. The oil supply shock is normalized to raise the real oil price by a maximum of 10 percent. The correlations between real GDP and CPI conditional on oil supply shocks for each draw of the structural parameters are computed from 200 simulated series of real GDP and CPI of 250 periods each (50 periods as burn-in). The posterior median of the correlation is -0.81, with a 68 percent credible interval of [-0.99, 0.68], when only `Sign Restriction` 2 is imposed. In contrast, when `Sign Restriction` 2 and `Correlation Restriction` 2.1 are imposed, it is -0.96, with

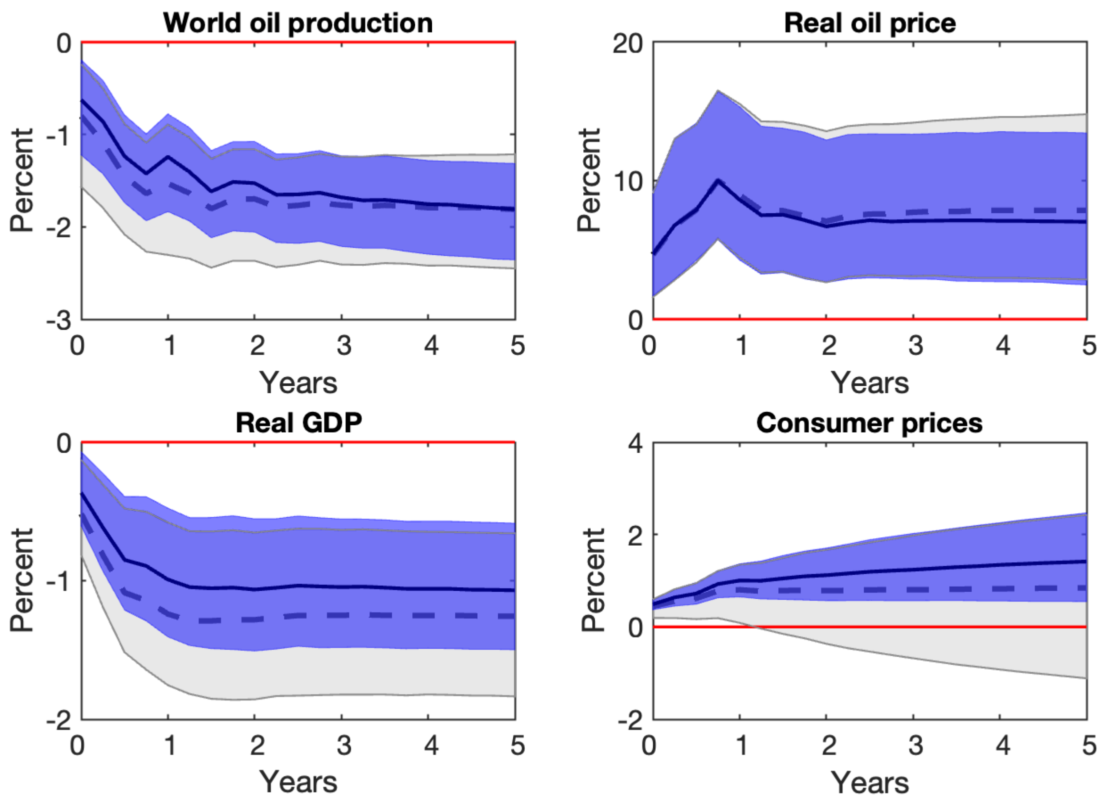

Figure A.2: IRFs to an oil supply shock (`Correlation Restriction` 2.1)

Notes: The lighter shaded area and dashed lines are the 68 percent point-wise credible sets and the median IRFs, respectively, using Sign Restriction 2. The darker shaded area and solid lines are the 68 percent point-wise credible sets and the median IRFs, respectively, using Sign Restriction 2 and Correlation Restriction 2.1. The IRFs of all variables are accumulated to the level.

a 68 percent credible interval of [-0.99, -0.80].

There are no significant differences between the IRFs in Figure A.2 and Figure 2 in the main text. Hence, regardless of whether it is strong or weak, the correlation restriction is very useful in identifying oil supply shocks and enables prices to increase significantly in response to an oil supply shock, consistent with economic theory.

B Adding zero restrictions to the SVAR for monetary policy shocks

In this appendix, we add the zero restrictions used in Arias, Caldara, and Rubio-Ramírez (2019) to the SVAR for monetary policy shocks in the main text. Specifically, the following zero restrictions are imposed:

Zero Restriction 1: *The federal funds rate only reacts contemporaneously to output, prices, and commodity prices.*

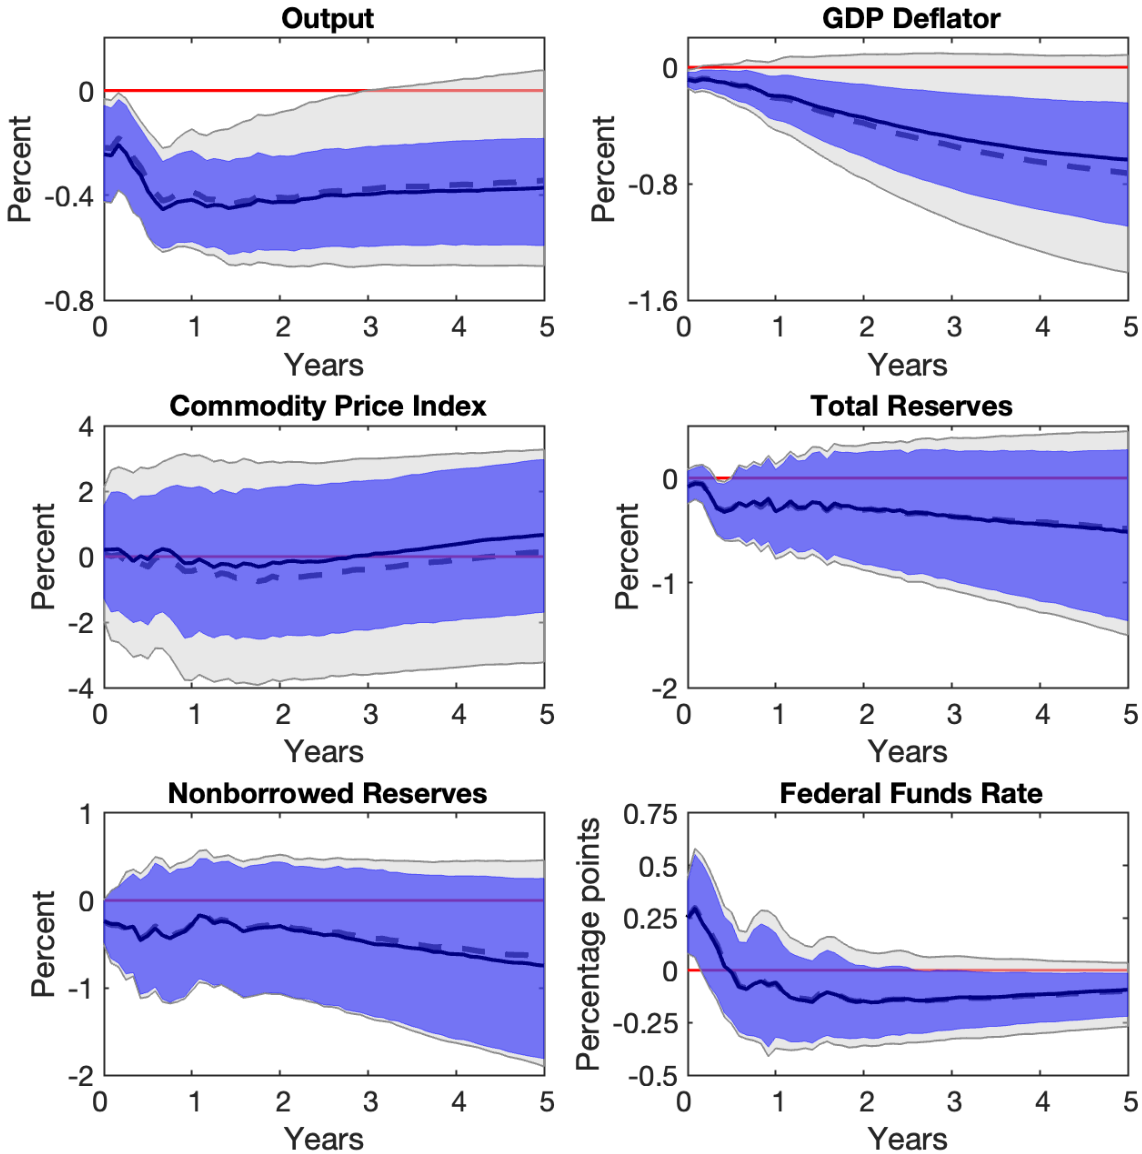

Figure B: IRFs to a contractionary monetary policy shock (`Zero Restriction` 1)

Notes: The lighter shaded area and dashed lines are the 68 percent point-wise credible sets and the median IRFs, respectively, using `Sign Restriction` 1 and `Zero Restriction` 1. The darker shaded area and solid lines are the 68 percent point-wise credible sets and the median IRFs, respectively, using `Sign Restriction` 1, `Zero Restriction` 1, and `Correlation Restriction` 1. The IRFs of all variables, except for the federal funds rate, are accumulated to the level.

`Zero Restriction` 1 implies that the federal funds rate does not react to total and nonborrowed reserves, which means that the two contemporaneous reserves are not included in the equation for the federal funds rate.

Figure B shows the IRFs of the six variables to a contractionary monetary policy shock with `Sign Restriction` 1 and `Zero Restriction` 1, and with the two restrictions and `Correlation Restriction` 1.[3] The lighter shaded area and dashed lines are the 68 percent

[3] The IRFs are computed from 7,018 draws that satisfy `Sign Restriction` 1 and `Zero Restriction` 1 and 3,435 draws that satisfy the two restrictions and `Correlation Restriction` 1.

point-wise credible sets and the median IRFs, respectively, using only `Sign Restriction` 1 and `Zero Restriction` 1. The darker shaded area and solid lines show the equivalent quantities when `Sign Restriction 1`, `Zero Restriction` 1, and `Correlation Restriction` 1 are used. In both cases, the monetary policy shock is normalized to raise the federal funds rate by 25 basis points in period 0. The correlations between real GDP and the GDP deflator conditional on monetary policy shocks for each draw of the structural parameters are computed from 200 simulated series of real GDP and the GDP deflator of 250 periods each (50 periods as burn-in). The posterior median of the correlation is -0.03, with a 68 percent credible interval of [-0.79, 0.60], when `Sign Restriction` 1 and `Zero Restriction` 1 are used. In contrast, when `Sign Restriction` 1, `Zero Restriction` 1, and `Correlation Restriction` 1 are used, it is 0.47, with a 68 percent credible interval of [0.18, 0.74].

The overall contour of the IRFs is very similar to those in Figure 1 in the main text. As Figure B clearly shows, the correlation restriction sharpens the identification of monetary policy shocks and allows the price puzzle to disappear, even when `Zero Restriction` 1 is additionally imposed.

C Adding narrative sign restrictions to the SVAR for oil supply shocks

In this appendix, we add the narrative sign restrictions used in Antolín-Díaz and Rubio-Ramírez (2018) to the SVAR for oil supply shocks in the main text. Specifically, the following narrative sign restrictions are imposed:

Narrative Sign Restriction 1: *The oil supply shock must take positive values in Q1 1979, Q4 1980, Q3 1990, Q1 2003, and Q1 2011.*

`Narrative Sign Restriction` 1 is slightly different from those in Antolín-Díaz and Rubio-Ramírez (2018). This is mainly because a monthly SVAR is estimated in Antolín-Díaz and Rubio-Ramírez (2018), while we estimate a quarterly SVAR. The associated events are the Iranian Revolution (December 1978–January 1979), the Iran-Iraq War (September–October 1980), the Persian Gulf War (August 1990), the Venezuela oil strike of December 2002, the start of the Iraq War (March 2003), and the Libyan Civil War (February 2011). Since we estimate a quarterly SVAR, for an event which occurs in the last month of a certain quarter, we impose a narrative sign restriction in the next quarter. For example, for the Iranian Revolution, which occurs December 1978–January 1979, we impose a narrative sign restriction in Q1 1979. Note also that since we identify only oil supply shocks, we do not impose narrative sign restrictions on the contribution of a certain shock in the quarter when the relevant event

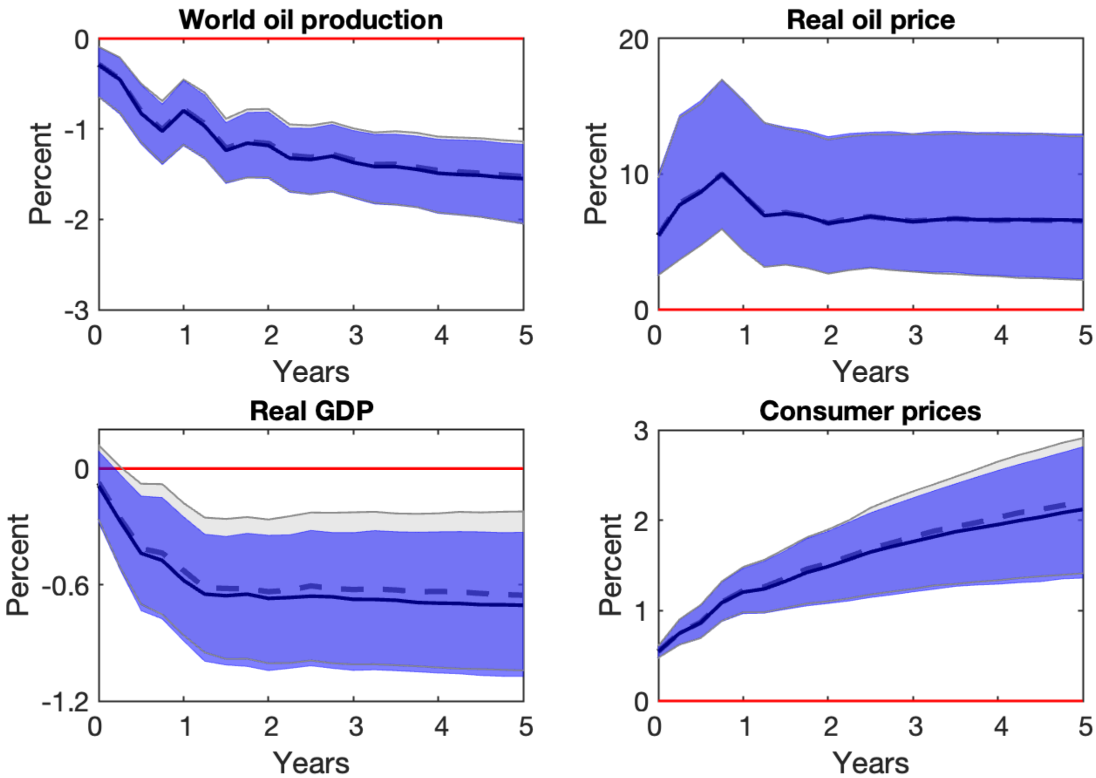

Figure C: IRFs to an oil supply shock (`Narrative Sign Restriction` 1)

Notes: The lighter shaded area and dashed lines are the 68 percent point-wise credible sets and the median IRFs, respectively, using `Sign Restriction` 2 and `Narrative Sign Restriction` 1. The darker shaded area and solid lines are the 68 percent point-wise credible sets and the median IRFs, respectively, using `Sign Restriction` 2, `Narrative Sign Restriction` 1, and `Correlation Restriction` 2. The IRFs of all variables are accumulated to the level.

occurs. That is, we do not impose restrictions on the historical decomposition around relevant events.

Figure C presents the IRFs of the four variables to an oil supply shock with `Sign Restriction` 2 and `Narrative Sign Restriction` 1, and with the two restrictions and `Correlation Restriction` 2.[4] The lighter shaded area and dashed lines are the 68 percent point-wise credible sets and the median IRFs with `Sign Restriction` 2 and `Narrative Sign Restriction` 1. The darker shaded area and solid lines show the equivalent quantities with `Sign Restriction` 2, `Narrative Sign Restriction` 1, and `Correlation Restriction` 2. The oil supply shock is normalized to raise the real oil price by a maximum of 10 percent. The correlations between real GDP and CPI conditional on oil supply shocks for each draw of the structural parameters are computed from 200 simulated series of real GDP and CPI of 250 periods each (50 periods as burn-in). The posterior median of the correlation is -0.95, with a 68 percent credible interval of [-0.99, -0.46], when `Sign Restriction` 2 and `Narrative`

[4] The IRFs are computed from 1,562 draws that satisfy `Sign Restriction` 2 and `Narrative Sign Restriction` 1, and 1,391 draws that satisfy `Sign Restriction` 2, `Narrative Sign Restriction` 1, and `Correlation Restriction` 2.

`Sign Restriction` 1 are imposed. In contrast, when `Sign Restriction` 2, `Narrative Sign Restriction` 1, and `Correlation Restriction` 2 are imposed, it is -0.96, with a 68 percent credible interval of [-0.99, -0.77].

There are no notable differences between the median IRFs of the four variables in the two cases, and between their credible sets in the two cases. The credible sets for the IRFs of output and prices are only slightly tighter when `Correlation Restriction` 2 is imposed. In other words, `Narrative Sign Restriction` 1 is as helpful as `Correlation Restriction` 2 when identifying oil supply shocks.

D Results under robust Bayesian inference

In set-identified SVARs, a given reduced form can be associated with multiple structural paths that satisfy the identifying restrictions and achieve the same likelihood value. As emphasized by Baumeister and Hamilton (2024), the conventional practice based on a uniform prior over the rotation matrix can place some admissible paths inside a reported credible set and others outside it, even though the likelihood does not distinguish among them. One way to address this issue is to specify and defend an explicit prior over the rotation matrix when ranking admissible structural paths. Another way is to use the robust Bayesian approach of Giacomini and Kitagawa (2021), which avoids conditioning on a particular prior over rotations. Since the SVARs in this paper remain set-identified after imposing the correlation restriction, it is useful to examine, within this framework, whether the correlation restriction adds identifying content beyond the sign restrictions. We therefore compare the results under the sign restrictions only with those under the sign and correlation restrictions.

For each reduced-form posterior draw $(\mathbf{B}, \mathbf{\Sigma})$, let $\mathcal{Q}_R(\mathbf{B}, \mathbf{\Sigma})$ be the set of rotations that satisfy a restriction set R, such as sign or correlation restrictions. For the shock of interest k, variable i, and horizon h, define

$$\ell^R_{i,h}(\mathbf{B}, \mathbf{\Sigma}) = \inf_{\mathbf{Q} \in \mathcal{Q}_R(\mathbf{B}, \mathbf{\Sigma})} \mathbf{IR}^{(i,k)}_h(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) \tag{D.1}$$

$$\text{and } u^R_{i,h}(\mathbf{B}, \mathbf{\Sigma}) = \sup_{\mathbf{Q} \in \mathcal{Q}_R(\mathbf{B}, \mathbf{\Sigma})} \mathbf{IR}^{(i,k)}_h(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}). \tag{D.2}$$

The set of posterior means is

$$\left[E_{\mathbf{B}, \mathbf{\Sigma}|Y} \left[\ell^R_{i,h}(\mathbf{B}, \mathbf{\Sigma}) \right], \ E_{\mathbf{B}, \mathbf{\Sigma}|Y} \left[u^R_{i,h}(\mathbf{B}, \mathbf{\Sigma}) \right] \right]. \tag{D.3}$$

This interval gives the lower and upper bounds on the posterior mean IRF that can be attained by rotations satisfying the restriction set R, and it is the object reported in this section.

We approximate the bounds by simulation. For each reduced-form draw, we draw a fixed number of Haar rotations and retain those that satisfy the restrictions. We then compute the minimum and maximum IRFs over the retained rotations, and average these bounds across reduced-form draws.[5] This follows the simulation approach in Giacomini and Kitagawa (2021).

For the monetary policy application, we retain 1,000 reduced-form draws and use 10,000 rotations for each draw. For all draws, $\mathcal{Q}_R(\mathbf{B}, \mathbf{\Sigma})$ is nonempty under both the sign restrictions and the sign and correlation restrictions. Among the rotations satisfying the sign restrictions, 54.2% do not satisfy the correlation restriction. For the oil supply application, we process 3,300 reduced-form draws and use 150,000 rotations for each draw. The set $\mathcal{Q}_R(\mathbf{B}, \mathbf{\Sigma})$ is nonempty for 1,178 draws under the sign restrictions and 1,007 draws under the sign and correlation restrictions. Among the rotations satisfying the sign restrictions, 38.0% do not satisfy the correlation restriction. In addition, 14.5% of the reduced-form draws with a nonempty sign-restricted rotation set have no admissible rotation once the correlation restriction is added. These results show that the correlation restriction substantially reduces the set of rotations admitted by the sign restrictions, and in some cases also adds identifying content for the reduced-form parameters.

Figure D.1 reports the results for monetary policy shocks. Adding the correlation restriction narrows the set of posterior means. The reduction is clear for output and the GDP deflator. For output, the correlation restriction removes much of the positive response allowed under the sign restrictions only, making the contractionary output response more tightly identified. For the GDP deflator, it also removes much of the positive response, indicating that the price-puzzle pattern is substantially weakened.

Figure D.2 reports the results for oil supply shocks. The main difference appears in consumer prices. Under the sign restrictions only, the set of posterior means includes negative price responses at longer horizons. Once the correlation restriction is imposed, the lower bound remains positive throughout the horizon. This supports the main result that the correlation restriction helps identify oil supply shocks in a way that is consistent with the supply shock interpretation. Taken together, the monetary policy and oil supply results show that the correlation restriction adds identifying content beyond the sign restrictions and sharpens the IRFs.

[5] A reduced-form draw is excluded from the calculation of the set of posterior means if 10 or fewer admissible rotations are found among the finite draws of $\mathbf{Q}$. This is done only to avoid unstable approximations of the minimum and maximum based on too few admissible rotations.

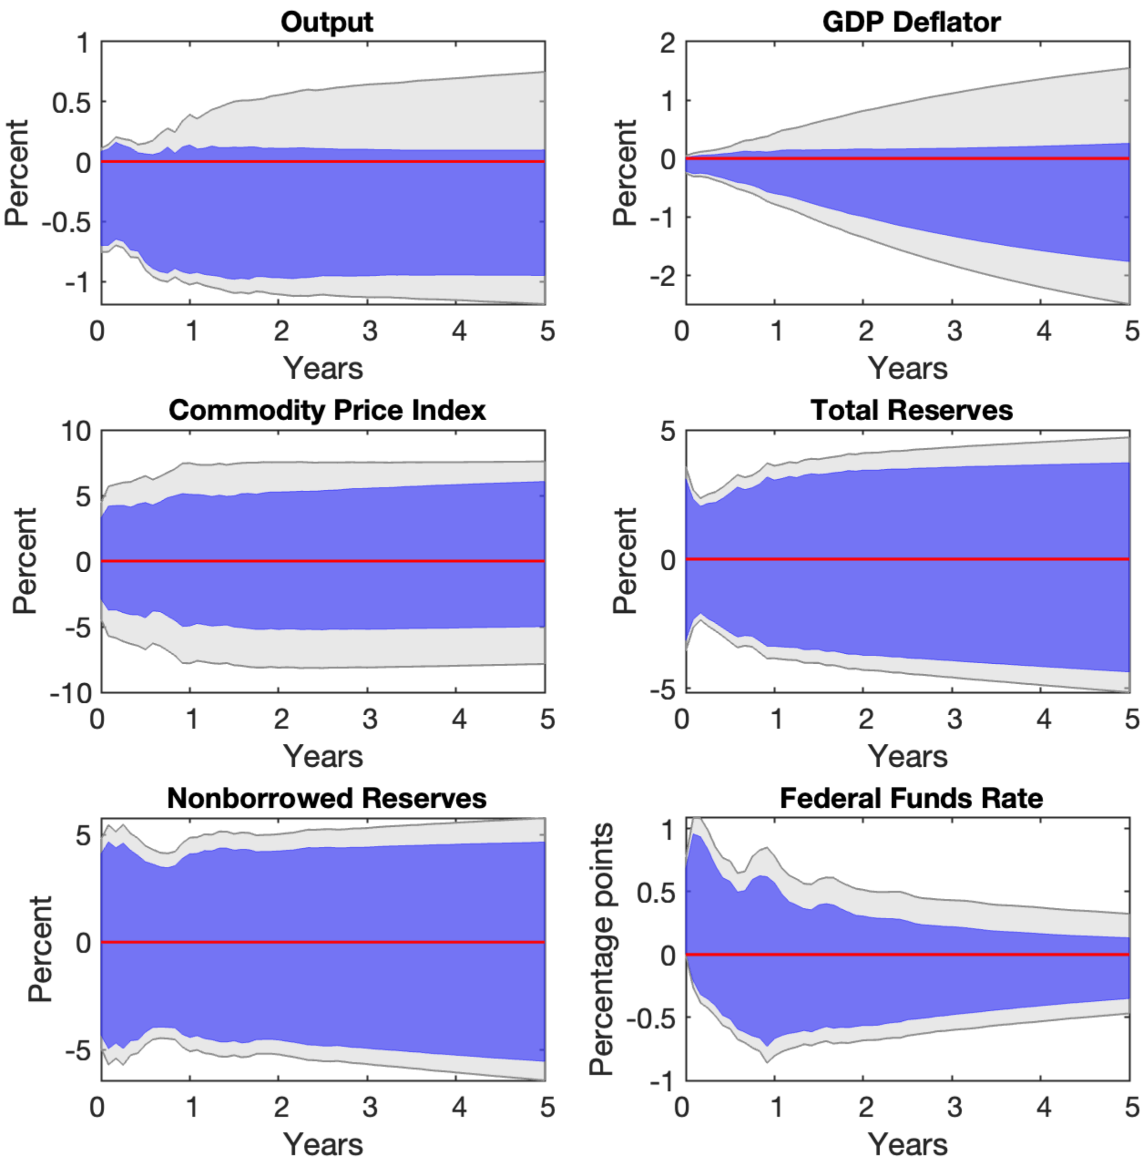

Figure D.1: Set of posterior means for a contractionary monetary policy shock

Notes: The lighter shaded area is the set of posterior means using `Sign Restriction` 1. The darker blue shaded area is the set of posterior means using `Sign Restriction` 1 and `Correlation Restriction` 1. The IRFs of all variables, except for the federal funds rate, are accumulated to the level.

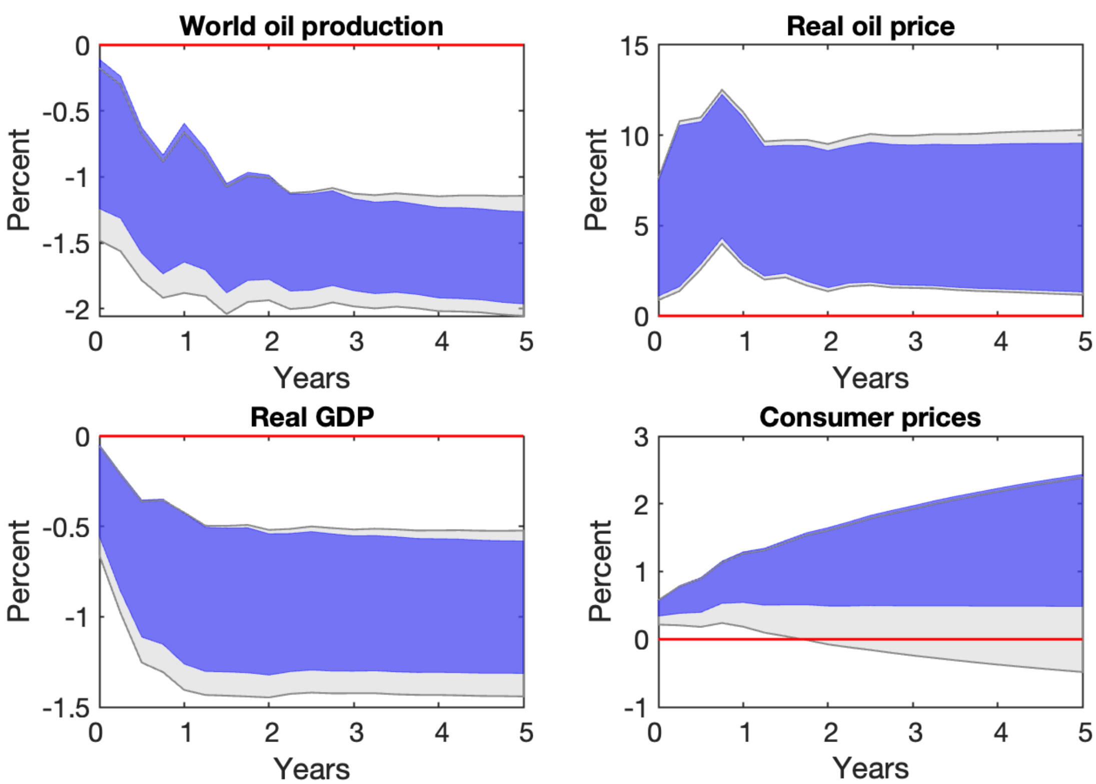

Figure D.2: Set of posterior means for an oil supply shock

Notes: The lighter shaded area is the set of posterior means using `Sign Restriction` 2. The darker blue shaded area is the set of posterior means using `Sign Restriction` 2 and `Correlation Restriction` 2. The IRFs of all variables are accumulated to the level.