FEDERAL RESERVE BANK
PHILADELPHIA

Working Papers
RESEARCH DEPARTMENT

Large SVARs

Jonas E. Arias
Federal Reserve Bank of Philadelphia

Juan Rubio-Ramírez
Emory University and Federal Reserve Bank of Atlanta

Minchul Shin
Federal Reserve Bank of Philadelphia

WP 26-04

PUBLISHED
January 2026

ISSN: 1962-5361

DOI: https://doi.org/10.21799/frbp.wp.2026.04

Large SVARs

Jonas E. Arias

Federal Reserve Bank of Philadelphia

Juan Rubio-Ramírez*

Emory University

Federal Reserve Bank of Atlanta

Minchul Shin

Federal Reserve Bank of Philadelphia

January 11, 2026

Abstract

We develop a new algorithm for inference in structural vector autoregressions (SVARs) identified with sign restrictions that can accommodate big data and modern identification schemes. The key innovation of our approach is to move beyond the traditional accept-reject framework commonly used in sign-identified SVARs. We show that embedding the elliptical slice sampling within a Gibbs sampler can deliver dramatic gains in computational speed and render previously infeasible applications tractable. To illustrate the approach in the context of sign-identified SVARs, we use a tractable example. We further assess the performance of our algorithm through two applications: a well-known small-SVAR model of the oil market featuring a tight identified set, and a large SVAR model with more than ten shocks and 100 sign restrictions.

JEL classification: C11, C15, C32.

Keywords: large structural vector autoregressions, sign restrictions, elliptical slice sampling.

We thank Frank Schorfheide for helpful comments. The views expressed in this paper are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Atlanta, the Federal Reserve Bank of Philadelphia, or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

*Corresponding author: Juan F. Rubio-Ramírez <juan.rubio-ramirez@emory.edu>, Economics Department, Emory University, Rich Memorial Building, Room 306, Atlanta, Georgia 30322-2240.

1 Introduction

The growing availability of large datasets has led to a renewed interest in the use of large-scale time-series models in economics. In univariate settings, Giannone, Lenza, and Primiceri (2021) show that densely parameterized models equipped with appropriate shrinkage priors typically outperform sparse alternatives in terms of predictive accuracy. In multivariate settings, Bańbura, Giannone, and Reichlin (2010) and Koop (2013) have demonstrated that Bayesian shrinkage enables the estimation of large vector autoregressions (VARs) without compromising out-of-sample performance, and more recently, Crump, Eusepi, Giannone, Qian, and Sbordone (2025) show how this approach can be used for policy analysis. These insights have direct implications for structural vector autoregressions (SVARs)—one of the major workhorses for studying the propagation of structural shocks in macroeconomics. If prediction tasks are better handled by large models, then inference on structural shocks using SVARs should similarly benefit from broader information sets, in line with early arguments in favor of using large SVARs to understand the macroeconomic effects of monetary policy (e.g., Leeper, Sims, and Zha, 1996; Bernanke, Boivin, and Eliasz, 2005).

Within the SVAR paradigm, sign restrictions have become a particularly popular method for identifying the parameters of interest, typically impulse responses. The conventional Bayesian approach to implementing sign restrictions—pioneered by Faust (1998), Canova and De Nicoló (2002), and Uhlig (2005) and extended by Rubio-Ramírez, Waggoner, and Zha (2010)—relies on sampling from the reduced-form posterior and a uniform prior over the set of orthogonal matrices, combined with an accept-reject approach to impose the restrictions. The method is straightforward to implement and produces independent draws. However, it becomes increasingly infeasible in large systems due to the vanishing probability of sampling admissible orthogonal matrices—especially as the number of sign restrictions needed to identify multiple structural shocks increases, tightening the identified set.[1] Recent

[1]We use the term admissible to refer to orthogonal matrices that satisfy the sign restrictions given the reduced-form parameters.

work by Chan, Matthes, and Yu (2025) introduces algorithmic refinements by exploiting symmetry and permutation invariance in the space of orthogonal matrices, but even these improved methods face limitations under tight identification. Notably, such computational burdens are not unique to high-dimensional models. These challenges have coincided with advances in identification strategies that also lead to tighter identified sets—even in low-dimensional SVARs—such as ranking restrictions and elasticity bounds (Kilian and Murphy, 2012; Amir-Ahmadi and Drautzburg, 2021), as well as narrative sign restrictions (Antolín-Díaz and Rubio-Ramírez, 2018; Ludvigson, Ma, and Ng, 2021). These approaches go beyond traditional sign restrictions while preserving their intuitive appeal and further reducing the volume of the admissible space of orthogonal matrices. Together, these trends—the adoption of larger information sets and the use of tighter identification strategies—and the limitations of accept-reject sampling methods underscore the need for alternative algorithms for Bayesian inference in SVARs identified with sign restrictions.

In this paper, we break with the accept-reject tradition and show that embedding the elliptical slice sampling method of Murray, Adams, and MacKay (2010) within a Gibbs sampler delivers substantial gains in computational speed, rendering previously infeasible applications tractable. Like in the conventional approach, using the uniform prior over the set of orthogonal matrices, our goal is to draw from the posterior distribution of the orthogonal reduced-form parameters conditional on sign restrictions. However, by eliminating the accept-reject step and directly conditioning on the sign restrictions within the Gibbs sampler, our algorithm overcomes the bottlenecks that arise under tight identification—thus enabling dynamic structural analysis with big data and rich identification schemes. To illustrate the advantages of our approach, we consider a very simple example similar to the one in Granziera, Moon, and Schorfheide (2018) and demonstrate that the efficiency of the accept-reject algorithm hinges critically on the size of the identified set. As the identified set becomes tighter, the accept-reject algorithm slows down dramatically. In contrast, our Gibbs sampler efficiently shrinks the support of the candidate impulse responses toward the

identified set, maintaining speed even under stringent restrictions.

For clarity and comparability with the literature, when describing our proposed Gibbs sampler algorithm, we adopt the conjugate normal-inverse-Wishart prior for the reduced-form parameters as our baseline. While this prior is popular due to its analytical convenience, it precludes cross-variable shrinkage. To address this limitation, we extend our algorithm to accommodate alternative priors, including the independent normal-inverse-Wishart and the asymmetric prior of Chan (2022), both of which support cross-variable shrinkage. We evaluate the performance of our approach using two applications. In the first, we replicate Kilian and Murphy (2014), a model of the world oil market in which the standard accept-reject algorithm fails. To address this infeasibility, Kilian and Murphy (2014) adopt an approach similar to that of Chan, Matthes, and Yu (2025), exploiting permutations and sign alternations. Our algorithm handles this application multiple times faster than the accept-reject approach, though the computational times of both approaches are within a range most practitioners would find acceptable. However, once we tighten the identified set by adding a restriction on the price elasticity of oil demand—motivated by Caldara, Cavallo, and Iacoviello (2019)—the difference in performance becomes substantial: The accept-reject algorithm moves from requiring about 20 minutes to produce 1,000 draws to nearly eight hours, whereas the computational time of our Gibbs sampler remains roughly constant, increasing only from about 2 minutes to 5 minutes for the same number of effective draws. In the second application, we revisit the structural analysis in Chan, Matthes, and Yu (2025), who use Crump et al.'s (2025) large SVAR model of the U.S. economy to identify eight structural shocks. While the latter uses the Minnesota prior, the former relies on the asymmetric prior of Chan (2022). To simplify the comparison, we revert to the Minnesota prior when applying both the accept-reject algorithm and the Gibbs sampler. We show that as the number of shocks under analysis increases, the efficiency of the algorithm in Chan, Matthes, and Yu (2025) declines markedly, eventually becoming impractical. With ten shocks, it would take several days to obtain 1,000 draws. In contrast, the computational time of

our Gibbs sampler is largely insensitive to the number of identified structural shocks. Even with ten shocks, it would take only a few minutes to obtain 1,000 effective draws. We also show that these striking differences are robust to using the asymmetric prior instead of the Minnesota prior.

We need to highlight the contemporaneous contribution of Read and Zhu (2025), who also propose an algorithm based on the slice sampling method but it is limited to the use of the conditionally uniform prior described in Uhlig (2017) and Amir-Ahmadi and Drautzburg (2021). While this prior can deliver substantial speed gains without the need of a Gibbs sampler, it does not satisfy the requirements set out in Arias, Rubio-Ramírez, and Waggoner (2025). As we explain in Section 8, this has important consequences: it implicitly alters the prior over the impulse responses in a way that depends on the identification scheme. This entanglement of inference and identification makes it difficult to know whether differences in posterior inference reflect genuine differences in identification or are merely artifacts of unintended changes in the prior distribution. By contrast, the uniform prior over orthogonal matrices satisfies the requirements in Arias, Rubio-Ramírez, and Waggoner (2025) and, therefore, guarantees that inference remains invariant to the set of imposed restrictions and allows researchers to cleanly separate the role of prior beliefs from the role of identification assumptions.

The remainder of the paper is organized as follows. Sections 2 through 4 introduce the SVAR model, the sign restrictions, and the baseline conjugate uniform-normal-inverse-Wishart (UNIW) prior. Section 5 describes the problem in a simple environment. Section 6 presents our Gibbs sampler featuring the elliptical slice sampling and outlines its theoretical properties. Section 7 applies the algorithm to two empirical settings: a small SVAR model of the world oil market and a large SVAR model of the U.S. economy. Section 8 shows the shortcomings of the conditionally uniform prior. Section 9 concludes. The Appendix adapts the algorithm to two popular priors in SVAR analysis, the independent normal-inverse-Wishart prior and the asymmetric prior proposed by Chan (2022), and it does some robustness analysis.

2 The Model

Consider the SVAR with the general form,

$$\mathbf{y}_t' \mathbf{A}_0 = \mathbf{x}_t' \mathbf{A}_+ \boldsymbol{\varepsilon}_t', \quad 1 \le t \le T, \tag{1}$$

where $\mathbf{A}_+' = \begin{bmatrix} \mathbf{A}_1' & \cdots & \mathbf{A}_p' & \mathbf{c}' \end{bmatrix}$ and $\mathbf{x}_t' = \begin{bmatrix} \mathbf{y}_{t-1}' & \cdots & \mathbf{y}_{t-p}' & 1 \end{bmatrix}$ for $1 \le t \le T$, and where $\mathbf{y}_t$ is an $n \times 1$ vector of endogenous variables, $\boldsymbol{\varepsilon}_t$ is an $n \times 1$ vector of exogenous structural shocks, $\mathbf{A}_\ell$ is an $n \times n$ matrix of parameters for $0 \le \ell \le p$ with $\mathbf{A}_0$ invertible, $\mathbf{c}$ is a $1 \times n$ vector of parameters, p is the lag length, and T is the sample size. Hence, the dimension of $\mathbf{A}_+$ is $m \times n$, where $m = np + 1$. The vector $\boldsymbol{\varepsilon}_t$, conditional on past information and the initial conditions $\mathbf{y}_0, \ldots, \mathbf{y}_{1-p}$, is Gaussian with mean zero and covariance matrix $\mathbf{I}_n$, the $n \times n$ identity matrix.

The reduced-form representation implied by Equation (1) is

$$\mathbf{y}_t' = \mathbf{x}_t' \mathbf{B} + \mathbf{u}_t', \text{ for } 1 \le t \le T, \tag{2}$$

where $\mathbf{B} = \mathbf{A}_+ \mathbf{A}_0^{-1}$, $\mathbf{u}_t' = \boldsymbol{\varepsilon}_t' \mathbf{A}_0^{-1}$, and $\mathbb{E}\left[\mathbf{u}_t \mathbf{u}_t'\right] = \boldsymbol{\Sigma} = \left(\mathbf{A}_0 \mathbf{A}_0'\right)^{-1}$. The matrices $\mathbf{B}$ and $\boldsymbol{\Sigma}$ are the reduced-form parameters, while $\mathbf{A}_0$ and $\mathbf{A}_+$ are the structural parameters. While $\mathbf{B}$ is an $m \times n$ matrix, $\boldsymbol{\Sigma}$ belongs to the set $\mathcal{S}(n)$, which is the set of $n \times n$ positive definite matrices. It will be useful to partition $\mathbf{B}$ as follows: $\mathbf{B} = \begin{bmatrix} \mathbf{B}_1' & \cdots & \mathbf{B}_p' & \mathbf{d}' \end{bmatrix}'$ where $\mathbf{B}_\ell$ is an $n \times n$ matrix of parameters for $1 \le \ell \le p$, and $\mathbf{d}$ is a $1 \times n$ vector of parameters.

It is well known that for linear Gaussian models of the type studied in this paper, $(\mathbf{A}_0, \mathbf{A}_+)$ and $(\tilde{\mathbf{A}}_0, \tilde{\mathbf{A}}_+)$ are observationally equivalent if and only if they have the same reduced-form representation. This implies that the structural parameters $(\mathbf{A}_0, \mathbf{A}_+)$ and $(\tilde{\mathbf{A}}_0, \tilde{\mathbf{A}}_+)$ are observationally equivalent if and only if $\mathbf{A}_0 = \tilde{\mathbf{A}}_0 \mathbf{Q}$ and $\mathbf{A}_+ = \tilde{\mathbf{A}}_+ \mathbf{Q}$ for some $\mathbf{Q} \in \mathcal{O}(n)$, where $\mathcal{O}(n)$ is the set of all $n \times n$ orthogonal matrices. To solve the identification problem, one often imposes sign restrictions on either the structural parameters or some function of the structural parameters, such as the impulse responses. To simplify the notation, we summarize

the sign restrictions by $\boldsymbol{S}_S(\mathbf{A}_0, \mathbf{A}_+) > 0$, and let $[\boldsymbol{S}_S(\mathbf{A}_0, \mathbf{A}_+) > 0]$ be an indicator function that equals one if the sign restrictions are satisfied and zero otherwise.

3 The Orthogonal Reduced-Form Parameterization

Equation (1) represents the SVAR in terms of the structural parameterization, which is characterized by $(\mathbf{A}_0, \mathbf{A}_+)$. The SVAR can alternatively be written in what we call the orthogonal reduced-form parameterization; see Arias, Rubio-Ramírez, and Waggoner (2018). This parameterization is characterized by the reduced-form parameters $(\mathbf{B}, \mathbf{\Sigma})$ together with an orthogonal matrix $\mathbf{Q}$, and is given by the following equation:

$$\mathbf{y}_t' = \mathbf{x}_t' \mathbf{B} + \boldsymbol{\varepsilon}_t' \mathbf{Q}' h(\mathbf{\Sigma}), \text{ for } 1 \leq t \leq T, \tag{3}$$

where the $n \times n$ matrix $h(\mathbf{\Sigma})$ is any decomposition of the covariance matrix $\mathbf{\Sigma}$ satisfying $h(\mathbf{\Sigma})' h(\mathbf{\Sigma}) = \mathbf{\Sigma}$. We take h to be the Cholesky decomposition, though any differentiable decomposition would suffice.

Given Equations (1) and (3), we can define a mapping between $(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})$ and $(\mathbf{A}_0, \mathbf{A}_+)$ by

$$f(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) = \big(\underbrace{h(\mathbf{\Sigma})^{-1} \mathbf{Q}}_{\mathbf{A}_0}, \underbrace{\mathbf{B}\, h(\mathbf{\Sigma})^{-1} \mathbf{Q}}_{\mathbf{A}_+} \big).$$

This mapping makes clear how the structural parameters depend on the reduced-form parameters and orthogonal matrices. Given the reduced-form parameters, each value of $\mathbf{Q} \in \mathcal{O}(n)$ can be viewed as a particular choice among observationally equivalent structural parameters. Thus, we can always write the sign restrictions in terms of the orthogonal reduced-form parameterization. Hence, let $[\boldsymbol{S}_R(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) > 0]$ be an indicator function in terms of the orthogonal reduced-form parameterization that equals one if the sign restrictions are satisfied and zero otherwise, where $\boldsymbol{S}_R(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) = \boldsymbol{S}_S(f(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}))$.

We can also define the impulse responses. Let $\mathbf{u}_t = \mathbf{L}_0 \boldsymbol{\varepsilon}_t$ for $1 \leq t \leq T$, where $\mathbf{L}_0$ is an

$n \times n$ invertible matrix that represents impulse responses at horizon zero. Given $\mathbf{L}_0$ and $\mathbf{B}$, it is possible to obtain the impulse responses beyond horizon zero recursively as

$$\mathbf{L}_\ell = \sum_{k=1}^{\min\{\ell,p\}} \mathbf{B}_k' \, \mathbf{L}_{\ell-k}, \quad \ell > 0. \tag{4}$$

We combine the impulse responses from horizons one through p and the constant term $\mathbf{c}$ into a single matrix, $\mathbf{L}_+ = \begin{bmatrix} \mathbf{L}_1' & \cdots & \mathbf{L}_p' & \mathbf{c}' \end{bmatrix}'$, where the maximum horizon of the impulse response in $\mathbf{L}_+$ matches the lag length in Equation (1). The impulse response parameterization is characterized by $(\mathbf{L}_0, \mathbf{L}_+)$. Given the function f and Equation (4), we can also define a mapping from $(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})$ to $(\mathbf{L}_0, \mathbf{L}_+)$ by

$$\phi(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) = \Big(\underbrace{h(\mathbf{\Sigma})' \mathbf{Q}}_{\mathbf{L}_0}, \underbrace{\begin{bmatrix} \mathbf{L}_1(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})' & \cdots & \mathbf{L}_p(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})' & \mathbf{Q}'(h(\mathbf{\Sigma})^{-1})' \, \mathbf{d}' \end{bmatrix}'}_{\mathbf{L}_+} \Big), \tag{5}$$

where $\mathbf{L}_\ell(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})$ for $1 \leq \ell \leq p$ is implicitly defined by Equation (4). The functions f and ϕ are invertible, and f, ϕ, and their inverses are differentiable.

4 Conjugate Priors and Posteriors

For the reduced-form representation in Equation (2), the normal-inverse-Wishart family of distributions is conjugate. A conjugate normal-inverse-Wishart distribution over the reduced-form parameters is characterized by four parameters: a scalar $\nu \geq n$, an $n \times n$ symmetric and positive definite matrix $\mathbf{\Phi}$, an $m \times n$ matrix $\mathbf{\Psi}$, and an $m \times m$ symmetric and positive definite matrix $\mathbf{\Omega}$. We denote this distribution by $NIW(\nu, \mathbf{\Phi}, \mathbf{\Psi}, \mathbf{\Omega})$ and its density by $NIW_{(\nu,\mathbf{\Phi},\mathbf{\Psi},\mathbf{\Omega})}(\mathbf{B}, \mathbf{\Sigma})$. Furthermore,

$$NIW_{(\nu,\mathbf{\Phi},\mathbf{\Psi},\mathbf{\Omega})}(\mathbf{B}, \mathbf{\Sigma}) \propto \underbrace{|\det(\mathbf{\Sigma})|^{-\frac{\nu+n+1}{2}} e^{-\frac{1}{2}\operatorname{tr}(\mathbf{\Phi}\mathbf{\Sigma}^{-1})}}_{\text{inverse-Wishart}} \underbrace{|\det(\mathbf{\Sigma})|^{-\frac{m}{2}} e^{-\frac{1}{2}\operatorname{vec}(\mathbf{B}-\mathbf{\Psi})'(\mathbf{\Sigma}\otimes\mathbf{\Omega})^{-1}\operatorname{vec}(\mathbf{B}-\mathbf{\Psi})}}_{\text{conditionally normal}}.$$

If the prior distribution over the reduced-form parameters is $NIW(\bar{\nu}, \bar{\mathbf{\Phi}}, \bar{\mathbf{\Psi}}, \bar{\mathbf{\Omega}})$, then the posterior distribution over the reduced-form parameters is $NIW(\tilde{\nu}, \tilde{\mathbf{\Phi}}, \tilde{\mathbf{\Psi}}, \tilde{\mathbf{\Omega}})$, where

$$\begin{aligned}
\tilde{\nu} &= T + \bar{\nu}, \\
\tilde{\mathbf{\Omega}} &= (\mathbf{X}'\mathbf{X} + \bar{\mathbf{\Omega}}^{-1})^{-1}, \\
\tilde{\mathbf{\Psi}} &= \tilde{\mathbf{\Omega}}(\mathbf{X}'\mathbf{Y} + \bar{\mathbf{\Omega}}^{-1}\bar{\mathbf{\Psi}}), \\
\tilde{\mathbf{\Phi}} &= \mathbf{Y}'\mathbf{Y} + \bar{\mathbf{\Phi}} + \bar{\mathbf{\Psi}}'\bar{\mathbf{\Omega}}^{-1}\bar{\mathbf{\Psi}} - \tilde{\mathbf{\Psi}}'\tilde{\mathbf{\Omega}}^{-1}\tilde{\mathbf{\Psi}},
\end{aligned}$$

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

The conjugate normal-inverse-Wishart prior is widely used in Bayesian VARs due to its computational convenience and desirable properties (see Uhlig, 1994; Faust, 1998; Uhlig, 2005; Sims and Zha, 1998; Rubio-Ramírez, Waggoner, and Zha, 2010; Kilian and Murphy, 2012, 2014). When combined with the conventional accept-reject approach, it produces independent draws from the posterior, which makes it especially attractive. However, it also imposes a Kronecker structure on the prior distribution of $\mathbf{B}$, thereby constraining its covariance matrix, and it rules out cross variable shrinkage. Consequently, researchers oftentimes consider: (i) the independent normal-inverse-Wishart prior, which avoids the Kronecker covariance structure and allows for greater flexibility; and (ii) the asymmetric priors proposed by Chan (2022), which is becoming popular because it accommodates cross-variable shrinkage. Even though, we will present the methodology using the conjugate normal-inverse-Wishart prior, due to its prevalence in the literature, in Appendix II.2 we also adapt the algorithm to these alternative priors.

Given the results in Arias, Rubio-Ramírez, and Waggoner (2025), we will combine the conjugate prior with the following uniform density over the set of orthogonal matrices:

$$\pi(\mathbf{Q} \mid \mathbf{B}, \mathbf{\Sigma}) = \begin{cases} \kappa & \text{if } \mathbf{Q} \in \mathcal{O}(n), \\ 0 & \text{otherwise,} \end{cases}$$

where $\int_{\mathcal{O}(n)} \kappa d\mathbf{Q} = 1$. This choice can be motivated by the fact that it assigns equal prior weight to both observationally equivalent vectors of impulse responses and observationally equivalent structural parameters (see Arias, Rubio-Ramírez, and Waggoner, 2025). We call this combination the conjugate uniform-normal-inverse-Wishart distribution over the orthogonal reduced-form parameterization; denote it by $UNIW(\nu, \mathbf{\Phi}, \mathbf{\Psi}, \mathbf{\Omega})$, and denote its density over the orthogonal reduced-form parameterization by $UNIW_{(\nu,\mathbf{\Phi},\mathbf{\Psi},\mathbf{\Omega})}(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})$. It is the case that

$$UNIW_{(\nu,\mathbf{\Phi},\mathbf{\Psi},\mathbf{\Omega})}(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) = \begin{cases} \kappa\, NIW_{(\nu,\mathbf{\Phi},\mathbf{\Psi},\mathbf{\Omega})}(\mathbf{B}, \mathbf{\Sigma}) & \text{if } \mathbf{Q} \in \mathcal{O}(n), \\ 0 & \text{otherwise.} \end{cases} \tag{6}$$

Inference Based on Sign Restrictions

Our objective will be to draw from the posterior of the orthogonal reduced-form parameters conditional on the sign restrictions,

$$p(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q} \mid (\mathbf{y}_t)_{t=1}^T, \boldsymbol{S}_R(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) > 0) = \frac{[\boldsymbol{S}_R(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) > 0]\, UNIW(\tilde{\nu}, \tilde{\mathbf{\Phi}}, \tilde{\mathbf{\Psi}}, \tilde{\mathbf{\Omega}})}{Pr\left(\boldsymbol{S}_R(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) > 0 \mid (\mathbf{y}_t)_{t=1}^T\right)}, \tag{7}$$

and then use f and ϕ to transform the draws to the desired vector of objects of interests such as the structural parameters or impulse responses. The traditional approach to obtain draws from Equation (7) uses the following accept-reject algorithm:

Algorithm 1. *This algorithm independently draws from* $p(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q} \mid (\mathbf{y}_t)_{t=1}^T, \boldsymbol{S}_R(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) > 0)$ *as described in Equation* (7).

1. *Draw* $(\mathbf{B}, \mathbf{\Sigma})$ *independently from the* $NIW(\tilde{\nu}, \tilde{\mathbf{\Phi}}, \tilde{\mathbf{\Psi}}, \tilde{\mathbf{\Omega}})$ *distribution.*

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

3. *Keep* $(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})$ *if* $[\boldsymbol{S}_R(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) > 0] = 1$.

4. *Return to Step 1 until the required number of draws has been obtained.*

As mentioned above, while this algorithm has been widely adopted, it is well known that there are cases in which the identified set is narrow, limiting the efficiency of the algorithm (see e.g., Kilian and Murphy, 2014; Baumeister and Hamilton, 2024; Chan, Matthes, and Yu, 2025; Read and Zhu, 2025). In the next section, we use a simple example to show its shortcomings. We will also show how a carefully designed elliptical slice sampling algorithm is not subject to this limitation and delivers dramatic speed gains. Importantly, Chan, Matthes, and Yu (2025) show a new numerically efficient version of Algorithm 1 that facilitates the drawing for a large number of structural restrictions. Therefore, when comparing our algorithm to the traditional accept-reject approach, we will use this efficient version as the benchmark.

5 The Problem with Accept-Reject Sampling

For the purposes of demonstrating the limitations of the accept-reject approach, it suffices to work with a simple example similar to the one explored by Granziera, Moon, and Schorfheide (2018). Thus, consider the following SVAR, with $n = 2$ and $m = 0$, written under the orthogonal reduced-form parameterization:

$$\mathbf{y}_t' = (y_{t,1}, y_{t,2}) = \boldsymbol{\varepsilon}_t'(\boldsymbol{\Sigma}_{tr}\,\mathbf{Q})',$$

where we let $\boldsymbol{\Sigma}_{tr} = h(\boldsymbol{\Sigma})'$. Initially, we assume $\boldsymbol{\Sigma}_{tr}$ is known, but we will later relax this assumption. Let $\sigma_{tr,ij}$ denote the i-th row and j-th column entry of $\boldsymbol{\Sigma}_{tr}$. For simplicity, we set $\sigma_{tr,11} = \sigma_{tr,22} = 1$ and $\sigma_{tr,21} = -0.9$. Note that the contemporaneous impact matrix $\mathbf{L}_0$ is defined as $\mathbf{L}_0 = \boldsymbol{\Sigma}_{tr}\,\mathbf{Q}$. Henceforth, we focus on the impulse responses to the first shock—it is straightforward to extend our analysis to the second shock.

Given the above, it is easy to see that the impact of the first shock on $y_{t,1}$ and $y_{t,2}$ can be written as $\ell_{11} = q_{11}$ and $\ell_{21} = -0.9q_{11} + q_{21}$, where ℓ_{i1} and q_{i1} are the $(i,1)$ entries of $\mathbf{L}_0$ and $\mathbf{Q}$, respectively. We now impose sign restrictions requiring that ℓ_{11} and ℓ_{21} are nonnegative. These sign restrictions imply $q_{11} \geq 0$ and $q_{21} \geq 0.9q_{11}$. Figure 1a illustrates this

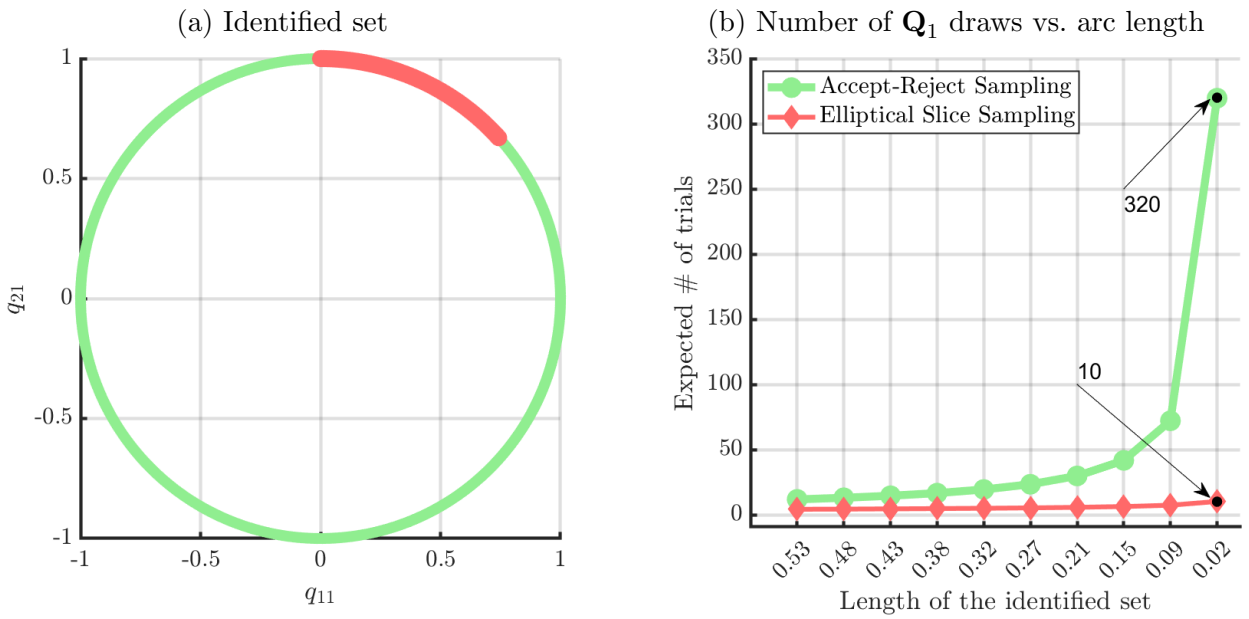

Figure 1: (a) Identified set (red) and domain of $(q_{11}, q_{21})'$ (green). (b) Expected number of draws required to meet the sign restrictions as a function of the identified set size (arc length).

setup graphically. The green circle represents the domain of $\mathbf{Q}_1 = (q_{11}, q_{21})'$, while the red arc highlights the identified set that satisfies the imposed sign restrictions.

When using the popular accept-reject sampling approach described in Algorithm 1, obtaining a draw from the posterior distribution of impulse responses satisfying the sign restrictions involves drawing a 2×1 vector $\mathbf{x}_1$ from a $N(\mathbf{0}, \mathbf{I}_2)$ distribution and converting it into a unit vector $\mathbf{q}_1$ via the normalization $\mathbf{q}_1 = \mathbf{x}_1 / \|\mathbf{x}_1\|$. The draw is accepted only if $\mathbf{q}_1$ satisfies the sign restrictions. Unrestricted draws $(q_{11}, q_{21})'$ lie uniformly on the entire unit circle (depicted in green), whereas the accepted draws are uniformly distributed only over the subset of the unit circle that meets the sign restrictions (the red arc).

The efficiency of the posterior simulator based on this type of accept-reject algorithm depends heavily on the size of the identified set. As the identified set becomes tighter, we naturally expect to discard a larger number of draws. Indeed, the expected number of draws required to satisfy the sign restrictions is inversely proportional to the probability of meeting those restrictions. Figure 1b illustrates this relationship analytically: the green line plots

the expected number of draws needed to satisfy the sign restrictions as a function of the size of the identified set (i.e., the length of the red arc). More specifically, we generate smaller identified sets by gradually moving the left endpoint of the red arc toward its right endpoint. As shown in the figure, the expected number of draws required increases hyperbolically as the identified set shrinks. In realistic scenarios, as illustrated later in our empirical applications, the number of draws required can become quite large, rendering the algorithm inefficient.

In this paper, we propose a Gibbs sampling algorithm based on the elliptical slice sampling, which draws from the identified set more efficiently. This method can be viewed as an adaptive Metropolis-Hastings algorithm that transitions from the previous draw $\mathbf{x}_1^{(0)}$ to a new draw using the following elliptical proposal:

$$\mathbf{x}_1^{(\star)} = \boldsymbol{\nu} \sin(\theta) + \mathbf{x}_1^{(0)} \cos(\theta), \text{ where } \theta \in [0, 2\pi],$$

where $\boldsymbol{\nu}$ is a 2×1 vector drawn from $N(\mathbf{0}, \mathbf{I}_2)$. The scalar parameter θ controls the step size of the proposed move. For instance, when θ is close to 0, the proposal is closer to the previous draw $\mathbf{x}_1^{(0)}$, whereas when θ approaches $\pi/2$, the proposal is closer to the newly drawn random vector $\boldsymbol{\nu}$. Unlike a conventional Metropolis-Hastings algorithm, the elliptical slice sampling adaptively searches for a suitable step size to guarantee acceptance of the proposed draw at every iteration. Intuitively, given that the previous draw lies within the identified set, the elliptical slice sampling ensures that the new proposal $\mathbf{x}_1^{(\star)}$ also remains within the identified set by uniformly drawing θ from a candidate set that shrinks exponentially. Under appropriate regularity conditions, this procedure ensures the validity and convergence of the Gibbs sampling algorithm as long as the random variable of interest, in this case $\mathbf{q}_1$, can be written as a transformation of a normally distributed random variable (see Murray, Adams, and MacKay, 2010; Natarovskii, Rudolf, and Sprungk, 2021, for details).

The fact that the candidate set for θ shrinks exponentially is an appealing feature, as it significantly reduces the number of candidate draws of $\mathbf{q}_1$ needed to satisfy the restrictions.

This efficiency gain becomes particularly important as the dimension of the model increases, since generating new draws of $\mathbf{q}_1$ is computationally costly. Figure 1b (red line) displays the average number of trials required by the elliptical slice sampling to generate an accepted draw of $\mathbf{q}_1$ within the identified set as a function of the length of the identified set. The number of required trials for the conventional accept-reject sampler grows hyperbolically, whereas that for elliptical slice sampling increases at a much slower rate.[2]

In the following section, we extend this simple example into a more realistic and useful setting by: (1) identifying multiple shocks simultaneously rather than just a single shock; (2) allowing sign restrictions to take a general form; and (3) developing a Gibbs sampling algorithm that uses the elliptical slice sampling and targets the posterior of the orthogonal reduced-form parameters conditional on the sign restrictions.

6 An Algorithm

In this section, we propose a Gibbs sampler algorithm that employs the elliptical slice sampling to draw from the posterior of the orthogonal reduced-form parameters conditional on the sign restrictions defined in Section 4. The algorithm relies on Assumption 1.

Assumption 1. *The following conditions hold:*

1.1 *For almost all* $(\mathbf{B}, \mathbf{\Sigma})$*, the set* $\{\mathbf{Q} \in \mathcal{O}(n) : \boldsymbol{S}_R(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) > 0\}$ *has positive measure.*

1.2 *For almost all* $(\mathbf{B}, \mathbf{Q})$*, the set* $\{\mathbf{\Sigma} \in \mathcal{S}(n) : \boldsymbol{S}_R(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) > 0\}$ *has positive measure.*

1.3 *For almost all* $(\mathbf{\Sigma}, \mathbf{Q})$*, the set* $\{\mathbf{B} \in \mathbb{R}^{m \times n} : \boldsymbol{S}_R(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) > 0\}$ *has positive measure.*

As noted above, the algorithm is formulated under a conjugate prior over the reduced-form parameters. Given the choice of a conjugate uniform-normal-inverse-Wishart prior

[2]To make a fair comparison, one should account for the serial correlation introduced by the Gibbs sampling algorithm, since the accept-reject algorithm generates independent draws. As demonstrated in Section 7, we compute the effective sample size and find that in this example the number of draws required to obtain one effective draw ranges from 1.04 to 1.35. Therefore, this adjustment does not alter the main conclusion illustrated in the figure.

distribution over the orthogonal reduced-form parameters, the posterior can be expressed as in Equation (7). Using this equation, we derive the conditional posterior distributions satisfying the sign restrictions for each component of the orthogonal reduced-form parameterization. Crucially, as will be shown below, we sample from each of these conditional distributions using the elliptical slice sampling. This is feasible because each conditional posterior can be represented by a distribution featuring a Gaussian kernel, thereby permitting the use of the elliptical slice sampling.

Conditional Posterior for Q. We first derive the posterior for $\mathbf{Q}$ conditional on the reduced-form parameters and the sign restrictions. Equation (7) implies:

$$p(\mathbf{Q} \mid \mathbf{B}, \mathbf{\Sigma}, (\mathbf{y}_t)_{t=1}^T, \boldsymbol{S}_R(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) > 0) = \frac{[\boldsymbol{S}_R(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) > 0]\, NIW_{(\tilde{\nu}, \tilde{\mathbf{\Phi}}, \tilde{\mathbf{\Psi}}, \tilde{\mathbf{\Omega}})}(\mathbf{B}, \mathbf{\Sigma})}{\int_{\mathcal{O}(n)} [\boldsymbol{S}_R(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) > 0]\, NIW_{(\tilde{\nu}, \tilde{\mathbf{\Phi}}, \tilde{\mathbf{\Psi}}, \tilde{\mathbf{\Omega}})}(\mathbf{B}, \mathbf{\Sigma})\, d\mathbf{Q}}$$
$$= \frac{[\boldsymbol{S}_R(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) > 0]}{\int_{\mathcal{O}(n)} [\boldsymbol{S}_R(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) > 0]\, d\mathbf{Q}} \propto [\boldsymbol{S}_R(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) > 0].$$

The first equality follows from Bayes' rule. The second equality holds because $NIW_{(\tilde{\nu}, \tilde{\mathbf{\Phi}}, \tilde{\mathbf{\Psi}}, \tilde{\mathbf{\Omega}})}(\mathbf{B}, \mathbf{\Sigma})$ is independent of $\mathbf{Q}$. The proportionality follows from Assumption 1.1.

To sample from $p(\mathbf{Q} \mid \mathbf{B}, \mathbf{\Sigma}, (\mathbf{y}_t)_{t=1}^T, \boldsymbol{S}_R(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) > 0)$, we exploit the mapping from $\mathbf{X} \sim N(\mathbf{0}_{\mathbf{n}\times\mathbf{n}}, \mathbf{I}_\mathbf{n}, \mathbf{I}_\mathbf{n})$ to $\mathbf{Q}$ via the QR decomposition, denoted $\mathbf{Q} = \gamma(\mathbf{X})$. This ensures that $\mathbf{Q}$ is uniformly distributed with respect to the Haar measure. Hence, if we draw $\mathbf{X}$ from $[\boldsymbol{S}_R(\mathbf{B}, \mathbf{\Sigma}, \gamma(\mathbf{X})) > 0]p(\mathbf{X})$ and transform it using $\mathbf{Q} = \gamma(\mathbf{X})$, we obtain draws from the desired distribution. Since $\mathbf{X}$ is Gaussian, sampling can be performed using the elliptical slice sampling.

Conditional Posterior for Σ. We next derive the posterior of $\mathbf{\Sigma}$ conditional on $(\mathbf{B}, \mathbf{Q})$ and the sign restrictions:

$$p(\mathbf{\Sigma} \mid \mathbf{B}, \mathbf{Q}, (\mathbf{y}_t)_{t=1}^T, \boldsymbol{S}_R(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) > 0) = \frac{[\boldsymbol{S}_R(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) > 0]\, N_{(\tilde{\mathbf{\Psi}}, \tilde{\mathbf{\Omega}}, \mathbf{\Sigma})}(\mathbf{B}) IW_{(\tilde{\nu}, \tilde{\mathbf{\Phi}})}(\mathbf{\Sigma})}{\int_{\mathcal{S}(n)} [\boldsymbol{S}_R(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) > 0]\, N_{(\tilde{\mathbf{\Psi}}, \tilde{\mathbf{\Omega}}, \mathbf{\Sigma})}(\mathbf{B}) IW_{(\tilde{\nu}, \tilde{\mathbf{\Phi}})}(\mathbf{\Sigma})\, d\mathbf{\Sigma}}$$

$$\propto [\boldsymbol{S}_R(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}) > 0]\, N_{(\tilde{\boldsymbol{\Psi}}, \tilde{\boldsymbol{\Omega}}, \boldsymbol{\Sigma})}(\mathbf{B}) IW_{(\tilde{\nu}, \tilde{\boldsymbol{\Phi}})}(\boldsymbol{\Sigma}).$$

The factorization comes from the decomposition $NIW = N \times IW$, which is convenient because it highlights the Gaussian kernel. The proportionality follows from Assumption 1.2.

To sample from $p(\boldsymbol{\Sigma} \mid \mathbf{B}, \mathbf{Q}, (\mathbf{y}_t)_{t=1}^T, \boldsymbol{S}_R(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}) > 0)$, we exploit the mapping $\boldsymbol{\Sigma} = \varsigma(\mathbf{R}) = (\mathbf{R}\,\mathbf{R}')^{-1}$, where $\mathbf{R} \sim N(\mathbf{0}_{\mathbf{n} \times \tilde{\nu}}, \tilde{\boldsymbol{\Phi}}^{-1}, \mathbf{I}_{\tilde{\nu}})$. This ensures that $\boldsymbol{\Sigma}$ is inverse-Wishart. Hence, if we draw $\mathbf{R}$ from $[\boldsymbol{S}_R(\varsigma(\mathbf{R}), \boldsymbol{\Sigma}, \mathbf{Q}) > 0]\, N_{(\tilde{\boldsymbol{\Psi}}, \tilde{\boldsymbol{\Omega}}, \varsigma(\mathbf{R}))}(\mathbf{B}) p(\mathbf{R})$ and transform it using $\boldsymbol{\Sigma} = \varsigma(\mathbf{R})$, we obtain draws from the desired distribution. Since $\mathbf{R}$ is Gaussian, sampling can be performed using the elliptical slice sampling.

Conditional Posterior for B. Finally, the posterior of $\mathbf{B}$ conditional on $(\boldsymbol{\Sigma}, \mathbf{Q})$ and the sign restrictions is:

$$p(\mathbf{B} \mid \boldsymbol{\Sigma}, \mathbf{Q}, (\mathbf{y}_t)_{t=1}^T, \boldsymbol{S}_R(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}) > 0) = \frac{[\boldsymbol{S}_R(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}) > 0]\, N_{(\tilde{\boldsymbol{\Psi}}, \tilde{\boldsymbol{\Omega}}, \boldsymbol{\Sigma})}(\mathbf{B}) IW_{(\tilde{\nu}, \tilde{\boldsymbol{\Phi}})}(\boldsymbol{\Sigma})}{\int_{\mathbb{R}^{m \times n}} [\boldsymbol{S}_R(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}) > 0]\, N_{(\tilde{\boldsymbol{\Psi}}, \tilde{\boldsymbol{\Omega}}, \boldsymbol{\Sigma})}(\mathbf{B}) IW_{(\tilde{\nu}, \tilde{\boldsymbol{\Phi}})}(\boldsymbol{\Sigma})\, d\mathbf{B}}$$
$$\propto [\boldsymbol{S}_R(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}) > 0]\, N_{(\tilde{\boldsymbol{\Psi}}, \tilde{\boldsymbol{\Omega}}, \boldsymbol{\Sigma})}(\mathbf{B}).$$

The proportionality follows from Assumption 1.3. Notably, the conditional posterior contains a Gaussian kernel, making the elliptical slice sampling applicable.

Gibbs Sampler. Having defined the three conditional posteriors, we can now outline the Gibbs sampler:

Algorithm 2. *This algorithm draws from* $p(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q} \mid (\mathbf{y}_t)_{t=1}^T, \boldsymbol{S}_R(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q}) > 0)$ *as described in Equation* (7).

1. *Set* $I > 1$, *initialize* $i = 1$, *and assign initial values to* $(\mathbf{B}^{i-1}, \boldsymbol{\Sigma}^{i-1})$.

2. *Draw* $\mathbf{Q}^i$ *from*

$$p(\mathbf{Q} \mid \mathbf{B}^{i-1}, \boldsymbol{\Sigma}^{i-1}, (\mathbf{y}_t)_{t=1}^T, \boldsymbol{S}_R(\mathbf{B}^{i-1}, \boldsymbol{\Sigma}^{i-1}, \mathbf{Q}) > 0) \propto [\boldsymbol{S}_R(\mathbf{B}^{i-1}, \boldsymbol{\Sigma}^{i-1}, \mathbf{Q}) > 0].$$

3. *Draw* $\boldsymbol{\Sigma}^i$ *from*

$$p(\boldsymbol{\Sigma} \mid \mathbf{B}^{i-1}, \mathbf{Q}^i, (\mathbf{y}_t)_{t=1}^T, \boldsymbol{S}_R(\mathbf{B}^{i-1}, \boldsymbol{\Sigma}, \mathbf{Q}^i) > 0) \propto [\boldsymbol{S}_R(\mathbf{B}^{i-1}, \boldsymbol{\Sigma}, \mathbf{Q}^i) > 0] N_{(\tilde{\boldsymbol{\Psi}}, \tilde{\boldsymbol{\Omega}}, \boldsymbol{\Sigma})}(\mathbf{B}^{i-1}) IW_{(\tilde{\nu}, \tilde{\boldsymbol{\Phi}})}(\boldsymbol{\Sigma}).$$

4. *Draw* $\mathbf{B}^i$ *from*

$$p(\mathbf{B} \mid \boldsymbol{\Sigma}^i, \mathbf{Q}^i, (\mathbf{y}_t)_{t=1}^T, \boldsymbol{S}_R(\mathbf{B}, \boldsymbol{\Sigma}^i, \mathbf{Q}^i) > 0) \propto [\boldsymbol{S}_R(\mathbf{B}, \boldsymbol{\Sigma}^i, \mathbf{Q}^i) > 0] N_{(\tilde{\boldsymbol{\Psi}}, \tilde{\boldsymbol{\Omega}}, \boldsymbol{\Sigma}^i)}(\mathbf{B}).$$

5. *If* $i < I$*, increment* i *and return to Step 2.*

As mentioned above, we have used the conjugate normal-inverse-Wishart prior over the reduced-form parameters to describe the algorithm. Appendix I discusses how to set initial values for the algorithm, and Appendices II.1 and II.2 show that the approach can be easily adapted to two alternative priors: the independent normal-inverse-Wishart prior and the asymmetric conjugate priors of Chan (2022), respectively.

7 Applications

We now illustrate the performance of our algorithm using two empirical applications. The first is a small-scale SVAR of the global oil market, based on the model in Kilian and Murphy (2014), which identifies flow supply, flow demand, and speculative demand shocks using a combination of sign and elasticity bounds. The tight identifying assumptions in this model render traditional accept-reject methods computationally intensive, whereas our algorithm improves efficiency while replicating the main results. The second application revisits the large-scale SVAR model of the U.S. economy developed by Crump, Eusepi, Giannone, Qian, and Sbordone (2025) and analyzed structurally by Chan, Matthes, and Yu (2025), which includes 35 macroeconomic and financial variables and identifies up to eight structural shocks. We show that our algorithm remains computationally stable as the number of restrictions

increases, in contrast to the exponential rise in computation time exhibited by the accept-reject method. Both applications highlight the scalability of our approach in distinct empirical settings. For each application, we first demonstrate that our approach replicates the main results reported in the original papers and then analyze the computational timing to show that our method can be more efficient than the traditional accept-reject algorithm.

7.1 Small SVAR of the World Oil Market

In our first application, we replicate the results of Kilian and Murphy (2014), who extend the framework of Kilian and Murphy (2012) by incorporating oil inventories to identify speculative demand shocks. The identification strategy in Kilian and Murphy (2014) relies on tight sign and elasticity bound restrictions, which result in a small identified set and may render standard accept-reject algorithms slow. Notably, Kilian and Murphy (2014) adopt an approach similar to that of Chan, Matthes, and Yu (2025), relying on permutations and sign alternations. Therefore, when using the same set of sign and elasticity bound restrictions, the computation times we report for the accept-reject algorithm are comparable to those in the original study and similar to those obtained with our approach. However, when an additional elasticity bound is introduced, the accept-reject approach becomes computationally unfeasible, while the Gibbs sampler remains as fast, highlighting the speed advantages of our algorithm in such settings.

Model Specification and Impulse Responses

We begin by describing the model specification in Kilian and Murphy (2014). They model the global market for crude oil using a four-variable SVAR featuring the percent change in global crude oil production, a measure of global real activity, the real price of crude oil, and the change in global above-ground crude oil inventories. The SVAR is specified at a monthly frequency, with an estimation sample covering 1973:M2–2009:M8. The model includes 24 lags, a constant, and seasonal dummies to remove seasonal variation. Kilian and Murphy

(2014) adopt a weak conjugate normal-inverse-Wishart prior distribution (see, e.g., Uhlig, 2005) for the reduced-form parameters.

Turning to identification, the goal of Kilian and Murphy (2014) is to identify three structural shocks using a combination of sign restrictions on impact impulse responses, sign restrictions at horizons 1 through 12, and elasticity bounds. Table 1 summarizes the identifying assumptions. The structural shocks are labeled flow supply shock, flow demand shock, and speculative demand shock.

	Sign Restrictions on Impact Impulse Responses		
Variable/Shock	Flow supply	Flow demand	Speculative demand
Oil production	−1	+1	+1
Real activity	−1	+1	−1
Real price of oil	+1	+1	+1
Inventories			+1
	Elasticity Bounds		
	Flow supply shock	Flow demand shock	Speculative demand shock
Price Elasticity of Oil Supply		(0, 0.025)	(0, 0.025)
	Sign Restrictions on Impulse Responses at Horizons 1 through 12		
	Flow supply shock	Flow demand shock	Speculative demand shock
Real activity	−1		
Real price of oil	+1		

Table 1: Sign and Elasticity Bound Restrictions

Note: All shocks raise the real oil price. ±1 indicates positive or negative sign restrictions; blanks indicate no restriction.

We adopt the exact same specification and use Algorithm 2. We obtain one million draws, saving one every 10; hence, the figures are produced using one hundred thousand draws. Figure 2 presents the impulse responses to the three shocks. The results broadly match those in Kilian and Murphy (2014). In particular, a negative flow supply shock causes a persistent decline in global economic activity and oil inventories, and a persistent increase in the real price of oil. The response of oil production is persistently negative. A positive flow demand shock is associated with a persistent increase in global economic activity, a persistent increase in the real price of oil, and a positive response of oil production. Oil production increases sluggishly, given the imposed elasticity bounds, and peaks at about one year after the shock before declining to pre-shock levels. Finally, a positive speculative demand shock causes a persistent increase in the real price of oil and a large increase in inventories. Global real

activity and oil production decline persistently in response to this shock, although the effects are modest.

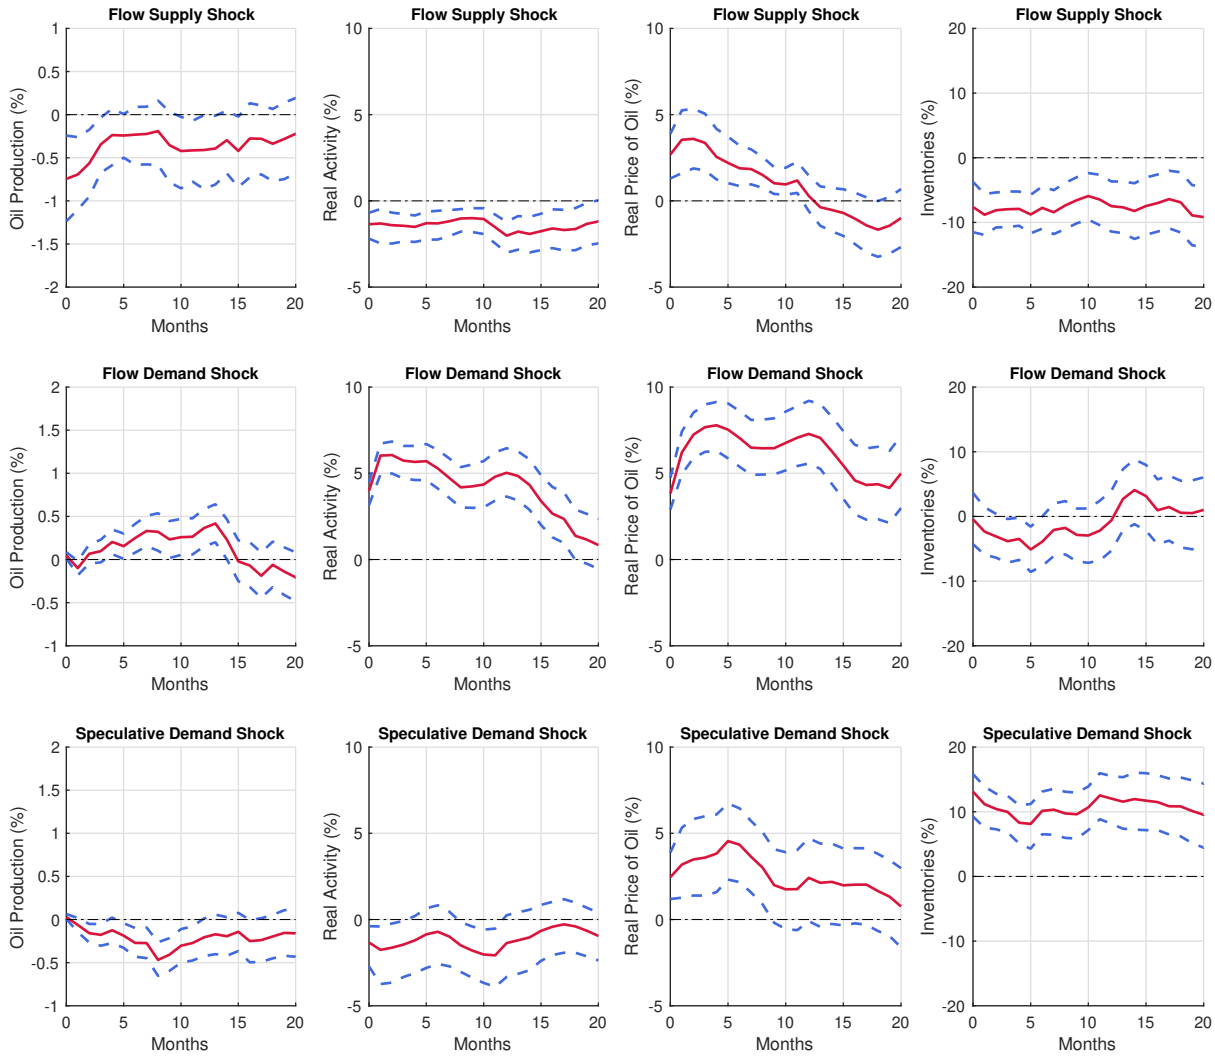

Figure 2: Impulse Responses

Note: The solid red lines depict the pointwise posterior median; the dashed blue lines depict the pointwise 68 percent posterior probability bands.

Timing

We next compare the computational time of our Gibbs sampler to that of the accept-reject algorithm. Table 2 reports the time (in hours) per 1,000 effective draws—defined as total computation time divided by the effective sample size and scaled by 1,000—using Algorithm 2

and the accept-reject method.[3] As highlighted above, we implement the efficient variant of the accept-reject algorithm proposed by Chan, Matthes and Yu. We approximate the effective sample size using the multivariate effective sample size metric of Vats, Flegal, and Jones (2019), which is well suited for SVAR analysis where inference often targets high-dimensional objects such as vectors of impulse responses. In particular, we estimate effective sample size under the impulse response parameterization using the multivariate batch means approach described in that work, with a batch size of $N^{1/4}$, where N denotes the number of stored draws.[4] Because the accept-reject algorithm produces independent draws, effective draws and sample size are the same.

The column "Benchmark Model" in Table 2 compares the time (in hours) required to obtain 1,000 effective draws under the specification of Kilian and Murphy (2014). As shown in the table, the Gibbs sampler requires less than 2 minutes to produce 1,000 effective draws. In contrast, the accept-reject algorithm takes approximately 20 minutes to achieve the same number of draws.[5] To further illustrate the gains of our proposed Gibbs sampler, we consider a scenario in which a researcher imposes an additional restriction on the price elasticity of oil demand in response to a flow supply shock. As emphasized by Caldara, Cavallo, and Iacoviello (2019), such a restriction is empirically important in SVAR models of the oil market. Following their work, we constrain the price elasticity of crude oil demand to lie within a narrow interval around the point estimate of −0.08 reported by Caldara, Cavallo, and Iacoviello (2019); specifically, we impose the restriction that the elasticity must lie in the interval $(-0.09, , -0.07)$. The column "Benchmark Model + Additional Restriction" in Table 2 reports the results under this added constraint. The time under the Gibbs sampler is under 6 minutes, while the performance of the accept-reject algorithm deteriorates sharply, requiring nearly 8 hours to obtain 1,000 draws.

[3]Since the accept-reject approach produces independent draws, the number of effective draws equals the total number of draws in that case.

[4]We consider the first three columns of the impulse response parameters, which correspond to the shocks of interest in this application. Results are robust to using impulse responses for all shocks.

[5]All computations were performed in MATLAB on an Intel Xeon Platinum 8488C processor with 16 active cores running at 2.4 GHz on an x86_64 architecture.

Specification	Benchmark Model	Benchmark Model + Additional Restriction
Gibbs Sampler	0.03	0.10
Accept-Reject	0.33	7.92

Table 2: Time (Hours) Per 1,000 Effective Draws

These results indicate that, for this model, the accept-reject algorithm is already near its computational limit under the benchmark specification, and that introducing even a single additional restriction dramatically reduces its efficiency. In contrast, the Gibbs sampler maintains its performance even as additional identifying restrictions are introduced.

7.2 Large SVAR of the U.S. Economy

In our second application, we replicate and extend the analysis of Chan, Matthes, and Yu (2025), who build on the large-scale SVAR framework of Crump, Eusepi, Giannone, Qian, and Sbordone (2025) to study the structural dynamics of the U.S. economy. Their model incorporates 35 macroeconomic and financial variables commonly monitored by the Federal Reserve and identifies eight structural shocks using an extensive set of sign and ranking restrictions. Chan, Matthes, and Yu (2025) employ an accept-reject algorithm, which becomes computationally intensive as the number of identifying restrictions increases.[6] As we show below, our algorithm is more efficient.

To demonstrate this, we extend the baseline model by identifying two additional shocks—an oil price shock and a consumer sentiment shock—bringing the total number of sign restrictions from 105 to 129. This provides a stringent test of our algorithm's performance relative to the accept-reject method. Importantly, Chan, Matthes, and Yu (2025) use the asymmetric priors defined in Chan (2022) for the reduced-form parameters, instead of the Minnesota prior used in Crump, Eusepi, Giannone, Qian, and Sbordone (2025). To simplify the comparison, we revert to the Minnesota prior when applying both the accept-reject algorithm and the Gibbs

[6]We thank Christian Matthes for sharing their replication files and data with us.

sampler. Appendix III compares both approaches under the asymmetric prior.

Model Specification and Impulse Responses

The SVAR used in this section is specified at a quarterly frequency, includes a constant and five lags, and uses an estimation sample that spans from 1973:Q2 until 2019:Q4. As mentioned, we assume a Minnesota prior for the reduced-form parameters and set the hyper-parameters following Giannone, Lenza, and Primiceri (2015). Turning to the identification, Table 3 summarizes the variables and the sign and ranking restrictions imposed on the contemporaneous impulse responses. Chan, Matthes, and Yu (2025) consider only the first eight shocks (demand, investment, financial, monetary, government spending, technology, labor supply, and wage bargaining). In total, 105 sign restrictions are imposed in their baseline specification. We have added two additional shocks (labeled oil price and consumer sentiment) to assess the performance of our algorithm. With the inclusion of these two shocks, the total number of sign restrictions increases to 129. When using the Gibbs sampler, we obtain one million draws and retain one every 10.

Let us begin by describing the selected impulse responses to a unit standard deviation expansionary demand shock, shown in Figure 3a. Red lines depict point-wise posterior medians, and shaded areas represent point-wise 68% posterior probability bands. The signs of the impact responses of real GDP, the PCE price index, the federal funds rate, and the unemployment rate are restricted. The remaining horizons, as well as the responses of non-residential investment and the real wage, are unrestricted. As can be seen, the demand shock causes a transient increase in output and prices, and a decrease in the unemployment rate. The federal funds rate increases in response to the shock. The restrictive stance of monetary policy eventually lowers economic activity, as seen, for example, in the decline of non-residential investment. The real wage decreases in the short run in response to the shock, as nominal wage increases are not sufficient to offset higher prices—possibly due to sluggish nominal wage adjustment.

Sign restrictions	Dem	Inv	Fin	Mon	Gov	Tec	Lab	Wag	Oil	Con
GDP	+1	+1	+1	-1	+1	+1	+1	+1	+1	+1
PCE	0	0	0	0	0	+1	0	0	+1	+1
Residential investment	0	0	0	0	0	0	0	0	0	+1
Nonresidential investment	0	+1	0	0	0	+1	0	0	+1	+1
Exports	0	0	0	0	0	0	0	0	0	0
Imports	0	0	0	0	0	0	0	0	0	0
Government spending	0	0	0	0	+1	0	0	0	0	0
Fed. budget surplus/deficit	0	0	0	0	-1	0	0	0	0	0
Fed. tax receipts	0	0	0	0	+1	0	0	0	0	0
GDP deflator	+1	+1	+1	-1	+1	-1	-1	-1	-1	+1
PCE index	+1	+1	+1	-1	+1	-1	-1	-1	-1	+1
PCE index less F&E	+1	+1	+1	-1	+1	-1	-1	-1	-1	+1
CPI index	+1	+1	+1	-1	+1	-1	-1	-1	-1	+1
CPI index less F&E	+1	+1	+1	-1	+1	-1	-1	-1	-1	+1
Hourly wage	0	0	0	0	0	+1	-1	-1	+1	0
Labor productivity	0	0	0	0	0	+1	0	0	+1	0
Utilization-adjusted TFP	0	0	0	0	0	+1	0	0	+1	0
Employment	0	0	0	-1	0	0	-1	0	0	0
Unemployment rate	-1	-1	-1	+1	-1	-1	+1	-1	+1	+1
Industrial production index	+1	+1	+1	-1	0	0	0	0	0	0
Capacity utilization	+1	+1	+1	-1	0	0	0	0	0	0
Housing starts	0	0	0	0	0	0	0	0	0	0
Disposable income	0	0	0	0	0	0	0	0	0	0
Consumer sentiment	0	0	0	0	0	0	0	0	0	0
Fed funds rate	+1	+1	+1	+1	+1	0	0	0	0	0
3-month T-bill rate	+1	+1	+1	+1	+1	0	0	0	0	0
2-year T-note rate	0	0	0	+1	0	0	0	0	0	0
5-year T-note rate	0	0	0	+1	0	0	0	0	0	0
10-year T-note rate	0	0	0	+1	0	0	0	0	0	0
Prime rate	+1	+1	+1	+1	+1	0	0	0	0	0
Aaa corporate bond yield	0	0	0	+1	0	0	0	0	0	0
Baa corporate bond yield	0	0	0	+1	0	0	0	0	0	0
Trade-weighted US index	0	0	0	0	0	0	0	0	0	0
S&P 500	0	-1	+1	-1	0	0	0	0	0	+1
Spot oil price	0	0	0	0	0	0	0	0	-1	0
Ranking restrictions										
Nonresidential investment/GDP	-1	+1	+1	0	0	0	0	0	0	0
Government spending/GDP	-1	-1	-1	0	+1	0	0	0	0	0
$\mathbb{N}^0$ of restrictions	14	15	15	19	14	12	8	8	13	11
Cum. $\mathbb{N}^0$ of restrictions	14	29	44	63	77	89	97	105	118	129

Table 3: Sign restrictions, ranking restrictions and identified shocks

Note: The mnemonics for the shocks are as follows. Dem: demand, Inv: Investment, Fin: Financial, Mon: Monetary Policy, Gov: Government Spending, Tec: Technology, Lab: Labor Supply, Wag: Wage Bargaining, Oil: Oil Price, Con: Consumer Sentiment.

The investment shock, shown in Figure 3b, appears similar to the demand shock in terms of economic consequences for real GDP, the federal funds rate, the price level, and the unemployment rate. However, the impulse response of non-residential investment is substantially different. In particular, the investment shock causes a short-run boom in

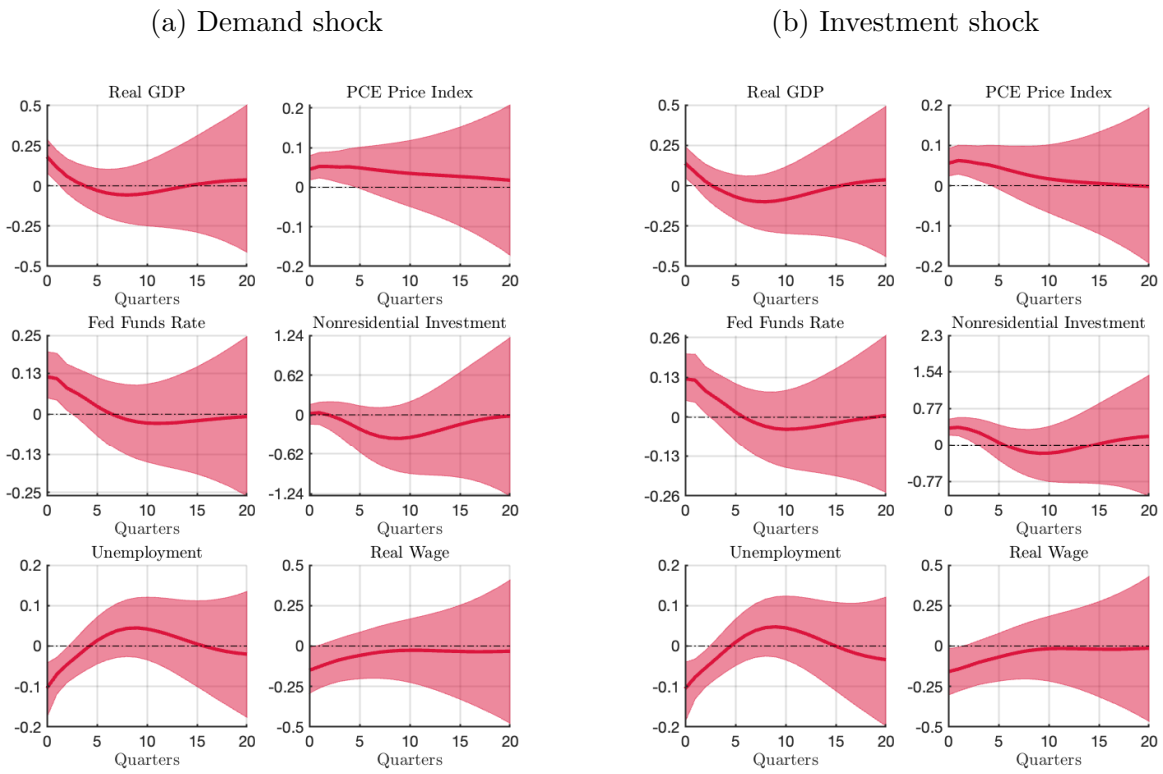

Figure 3: Impulse Responses

non-residential investment. This finding partly reflects the ranking restriction requiring that the impact response of non-residential investment be larger than the impact response of real GDP. As with the demand shock, the investment shock causes a persistently negative response of the real wage.

Turning to the financial shock, shown in Figure 4a, it is worth highlighting that this shock is identified using the same sign restrictions as the investment shock, except for the impact response of the S&P 500, which is assumed to be positive instead of negative. Overall, the impulse responses are similar, except that the decline in non-residential investment after five quarters is slightly less pronounced under the financial shock, consistent with the positive response of asset prices.

The impulse responses to a unit standard deviation contractionary monetary policy shock are depicted in Figure 4b. This shock causes the federal funds rate to remain above zero for more than two years, reflecting inertia in the conduct of monetary policy. Real GDP

(a) Financial shock

(b) Monetary policy shock

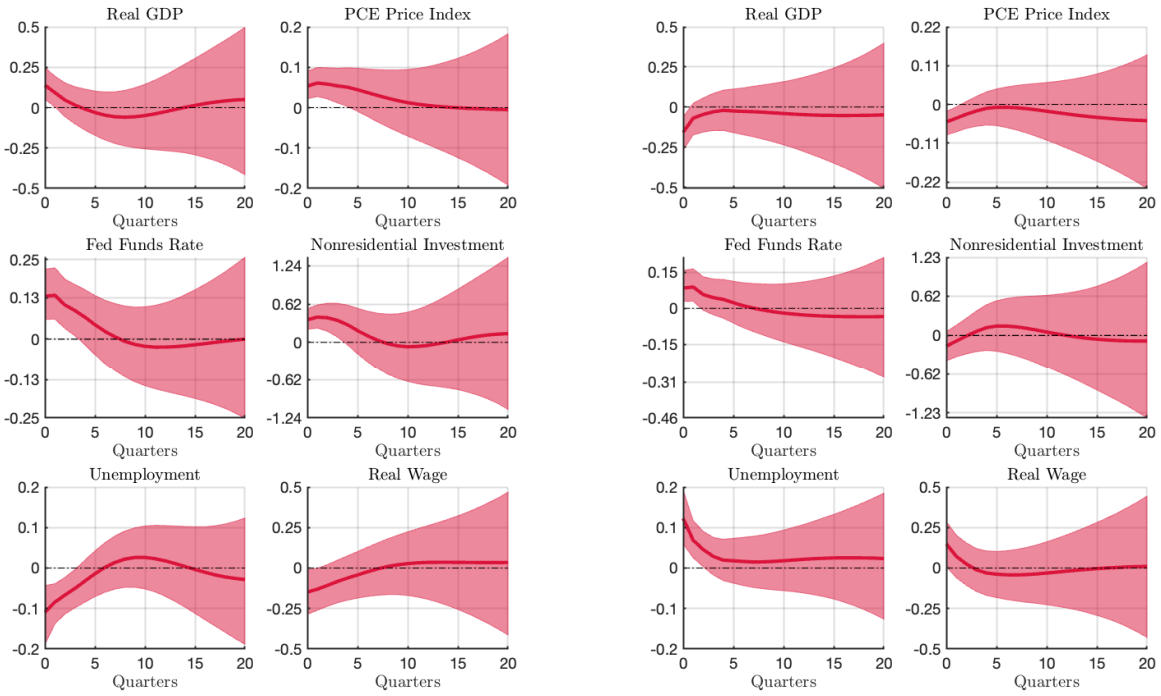

Figure 4: Impulse Responses

and prices decline persistently, and the unemployment rate jumps upon impact before slowly returning to baseline. Non-residential investment drops on impact and recovers after about one year, in line with a less restrictive monetary policy stance. The real wage increases, driven by a decrease in the price level. A notable aspect of these responses is that they suggest monetary policy can operate with shorter lags than traditionally assumed under the "long and variable lags" view.

The government spending shock is shown in Figure 5a. An expansionary one unit standard deviation government spending shock leads to an increase in real GDP for about two quarters and to a long-lasting increase in the price level. To conclude, we discuss the impulse responses to the supply-related structural shocks, that is, the technology, labor supply, and wage bargaining shocks. A unit standard deviation positive technology shock leads to a protracted increase in real GDP and non-residential investment (see Figure 5b). The higher level of output is accompanied by a sustained decline in the unemployment rate and a sustained

(a) Government spending shock

(b) Technology shock

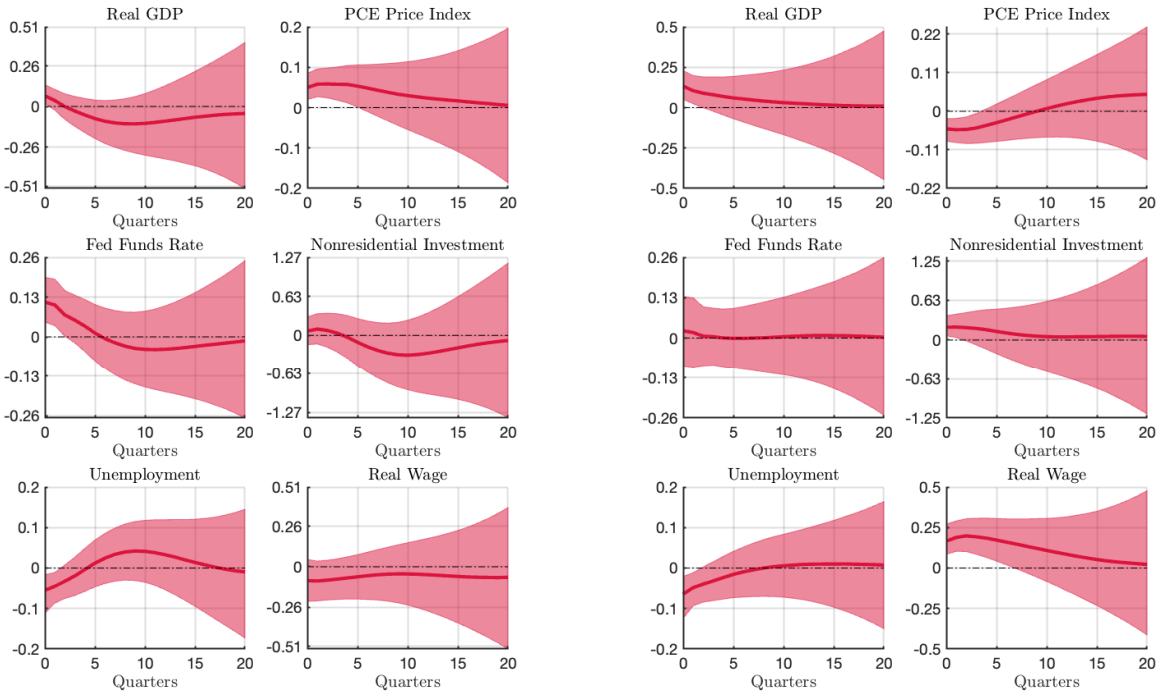

Figure 5: Impulse Responses

increase in the real wage. The federal funds rate rises marginally, indicating that monetary policy remains roughly neutral in response to technology shocks.

The responses to a unit standard deviation positive labor supply shock are shown in Figure 6a. This shock induces a hump-shaped response of real GDP and leads to persistently lower prices. The responses to a unit standard deviation negative wage bargaining shock are shown in Figure 6b. The identifying assumptions for this shock are identical to those of an expansionary labor supply shock, except that the unemployment rate is assumed to decrease upon impact. When a negative wage bargaining shock occurs, workers experience a decline in their nominal wage alongside a decrease in the unemployment rate. The real wage remains unaffected on impact, as the lower wages are offset by the assumed decrease in the price level. Subsequently, the price level remains below zero, inducing an increase in the real wage.

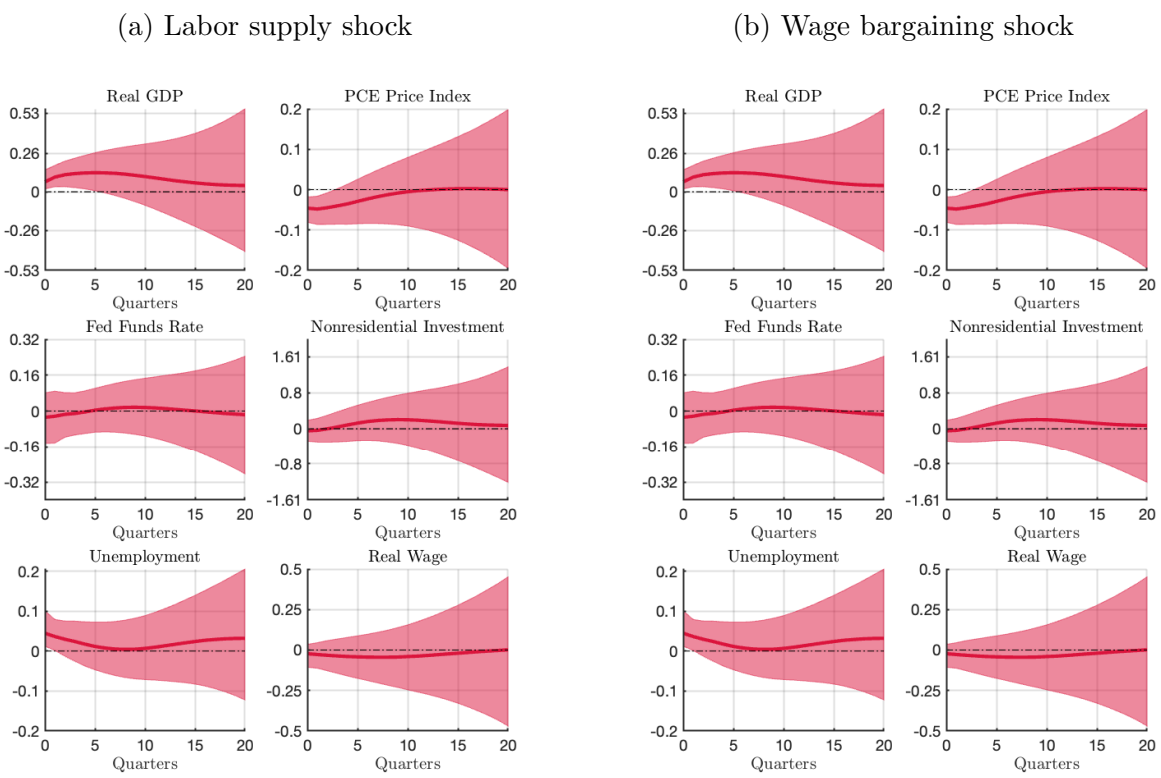

Figure 6: Impulse Responses

Timing

We begin by comparing the efficiency of the Gibbs sampler algorithm relative to the accept-reject algorithm when replicating the identification scheme in Chan, Matthes, and Yu (2025). Figure 7a reports the time (in minutes) per 1,000 effective draws using Algorithm 2 as a function of the number of identified shocks.[7]

To assess the computational time as a function of the size of the identified set, we proceed incrementally: we first obtain draws by identifying only the demand shock, then add the investment shock, the financial shock, and so on, until all eight shocks in Table 3 are included. As shown, the time per 1,000 effective draws remains computationally feasible even as the number of sign restrictions increases.

Figure 7b replicates the same figure but using the efficient accept-reject version of

[7]When computing the multivariate effective sample size, we only consider the columns of the impulse response parameters corresponding to the shocks of interest in this application.

Figure 7: Time Per 1,000 Effective Draws

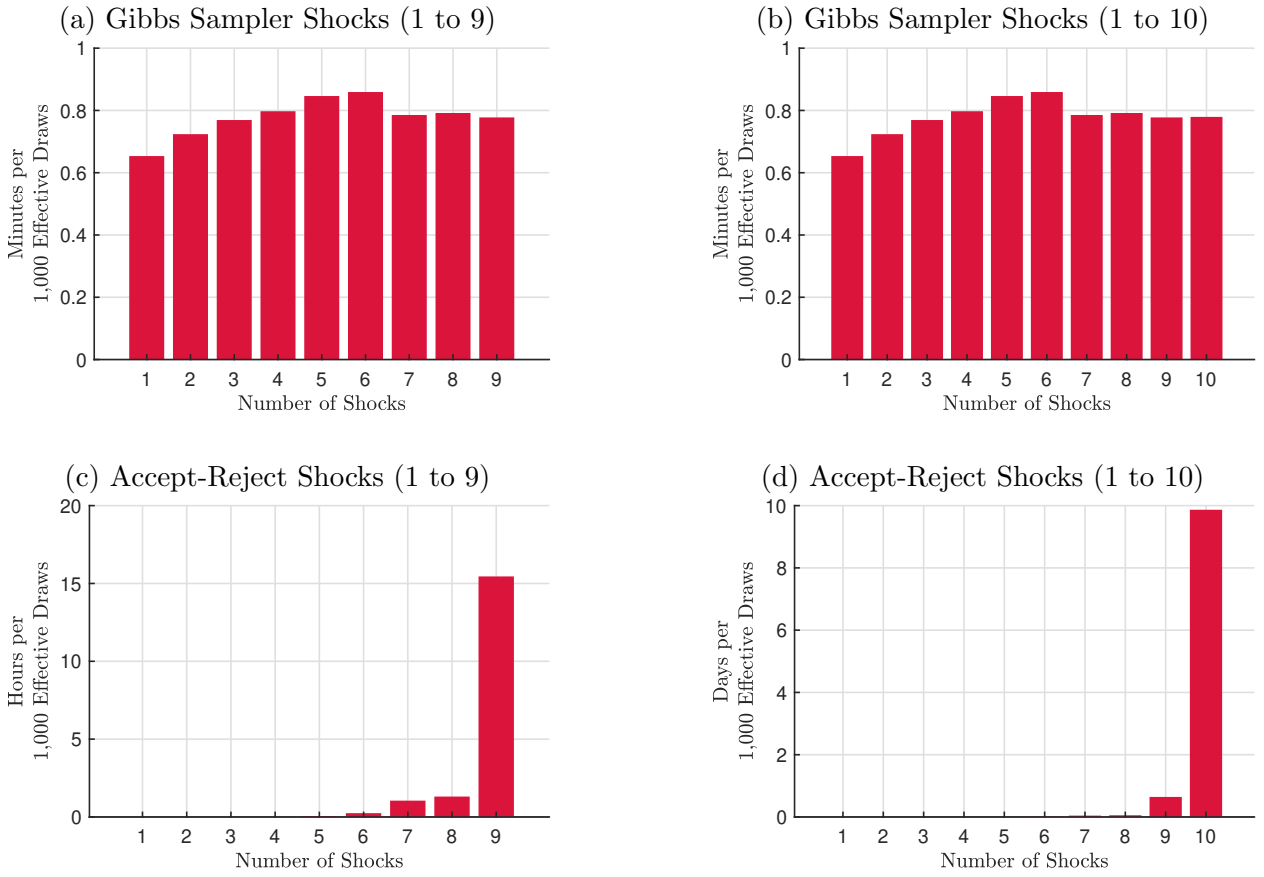

Figure 8: Gibbs Sampler vs. Accept-Reject

Note: The time of the accept-reject algorithm for shocks 9 and 10 is extrapolated based on 10 draws.

Algorithm 1 proposed by Chan, Matthes, and Yu (2025). In this case, the computation time increases dramatically, as shown in the figure. Although the runtime will vary depending on the hardware architecture and the number of variables, the main conclusion from comparing Figures 7a–7b remains unchanged: the performance of the accept-reject algorithm can

deteriorate sharply as the identified set narrows. Figure 7c combines the timings to facilitate visual comparison.

To further emphasize this point, we now consider additional shocks to illustrate that the accept-reject approach can eventually become impractical. Specifically, we extend the number of shocks identified in Chan, Matthes, and Yu (2025) by adding the oil price shock and the consumer sentiment shock described in Table 3. Figures 8a–8d replicate the exercise shown in Figures 7a–7b for the cases of nine and ten shocks.[8]

As the reader can see, computation time does not increase exponentially when we use the elliptical slice sampling approach. In contrast, when we consider nine shocks under the accept-reject approach, the times are now measured in hours, and when we consider ten shocks, the times are measured in days—Figures 8c and 8d provide the detailed timings. To facilitate the comparison, Figure 9 overlays both sets of timings to make clear that our algorithm can handle settings (in terms of the number of variables and shocks) that the traditional accept-reject approach cannot.

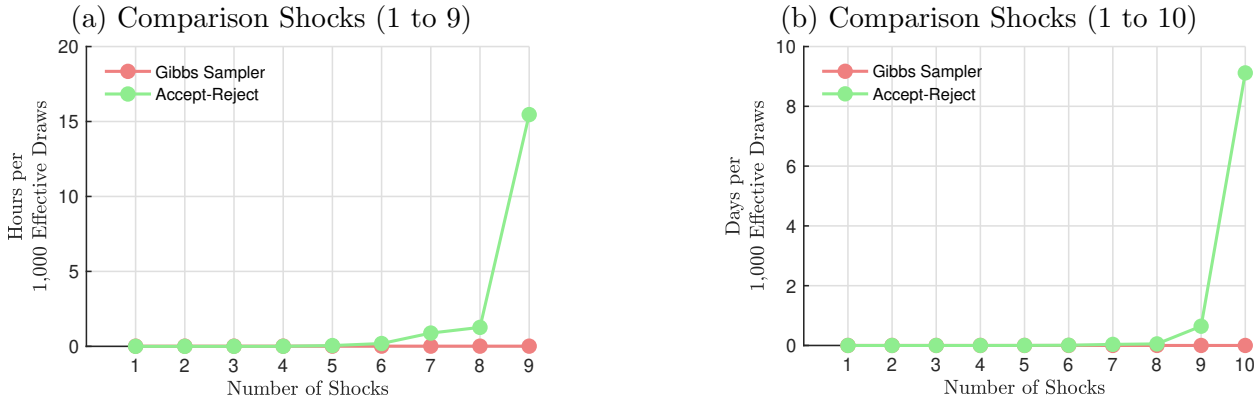

Figure 9: Gibbs Sampler vs. Accept-Reject

[8]The runtime of the accept-reject algorithm for the nine- and ten-shock cases is extrapolated based on ten draws.

8 Pitfalls of the Conditionally Uniform Prior

The main computational cost of our approach stems from running the Gibbs sampler, particularly in large models, since it produces auto-correlated draws. As a result, one might be tempted to bypass this cost and instead address the bottleneck issues inherent in the accept-reject approach by employing the conditionally uniform prior approach described in Uhlig (2017), Amir-Ahmadi and Drautzburg (2021), and Read and Zhu (2025), among others. The main appeal of this simpler approach is that, like the accept-reject algorithm, it generally yields independent draws. While such a simplification is indeed attractive due to its lower computational burden, it is essential for the researcher to be aware of a critical drawback, which we will explain in this section. Before turning to this pitfall, we first describe the conditionally uniform prior and outline the main algorithmic steps commonly used in its implementation.

Let $\mathbb{Q}_n(\mathbf{B},\mathbf{\Sigma}) = \{\mathbf{Q} \in \mathbb{Q}_n : \mathbf{S}_R(\mathbf{B},\mathbf{\Sigma},\mathbf{Q}) > 0\}$ and define $\int_{\mathbb{Q}_n(\mathbf{B},\mathbf{\Sigma})} \kappa(\mathbf{B},\mathbf{\Sigma}) d\mathbf{Q} = 1$. The prior underlying the conditionally uniform approach is then given by:

$$CUNIW_{(\nu,\mathbf{\Phi},\mathbf{\Psi},\mathbf{\Omega})}(\mathbf{B},\mathbf{\Sigma},\mathbf{Q}) = \begin{cases} \kappa(\mathbf{B},\mathbf{\Sigma})\, NIW_{(\nu,\mathbf{\Phi},\mathbf{\Psi},\mathbf{\Omega})}(\mathbf{B},\mathbf{\Sigma}) & \text{if } \mathbf{Q} \in \mathbb{Q}_n(\mathbf{B},\mathbf{\Sigma}), \\ 0 & \text{otherwise.} \end{cases} \tag{8}$$

Crucially, $\kappa(\mathbf{B},\mathbf{\Sigma})$ depends on $(\mathbf{B},\mathbf{\Sigma})$, while κ in Equation (6) does not. This is due to the fact that the conditionally uniform approach combines the conjugate normal-inverse-Wishart prior over the reduced-form parameters with $\pi(\mathbf{Q} \mid \mathbf{B},\mathbf{\Sigma})$ of the form:

$$\pi(\mathbf{Q} \mid \mathbf{B},\mathbf{\Sigma}) = \begin{cases} \kappa(\mathbf{B},\mathbf{\Sigma}) & \text{if } \mathbf{Q} \in \mathcal{O}(n)(\mathbf{B},\mathbf{\Sigma}), \\ 0 & \text{otherwise.} \end{cases}$$

This conditional uniform prior has the property that it overweights reduced-form parameters with smaller identified sets, measured by $\kappa^{-1}(\mathbf{B},\mathbf{\Sigma})$ (see Uhlig, 2017). Unlike the prior

in Equation (6), the prior in Equation (8) cannot be justified using the results of Arias, Rubio-Ramírez and Waggoner (2025), who aim to construct priors that, among other things, separate inference from identification, as it is undesirable for the prior to change when the restrictions are modified, since this makes it impossible to determine whether differences in results stem from changes in the prior or from changes in the identification restrictions.

Under the conditional uniform prior, the objective is to draw from the following posterior of the orthogonal reduced-form parameters conditional on the sign restrictions:

$$p(\mathbf{B},\mathbf{\Sigma},\mathbf{Q}\mid(\mathbf{y}_t)_{t=1}^T,\boldsymbol{S}_R(\mathbf{B},\mathbf{\Sigma},\mathbf{Q})>0)=\frac{[\boldsymbol{S}_R(\mathbf{B},\mathbf{\Sigma},\mathbf{Q})>0]\;CUNIW(\tilde{\nu},\tilde{\mathbf{\Phi}},\tilde{\mathbf{\Psi}},\tilde{\mathbf{\Omega}})}{Pr\left(\boldsymbol{S}_R(\mathbf{B},\mathbf{\Sigma},\mathbf{Q})>0\mid(\mathbf{y}_t)_{t=1}^T\right)},\tag{9}$$

and then use f and ϕ to transform the draws to the desired vector of objects of interests such as the structural parameters or impulse responses. It is straightforward to adapt the traditional approach described in Algorithm 1 to obtain draws from Equation (9) as follows:

Algorithm 3. *This algorithm independently draws from* $p(\mathbf{B},\mathbf{\Sigma},\mathbf{Q}\mid(\mathbf{y}_t)_{t=1}^T,\boldsymbol{S}_R(\mathbf{B},\mathbf{\Sigma},\mathbf{Q})>0)$ *as described in Equation* (9)*.*

1. *Draw* $(\mathbf{B},\mathbf{\Sigma})$ *independently from the* $NIW(\tilde{\nu},\tilde{\mathbf{\Phi}},\tilde{\mathbf{\Psi}},\tilde{\mathbf{\Omega}})$ *distribution.*

2. *Draw* $\mathbf{Q}$ *independently from the uniform distribution over* $\mathcal{O}(n)$ *until* $\boldsymbol{S}_R(\mathbf{B},\mathbf{\Sigma},\mathbf{Q})>0$*.*

3. *Repeat Steps 1 and 2 until the desired number of draws is obtained.*

While Algorithm 3 does not sample from the posterior distribution defined in Equation (7), it draws from the posterior distribution defined in Equation (9) and it can be justified under a prior different from that in Equation (6).

We now illustrate how, under the conditionally uniform prior approach, inference and identification become intertwined. Consider an SVAR with $n=3$ and $m=0$, that is, without lags or constant terms. Suppose that Researcher A aims to identify three structural shocks using identification scheme A, as defined in Table 4, while Researcher B employs identification

scheme B, also defined in the same table. Clearly, any set of impulse responses satisfying scheme B will also satisfy scheme A.

	Identification A		
	Shock 1	Shock 2	Shock 3
Variable 1	+1	+1	+1
Variable 2	+1	−1	+1
Variable 3	+1		−1

	Identification B		
	Shock 1	Shock 2	Shock 3
Variable 1	+1	+1	+1
Variable 2	+1	−1	+1
Variable 3	+1	−1	−1

Table 4: Sign restrictions: +1 and −1 indicate positive and negative sign restrictions, respectively; blanks indicate no restriction.

Define $\mathbb{Q}_n^j(\mathbf{\Sigma}) = \{\mathbf{Q} \in \mathbb{Q}_n : \mathbf{S}_R^j(\mathbf{\Sigma}, \mathbf{Q}) > 0\}$ as the set of valid rotation matrices for researcher $j \in \{A, B\}$, with associated $\kappa^j(\mathbf{\Sigma})$ implicitly defined by $\int_{\mathbb{Q}_n^j(\mathbf{\Sigma})} \kappa^j(\mathbf{\Sigma}) d\,\mathbf{Q} = 1$. Even if both researchers specify the same reduced-form prior, changing the identification scheme alters the implied prior over the impulse responses. To illustrate, set $\nu = 100$ and $\mathbf{\Phi} = \mathbf{I}_n$, and consider ten values of $\{\mathbf{\Sigma}^i\}_{i=1}^{10}$ such that when evaluated at the prior distribution $IW_{(\nu,\mathbf{\Phi})}$ they all have identical prior density. Let $\{\mathbf{L}_0^i\}_{i=1}^{10}$ denote ten corresponding impact impulse response matrices that happen to satisfy identification B. Because the volume element from $(\mathbf{\Sigma}, \mathbf{Q})$ to $\mathbf{L}_0$ only depends on the determinant of $\mathbf{\Sigma}$ and we have also restricted $\{\mathbf{\Sigma}^i\}_{i=1}^{10}$ to have the same determinant, all $\{\mathbf{L}_0^i\}_{i=1}^{10}$ are equally favored by the prior before any sign restrictions are imposed. After the sign restrictions are introduced, we have:

$$\frac{\pi^j(\mathbf{L}_0^i)}{\pi^j(\mathbf{L}_0^{i'})} = \frac{\kappa^j(\mathbf{\Sigma}^i)}{\kappa^j(\mathbf{\Sigma}^{i'})} \text{ for } i, i' = 1, \ldots, 10, \text{ and } j \in \{A, B\},$$

where $\pi^j(\mathbf{L}_0)$ is the prior density over impulse responses under researcher $j \in \{A, B\}$. This expression implies that any variation in this ratio is solely attributable to the additional

negative sign restriction on shock 2 to variable 3 associated with identification scheme B. As shown in Table 5, the difference in the identification schemes leads to differences in the implied priors. On the one hand, Researcher A's prior favors $\mathbf{L}_0^5$ 1.45 times as much as $\mathbf{L}_0^1$, while Researcher B's prior favors $\mathbf{L}_0^1$ about twice as much as $\mathbf{L}_0^4$. On other hand, Researcher B's prior favors $\mathbf{L}_0^3$ twice as much as $\mathbf{L}_0^1$, while Researcher A's prior favors $\mathbf{L}_0^1$ 1.13 times as much as $\mathbf{L}_0^3$. These findings demonstrate that, under the conditionally uniform approach, changing identification schemes alters the implied prior over parameters of interest (such as impulse responses), thereby entangling estimation and identification.

Draw	1	2	3	4	5	6	7	8	9	10
$\pi^A(\mathbf{L}_0^i)/\pi^A(\mathbf{L}_0^1)$	1.00	1.29	0.89	0.62	1.45	1.52	0.46	0.07	1.24	0.41
$\pi^B(\mathbf{L}_0^i)/\pi^B(\mathbf{L}_0^1)$	1.00	1.60	1.88	0.25	0.58	0.83	0.26	0.03	1.00	0.31

Table 5: Ratio of priors across draws of $\mathbf{L}_0^i$.

The reason for this unfortunate result is that $\kappa(\mathbf{\Sigma})$ varies across identification schemes in a way that disproportionately favors $\mathbf{L}_0$ values associated with $\mathbf{\Sigma}$ that induce smaller identified sets. Since different sign restrictions affect the size of the identified set differently for each $\mathbf{\Sigma}$, different identification restrictions will imply different priors over impulse response functions. This problem does not arise under the uniform prior described in Equation (6), since in that case κ does not depend on the reduced-form parameters. More broadly, this highlights the cost of not adopting a uniform prior over the orthogonal matrices as described in Arias, Rubio-Ramírez, and Waggoner (2025). As shown in that paper, specifying a uniform prior over the set of orthogonal matrices ensures disentangling inference from identification.

9 Conclusion

This paper proposes a Gibbs sampling algorithm for structural vector autoregressions identified with sign restrictions. We show that the algorithm effectively overcomes the computational bottlenecks associated with conventional accept-reject methods, especially as the number

of identifying restrictions increases or as the identified set becomes tight. Our empirical applications illustrate how the proposed algorithm can extend existing analyses in the literature, including SVARs with a large-number of macroeconomic and financial variables. Overall, the paper provides contributions to the implementation of sign-restricted SVARs, offering tools that are applicable across a wide range of empirical models.

References

Amir-Ahmadi, P. and T. Drautzburg (2021). Identification and Inference with Ranking Restrictions. *Quantitative Economics 12*(1), 1–39.

Antolín-Díaz, J. and J. F. Rubio-Ramírez (2018, October). Narrative sign restrictions for svars. *American Economic Review 108*(10), 2802–29.

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.

Arias, J. E., J. F. Rubio-Ramírez, and D. F. Waggoner (2025). Uniform Priors for Impulse Responses. *Econometrica 93*(2), 695–718.

Bańbura, M., D. Giannone, and L. Reichlin (2010). Large Bayesian Vector Auto Regressions. *Journal of Applied Econometrics 25*(1), 71–92.

Baumeister, C. and J. D. Hamilton (2024). Advances in Using Vector Autoregressions to Estimate Structural Magnitudes. *Econometric Theory 40*(3), 472–510.

Bernanke, B. S., J. Boivin, and P. Eliasz (2005, 02). Measuring the Effects of Monetary Policy: A Factor-Augmented Vector Autoregressive (FAVAR) Approach. *The 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.

Chan, J. C. C. (2022). Asymmetric Conjugate Priors for Large Bayesian VARs. *Quantitative Economics 13*(3), 1145–1169.

Chan, J. C. C., C. Matthes, and X. Yu (2025). Large Structural VARs with Multiple Sign and Ranking Restrictions.

Crump, R. K., S. Eusepi, D. Giannone, E. Qian, and A. M. Sbordone (2025). A Large Bayesian VAR of the United States Economy. *International Journal of Central Banking (Forthcoming)*.

Faust, J. (1998). The Robustness of Identified VAR Conclusions about Money. *Carnegie-Rochester Conference Series on Public Policy 49*, 207–244.

Giannone, D., M. Lenza, and G. E. Primiceri (2015). Prior Selection for Vector Autoregressions. *Review of Economics and Statistics 97*(2), 436–451.

Giannone, D., M. Lenza, and G. E. Primiceri (2021). Economic Predictions With Big Data: The Illusion of Sparsity. *Econometrica 89*(5), 2409–2437.

Granziera, E., H. R. Moon, and F. Schorfheide (2018). Inference for VARs Identified with Sign Restrictions. *Quantitative Economics 9*(3), 1087–1121.

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.

Kilian, L. and D. P. Murphy (2014). The Role of Inventories and Speculative Trading in the Global Market for Crude Oil. *Journal of Applied Econometrics 29*(3), 454–478.

Koop, G. (2013). Forecasting with Medium and Large Bayesian VARs. *Journal of Applied Econometrics 28*(2), 177–203.

Leeper, E. M., C. A. Sims, and T. Zha (1996). What Does Monetary Policy Do? *Brookings Papers on Economic Activity 1996*(2), 1–78.

Ludvigson, S. C., S. Ma, and S. Ng (2021, October). Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response? *American Economic Journal: Macroeconomics 13*(4), 369–410.

Murray, I., R. Adams, and D. MacKay (2010). Elliptical Slice Sampling. *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics 9*, 541–548.

Natarovskii, V., D. Rudolf, and B. Sprungk (2021). Geometric Convergence of Elliptical Slice Sampling. *Proceedings of the 38th International Conference on Machine Learning 139*, 7969–7978.

Read, M. and D. Zhu (2025). Fast Posterior Sampling in Tightly Identified SVARs Using "Soft" Sign Restrictions. *Federal Reserve Bank of Australia Working Paper* (2025-03).

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.

Sims, C. A. and T. Zha (1998). Bayesian Methods for Dynamic Multivariate Models. *International Economic Review 39*(4), 949–968.

Uhlig, H. (1994). What Macroeconomists Should Know About Unit Roots: A Bayesian Perspective. *Econometric Theory 10*(3-4), 645–671.

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.

Uhlig, H. (2017). Shocks, Sign Restrictions, and Identification. *Advances in Economics and Econometrics 2*, 95.

Vats, D., J. M. Flegal, and G. L. Jones (2019). Multivariate Output Analysis for Markov Chain Monte Carlo. *Biometrika 106*(2), 321–337.

Appendix

I Initialization

Our algorithm requires initial values for $(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})$ that satisfy the sign restrictions imposed to identify the SVAR. We accomplish this by setting $(\mathbf{B}, \mathbf{\Sigma})$ to the maximum likelihood estimates and by constructing $\mathbf{Q}$ one column at a time such that the sign restrictions hold, as described in Algorithm 4. This strategy works as long as the researcher does not imposes cross equation sign restrictions.

Algorithm 4. *The following algorithm sequentially builds an orthogonal matrix* $\mathbf{Q}$ *subject to sign restrictions. Let* $\mathbf{\Sigma} \in \mathcal{S}(n)$ *and define* $\mathbf{Q}_j = \begin{bmatrix} \mathbf{Q}_1 & \cdots & \mathbf{Q}_j \end{bmatrix}$ *for* $1 \leq j \leq n$. *Let* $[\boldsymbol{S}_j(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}_j) > 0]$ *denote sign restrictions associated with shock* j *for* $j = 1, \ldots, n$, *and assume the researcher imposes sign restrictions for the first* $m \leq n$ *shocks. The restrictions are satisfied when* $\prod_{j=1}^{m} [\boldsymbol{S}_j(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}_j) > 0] = 1$.

1. *Let* $j = 1$.

2. *Draw* $\mathbf{x}_j \in \mathbb{R}^n$ *independently from a standard normal distribution.*

3. *If* $j \leq m$, *let*

$$\mathbf{x}_j = \mathbf{x}_j - \mathbf{Q}_{j-1} \mathbf{Q}_{j-1}' \mathbf{x}_j.$$

   and set $\mathbf{Q}_j = \mathbf{x}_j / \| \mathbf{x}_j \|$ *and return to Step 2 until* $[\boldsymbol{S}_j(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}_j) > 0] = 1$.

4. *If* $j > m$,

$$\mathbf{Q}_j = \mathbf{x}_j / \| \mathbf{x}_j \| .$$

5. *If* $j < n$, *set* $j = j + 1$ *and return to Step 2.*

The Gaussian draw in Step 2 is invariant to rotations; projecting onto the orthogonal complement and re-normalizing yields a vector uniformly distributed on the unit sphere

of that subspace. Rejection sampling in Step 3 enforces user-defined sign restrictions. For columns $j > m$ no restrictions apply.

We note that a similar idea is employed by Amir-Ahmadi and Drautzburg (2021), who generate draws from the posterior distribution $p(\mathbf{Q} \mid \mathbf{B}, \mathbf{\Sigma}, (\mathbf{y}_t)_{t=1}^T, \boldsymbol{S}_R(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) > 0)$ by sequentially sampling each column $\mathbf{q}_j$ of $\mathbf{Q}$ conditional on the remaining columns, that is, from $p(\mathbf{q}_j \mid \mathbf{Q}_{-j}, \mathbf{B}, \mathbf{\Sigma}, (\mathbf{y}_t)_{t=1}^T, \boldsymbol{S}_R(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q}) > 0)$ for $j = 1, \ldots, n$, where $\mathbf{Q}_{-j}$ denotes the matrix $\mathbf{Q}$ with the j-th column removed. In such cases, the projection step in Algorithm 4 (Step 3) is modified accordingly. In both approaches, if the sign restrictions are linear in $\mathbf{q}_j$ (that is, of the form $\mathbf{A}\,\mathbf{q}_j \geq 0$) sampling efficiency can be improved. Specifically, these restrictions are satisfied as long as $\mathbf{A}(\mathbf{I} - \mathbf{Q}_{j-1}\,\mathbf{Q}_{j-1}')\,\mathbf{x}_j \geq 0$, where $\mathbf{x}_j$ is a draw from a standard multivariate normal distribution. Hence, one can generate draws of $\mathbf{x}_j$ that satisfy the sign restrictions using efficient routines for sampling from truncated multivariate normal distributions with linear inequality constraints.

II Alternative Prior Specifications

Our Gibbs sampler algorithm can be adapted to work with two popular priors in SVAR analysis: the independent normal-inverse-Wishart prior, and the asymmetric priors proposed by Chan (2022).

II.1 Independent normal-inverse-Wishart Prior

Let us begin showing how to adapt our Gibbs sampler to the case in which a researcher aims to use an independent normal-inverse-Wishart prior for $(\mathbf{B}, \mathbf{\Sigma})$ of the form $IW_{(\bar{\nu}, \bar{\mathbf{\Phi}})}(\mathbf{\Sigma}) N_{(\mathrm{vec}(\bar{\boldsymbol{\mu}}_{\mathbf{B}}), \bar{\mathbf{V}}_{\mathbf{B}})}(\mathrm{vec}(\mathbf{B}))$. When using this prior, the final objective of the researcher is to sample from the following

posterior:

$$p(\mathbf{B},\boldsymbol{\Sigma},\mathbf{Q}\mid(\mathbf{y}_t)_{t=1}^T,\boldsymbol{S}_R(\mathbf{B},\boldsymbol{\Sigma},\mathbf{Q})>0)\propto[\boldsymbol{S}_R(\mathbf{B},\boldsymbol{\Sigma},\mathbf{Q})>0]$$
$$\times NIW_{(\hat{\nu},\hat{\boldsymbol{\Phi}},\hat{\boldsymbol{\Psi}},\hat{\boldsymbol{\Omega}})}(\mathbf{B},\boldsymbol{\Sigma})IW_{(\bar{\nu},\bar{\boldsymbol{\Phi}})}(\boldsymbol{\Sigma})N_{(\mathrm{vec}(\bar{\boldsymbol{\mu}}_{\mathbf{B}}),\bar{\mathbf{V}}_{\mathbf{B}})}(\mathrm{vec}(\mathbf{B})) \tag{A.1}$$

where $NIW_{(\hat{\nu},\hat{\boldsymbol{\Phi}},\hat{\boldsymbol{\Psi}},\hat{\boldsymbol{\Omega}})}(\mathbf{B},\boldsymbol{\Sigma})$ denotes the likelihood, with $\hat{\nu}=T-m-(n+1)$, $\hat{\boldsymbol{\Omega}}=(\mathbf{X}'\mathbf{X})^{-1}$, $\hat{\boldsymbol{\Psi}}=\hat{\boldsymbol{\Omega}}\mathbf{X}'\mathbf{Y}$, and $\hat{\boldsymbol{\Phi}}=\mathbf{Y}'\mathbf{Y}-\hat{\boldsymbol{\Psi}}'\hat{\boldsymbol{\Omega}}^{-1}\hat{\boldsymbol{\Psi}}$, where $\mathbf{Y}'=(\mathbf{y}_1,\ldots,\mathbf{y}_t)$ and $\mathbf{X}'=(\mathbf{x}_1,\ldots,\mathbf{x}_t)$. Table 6 describes the conditional posterior distributions obtained from Equation (A.1)—using similar arguments than those given in Section 6—where

$$\tilde{\nu}=\bar{\nu}+T,$$
$$\tilde{\boldsymbol{\Phi}}=\bar{\boldsymbol{\Phi}}+(\mathbf{Y}-\mathbf{X}\mathbf{B})'(\mathbf{Y}-\mathbf{X}\mathbf{B}),$$
$$\tilde{\mathbf{V}}_{\mathbf{B}}^{-1}=\bar{\mathbf{V}}_{\mathbf{B}}^{-1}+(\boldsymbol{\Sigma}^{-1}\otimes\mathbf{X}'\mathbf{X}),\text{and}$$
$$\tilde{\mu}_{\mathbf{B}}=\tilde{\mathbf{V}}_{\mathbf{B}}(\bar{\mathbf{V}}_{\mathbf{B}}^{-1}\bar{\boldsymbol{\mu}}_{\mathbf{B}}+(\boldsymbol{\Sigma}^{-1}\otimes\mathbf{X}'\mathbf{X})\hat{\boldsymbol{\Psi}}).$$

$p(\mathbf{Q}\mid\mathbf{B},\boldsymbol{\Sigma},(\mathbf{y}_t)_{t=1}^T,\boldsymbol{S}_R(\mathbf{B},\boldsymbol{\Sigma},\mathbf{Q})>0)$	$\propto[\boldsymbol{S}_R(\mathbf{B},\boldsymbol{\Sigma},\mathbf{Q})>0]$
$p(\boldsymbol{\Sigma}\mid\mathbf{B},\mathbf{Q},(\mathbf{y}_t)_{t=1}^T,\boldsymbol{S}_R(\mathbf{B},\boldsymbol{\Sigma},\mathbf{Q})>0)$	$\propto[\boldsymbol{S}_R(\mathbf{B},\boldsymbol{\Sigma},\mathbf{Q})>0]IW_{(\tilde{\nu},\tilde{\boldsymbol{\Phi}})}(\boldsymbol{\Sigma})$
$p(\mathbf{B}\mid\boldsymbol{\Sigma},\mathbf{Q},(\mathbf{y}_t)_{t=1}^T,\boldsymbol{S}_R(\mathbf{B},\boldsymbol{\Sigma},\mathbf{Q})>0)$	$\propto[\boldsymbol{S}_R(\mathbf{B},\boldsymbol{\Sigma},\mathbf{Q})>0]N_{(\mathrm{vec}(\tilde{\boldsymbol{\mu}}_{\mathbf{B}}),\tilde{\mathbf{V}}_{\mathbf{B}})}(\mathrm{vec}(\mathbf{B}))$

Table 6: Conditional Posterior Distributions

Hence, we now can write a Gibbs Sampler of the following form:

Algorithm 5. *This algorithm draws from* $p(\mathbf{B},\boldsymbol{\Sigma},\mathbf{Q}\mid(\mathbf{y}_t)_{t=1}^T,\boldsymbol{S}_R(\mathbf{B},\boldsymbol{\Sigma},\mathbf{Q})>0)$ *described in Equation* (A.1).

1. *Let* $I>1$ *and set* $i=1$ *and assign initial values to* $(\mathbf{B}^{i-1},\boldsymbol{\Sigma}^{i-1})$.

2. *Draw* $\mathbf{Q}^i$ *from*

$$p(\mathbf{Q}\mid\mathbf{B}^{i-1},\boldsymbol{\Sigma}^{i-1},(\mathbf{y}_t)_{t=1}^T,\boldsymbol{S}_R(\mathbf{B}^{i-1},\boldsymbol{\Sigma}^{i-1},\mathbf{Q})>0)\propto[\boldsymbol{S}_R(\mathbf{B}^{i-1},\boldsymbol{\Sigma}^{i-1},\mathbf{Q})>0].$$

3. *Draw* $\boldsymbol{\Sigma}^i$ *from*

$$p(\boldsymbol{\Sigma} \mid \mathbf{B}^{i-1}, \mathbf{Q}^i, (\mathbf{y}_t)_{t=1}^T, \boldsymbol{S}_R(\mathbf{B}^{i-1}, \boldsymbol{\Sigma}, \mathbf{Q}^i) > 0) \propto [\boldsymbol{S}_R(\mathbf{B}^{i-1}, \boldsymbol{\Sigma}, \mathbf{Q}^i) > 0] IW_{(\tilde{\nu}, \tilde{\boldsymbol{\Phi}})}(\boldsymbol{\Sigma}).$$

4. *Draw* $\mathbf{B}^i$ *from*

$$p(\mathbf{B} \mid \boldsymbol{\Sigma}^i, \mathbf{Q}^i, (\mathbf{y}_t)_{t=1}^T, \boldsymbol{S}_R(\mathbf{B}, \boldsymbol{\Sigma}^i, \mathbf{Q}^i) > 0) \propto [\boldsymbol{S}_R(\mathbf{B}, \boldsymbol{\Sigma}^i, \mathbf{Q}^i) > 0] N_{(vec(\tilde{\boldsymbol{\mu}}_{\mathbf{B}}), \tilde{\mathbf{V}}_{\mathbf{B}})}(vec(\mathbf{B})).$$

5. *If* $i < I$*, let* $i = i + 1$ *and return to Step 2.*

II.2 Asymmetric Priors

Instead of directly working with the typical conjugate normal-inverse-Wishart prior over the reduced-form parameters $(\mathbf{B}, \boldsymbol{\Sigma})$, Chan (2022) works with structural parameters subject to a recursive identification scheme and denotes the resulting parameterization as $(\boldsymbol{\theta}, \boldsymbol{\sigma}^2)$. This parameterization implies a prior over the reduced-form parameters $(\mathbf{B}, \boldsymbol{\Sigma})$ that allows the shrinkage strength for coefficients on own lags to be different than the shrinkage strength for coefficients on the lag of other variables.

For completeness, we will describe the mapping from $(\boldsymbol{\theta}, \boldsymbol{\sigma}^2)$ to $(\mathbf{B}, \boldsymbol{\Sigma})$. To this end, it is useful to begin by considering the recursive SVAR proposed by Chan (2022):

$$\mathbf{A}^S \mathbf{y}_t = \mathbf{b}^S + \mathbf{B}_1^S \mathbf{y}_{t-1} + \cdots + \mathbf{B}_1^S \mathbf{y}_{t-p} + \boldsymbol{\varepsilon}_t^y, \quad \boldsymbol{\varepsilon}_t^y \sim N(\mathbf{0}, \boldsymbol{\Sigma}^S) \tag{A.2}$$

where $\mathbf{A}^S$ is a lower triangular matrix with ones along the diagonal, $\mathbf{b}^S$, $\mathbf{B}_1^S$,..., $\mathbf{B}_p^S$, denote a constant term and the slope coefficients, and $\boldsymbol{\Sigma}^S = \text{diag}(\boldsymbol{\sigma}^2) = \text{diag}([\sigma_1^2, \ldots, \sigma_n^2])$. Given this representation, we use the following notation: b_i^S denotes the i-th element of $\mathbf{b}^S$, $\mathbf{b}_{j,i}^S$ represents the i-th row of $\mathbf{B}_j^S$, $\boldsymbol{\beta}_i = (b_i^S, \mathbf{b}_{1,i}^S, \ldots, \mathbf{b}_{p,i}^S)'$ collects all the VAR slope coefficients associated with the i-th equation, and $\boldsymbol{\alpha}_i = (A_{i,1}, \ldots, A_{i,i-1})'$ where $A_{i,j}$ denotes the i-th row and the j-th column of $\mathbf{A}$.

For $i = 1, \ldots, n$, let $\boldsymbol{\theta}_i = [\boldsymbol{\beta}_i', \boldsymbol{\alpha}_i']'$, where $\boldsymbol{\beta}_i$ is a vector containing the slope coefficients of the i-th equation of a recursive SVAR and $\boldsymbol{\alpha}_i$ is a vector containing the unrestricted elements of the i-th row of the inverse of the impact impulse response matrix. Furthermore, set $\boldsymbol{\theta} = (\boldsymbol{\theta}_1', \ldots, \boldsymbol{\theta}_n')'$. If we define $\mathbf{w}_{i,t}^S = (-y_{1,t}, \ldots, -y_{i-1,t})$, $\mathbf{x}_t^S = (1, \mathbf{y}_{t-1}', \ldots, \mathbf{y}_{t-p}')$, $\mathbf{x}_{i,t} = (\mathbf{w}_{i,t}^S, \mathbf{x}_t^S)$, $\mathbf{y}_i = (y_{i,1}, \ldots, y_{i,T})'$, and $\mathbf{X}_i = (\mathbf{x}_{i,1}, \ldots, \mathbf{x}_{i,T})'$, then $\mathbf{y}_i = \mathbf{X}_i \boldsymbol{\theta}_i + \boldsymbol{\varepsilon}_i^y$, $\boldsymbol{\varepsilon}_i^y \sim N(\mathbf{0}, \sigma_i^2 \mathbf{I}_T)$. Importantly, equation (A.2) can be written as

$$\mathbf{y}_t = \mathbf{b} + \mathbf{B}_1 \mathbf{y}_{t-1} + \cdots + \mathbf{B}_p \mathbf{y}_{t-p} + \boldsymbol{\varepsilon}_t^y, \quad \boldsymbol{\varepsilon}_t^y \sim N(\mathbf{0}, \boldsymbol{\Sigma}) \tag{A.3}$$

where $\mathbf{b} = (\mathbf{A}^S)^{-1} \mathbf{b}^S$, $\mathbf{B}_j = (\mathbf{A}^S)^{-1} \mathbf{B}_j^S$ for $j = 1, \ldots, p$, $\mathbf{B} = [\mathbf{b}, \mathbf{B}_1, \ldots, \mathbf{B}_p]'$, and $\boldsymbol{\Sigma} = (\mathbf{A}^S)^{-1} \boldsymbol{\Sigma}^S ((\mathbf{A}^S)^{-1})'$. Hence, we have implicitly defined a mapping from $(\boldsymbol{\theta}, \boldsymbol{\sigma}^2)$ to $(\mathbf{B}, \boldsymbol{\Sigma})$.

The mapping is invertible. Given $(\mathbf{B}, \boldsymbol{\Sigma})$ we can obtain $(\boldsymbol{\theta}, \boldsymbol{\sigma}^2)$ by using the Cholesky composition to obtain $\mathbf{A}^S$ and $\boldsymbol{\Sigma}^S$, and then we can construct $\boldsymbol{\beta}_i$ by using $\mathbf{b}^S = \mathbf{A}^S \mathbf{b}$, $\mathbf{B}_1^S = \mathbf{A}^S \mathbf{B}_1$, ... $\mathbf{B}_p^S = \mathbf{A}^S \mathbf{B}_p$, for $i = 1, \ldots, n$. Putting the mapping defined above with the mapping ϕ, we obtain a one-to-mapping between the impulse responses, that is $(\mathbf{L}_0, \mathbf{L}_+)$, and $(\boldsymbol{\theta}, \boldsymbol{\sigma}^2, \mathbf{Q})$.

Exploiting Proposition 2 in Arias, Rubio-Ramírez, and Waggoner (2025), we will combine the product of normal-inverse-gamma over $(\boldsymbol{\theta}, \boldsymbol{\sigma}^2)$ proposed by Chan (2022), i.e.,

$$p(\boldsymbol{\theta}, \boldsymbol{\sigma}^2) = \prod_{i=1}^n p(\boldsymbol{\theta}_i, \sigma_i^2) = \prod_{i=1}^n IG_{(\nu_i, S_i)}(\sigma_i^2) N_{(\mathbf{m}_i, \sigma_i^2 \mathbf{V}_i)}(\boldsymbol{\theta}_i) \tag{A.4}$$

where $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_n)$ and $\boldsymbol{\sigma}^2 = (\sigma_1^2, \ldots, \sigma_n^2)$, with a uniform prior over $\mathbf{Q}$ under the Haar measure in order to induce a prior over $(\mathbf{L}_0, \mathbf{L}_+)$. Then, tracing a parallel to the case in which we work with the orthogonal reduced-form parameters $(\mathbf{B}, \boldsymbol{\Sigma}, \mathbf{Q})$, we will conduct posterior inference over $(\mathbf{L}_0, \mathbf{L}_+)$ subject to sign restrictions by sampling from the following posterior

$$p(\boldsymbol{\theta}, \boldsymbol{\sigma}^2, \mathbf{Q} \mid (\mathbf{y}_t)_{t=1}^T, \boldsymbol{S}(\boldsymbol{\theta}, \boldsymbol{\sigma}^2, \mathbf{Q}) > 0) \propto [\boldsymbol{S}(\boldsymbol{\theta}, \boldsymbol{\sigma}^2, \mathbf{Q}) > 0] \prod_i^n IG_{(\tilde{\nu}_i, \tilde{S}_i)}(\sigma_i^2) N_{(\tilde{\boldsymbol{\theta}}_i, \sigma_i^2 \mathbf{K}_{\tilde{\boldsymbol{\theta}}_i}^{-1})}(\boldsymbol{\theta}_i), \tag{A.5}$$

and then converting the draws to the desired parameterization; where $\tilde{\nu}_i = \nu_i + \frac{T}{2}$, $\tilde{S}_i = S_i + (\mathbf{y}_i'\mathbf{y}_i + \mathbf{m}_i'\mathbf{V}_i^{-1}\mathbf{m}_i - \tilde{\boldsymbol{\theta}}_i'\mathbf{K}_{\boldsymbol{\theta}_i}\tilde{\boldsymbol{\theta}}_i)/2$, $\mathbf{y}_i = (\mathbf{y}_{i1}, \ldots, \mathbf{y}_{iT})'$, $\mathbf{y}_{it}$ denotes the i-th entry of $\mathbf{y}_t$, $\mathbf{K}_{\boldsymbol{\theta}_i} = \mathbf{V}_i^{-1} + \mathbf{X}_i'\mathbf{X}_i$, and $\tilde{\boldsymbol{\theta}}_i = \mathbf{K}_{\boldsymbol{\theta}_i}^{-1}(\mathbf{V}_i^{-1}\mathbf{m}_i + \mathbf{X}_i'\mathbf{y}_i)$.

Conditional Posterior for Q. Let us begin by deriving the posterior for $\mathbf{Q}$ given the reduced-form parameters and the sign restrictions. Equation (A.5) implies that

$$p(\mathbf{Q} \mid \boldsymbol{\theta}, \boldsymbol{\sigma}^2, (\mathbf{y}_t)_{t=1}^T, \boldsymbol{S}(\boldsymbol{\theta}, \boldsymbol{\sigma}^2, \mathbf{Q}) > 0) = \frac{[\boldsymbol{S}(\boldsymbol{\theta}, \boldsymbol{\sigma}^2, \mathbf{Q}) > 0]}{\int_{\mathcal{O}(n)} [\boldsymbol{S}(\boldsymbol{\theta}, \boldsymbol{\sigma}^2, \mathbf{Q}) > 0] d\mathbf{Q}} \propto [\boldsymbol{S}(\boldsymbol{\theta}, \boldsymbol{\sigma}^2, \mathbf{Q}) > 0].$$

Hence, we can sample from this conditional posterior as described in Section 6.

Conditional Posterior for $\boldsymbol{\sigma}^2$. Next, we derive a useful expression for the posterior of $\boldsymbol{\sigma}^2$ given $(\boldsymbol{\theta}, \mathbf{Q})$ and the sign restrictions. In this case, Equation (A.5) implies that

$$p(\boldsymbol{\sigma}^2 \mid \boldsymbol{\theta}, \mathbf{Q}, (\mathbf{y}_t)_{t=1}^T, \boldsymbol{S}(\boldsymbol{\theta}, \boldsymbol{\sigma}^2, \mathbf{Q}) > 0) \propto [\boldsymbol{S}(\boldsymbol{\theta}, \boldsymbol{\sigma}^2, \mathbf{Q}) > 0] \prod_{i=1}^n IG_{(\tilde{\nu}_{\boldsymbol{\sigma}^2,i}, \tilde{S}_{\boldsymbol{\sigma}^2,i})}(\sigma_i^2),$$

where $\tilde{\nu}_{\boldsymbol{\sigma}^2,i} = \nu_i + \frac{T+k_i}{2}$ and $\tilde{S}_{\boldsymbol{\sigma}^2,i} = \tilde{S}_i + \frac{1}{2}(\boldsymbol{\theta}_i - \tilde{\boldsymbol{\theta}}_i)' K_{\boldsymbol{\theta}_i}(\boldsymbol{\theta}_i - \tilde{\boldsymbol{\theta}}_i)/2$. Drawing from the conditional posterior $p(\boldsymbol{\sigma} \mid \boldsymbol{\theta}, \mathbf{Q}, (\mathbf{y}_t)_{t=1}^T, \boldsymbol{S}(\boldsymbol{\theta}, \boldsymbol{\sigma}^2, \mathbf{Q}) > 0)$ follows straightforwardly from Section 6 once we notice that $IG_{(\tilde{\nu}_{\boldsymbol{\sigma}^2,i}, \tilde{S}_{\boldsymbol{\sigma}^2,i})}(\sigma_i^2)$ is equivalent to a univariate inverse-Wishart $IW_{(2\tilde{\nu}_{\boldsymbol{\sigma}^2,i}, 2\tilde{S}_{\boldsymbol{\sigma}^2,i})}(\sigma_i^2)$. Given the indicator function $[\boldsymbol{S}(\boldsymbol{\theta}, \boldsymbol{\sigma}^2, \mathbf{Q}) > 0]$, we draw σ_i^2 conditional on the $\boldsymbol{\sigma}_{-i}^2$ where $\boldsymbol{\sigma}_{-i}^2$ denotes the vector $\boldsymbol{\sigma}^2$ excluding its i-th entry.

Conditional Posterior for $\boldsymbol{\theta}$ The third and last conditional posterior corresponds to the posterior of $\boldsymbol{\theta}$ given $(\boldsymbol{\sigma}^2, \mathbf{Q})$ and the sign restrictions. In this case, we use Equation (A.5) to obtain:

$$p(\boldsymbol{\theta} \mid \boldsymbol{\sigma}^2, \mathbf{Q}, (\mathbf{y}_t)_{t=1}^T, \boldsymbol{S}(\boldsymbol{\theta}, \boldsymbol{\sigma}^2, \mathbf{Q}) > 0) \propto [\boldsymbol{S}(\boldsymbol{\theta}, \boldsymbol{\sigma}^2, \mathbf{Q}) > 0] \prod_{i=1}^n N_{(\hat{\boldsymbol{\theta}}_i, \sigma_i^2 K_{\boldsymbol{\theta}_i}^{-1})}(\boldsymbol{\theta}_i).$$

Given the indicator function $[\boldsymbol{S}(\boldsymbol{\theta}, \boldsymbol{\sigma}^2, \mathbf{Q}) > 0]$, we draw $\boldsymbol{\theta}_i$ conditional on the $\boldsymbol{\theta}_{-i}$ where $\boldsymbol{\theta}_{-i}$ denotes the vector $\boldsymbol{\theta}$ excluding its i-th entry. Having defined the three conditional posteriors described above, we now can write a Gibbs Sampler of the following form:

Algorithm 6. *This algorithm draws from* $p(\boldsymbol{\theta}, \boldsymbol{\sigma}^2, \mathbf{Q} \mid (\mathbf{y}_t)_{t=1}^T, \boldsymbol{S}(\boldsymbol{\theta}, \boldsymbol{\sigma}, \mathbf{Q}) > 0)$ *as described in Equation* (A.5).

1. *Let* $J > 1$ *and set* $j = 1$ *and assign initial values to* $(\boldsymbol{\theta}^{j-1}, \boldsymbol{\sigma}^{2(j-1)})$.

2. *Draw* $\mathbf{Q}^j$ *from*

$$p(\mathbf{Q} \mid \boldsymbol{\theta}^{j-1}, \boldsymbol{\sigma}^{2(j-1)}, (\mathbf{y}_t)_{t=1}^T, \boldsymbol{S}(\boldsymbol{\theta}^{j-1}, \boldsymbol{\sigma}^{2(j-1)}, \mathbf{Q}) > 0) \propto [\boldsymbol{S}(\boldsymbol{\theta}^{j-1}, \boldsymbol{\sigma}^{2(j-1)}, \mathbf{Q}) > 0].$$

3. *Draw* $\boldsymbol{\sigma}^{2(j)}$ *from*

$$p(\boldsymbol{\sigma}^2 \mid \boldsymbol{\theta}^{j-1}, \mathbf{Q}^j, (\mathbf{y}_t)_{t=1}^T, \boldsymbol{S}(\boldsymbol{\theta}^{j-1}, \boldsymbol{\sigma}^2, \mathbf{Q}^j) > 0) \propto [\boldsymbol{S}(\boldsymbol{\theta}^{j-1}, \boldsymbol{\sigma}^2, \mathbf{Q}^j) > 0] \prod_{i=1}^n IG_{(\tilde{\nu}_{\boldsymbol{\sigma}^2,i}, \tilde{S}_{\boldsymbol{\sigma}^2,i})}(\sigma_i^2).$$

4. *Draw* $\boldsymbol{\theta}^j$ *from*

$$p(\boldsymbol{\theta} \mid \boldsymbol{\sigma}^{2(j)}, \mathbf{Q}^j, (\mathbf{y}_t)_{t=1}^T, \boldsymbol{S}(\boldsymbol{\theta}, \boldsymbol{\sigma}^{2(j)}, \mathbf{Q}^j) > 0) \propto [\boldsymbol{S}(\boldsymbol{\theta}, \boldsymbol{\sigma}^{2(j)}, \mathbf{Q}^j) > 0] \prod_{i=1}^n N_{(\hat{\boldsymbol{\theta}}_i, \sigma_i^2 K_{\boldsymbol{\theta}_i}^{-1})}(\boldsymbol{\theta}_i).$$

5. *If* $j < J$, *let* $j = j + 1$ *and return to Step 2.*

III Large-SVAR with Asymmetric Priors

To conclude the appendix, we reproduce the analysis in Section 7.2 using the asymmetric priors proposed by Chan (2022).

III.1 Impulse Responses

We begin by reporting the impulse responses with each of the eight structural shocks identified by Chan, Matthes, and Yu (2025) obtained using the Gibbs sampler described in Algorithm 6. The results reported in Figures 10a-11d are nearly identical to those reported in Chan, Matthes, and Yu (2025). Notice that in this case we can extend the horizon of the impulse responses beyond 5 years without a large increase in posterior uncertainty. This is because the asymmetric priors used in Chan, Matthes, and Yu (2025) allow for more flexibility in terms of shrinkage of the slope coefficients of the reduced-form representation of the SVAR.

III.2 Timing

In this section, we compare the efficiency of our Gibbs sampler algorithm to the one of the accept-reject when using the asymmetric priors. We highlight three results. First, there are cases in which when using asymmetric priors with a relatively small number of sign restrictions, the accept-reject algorithm is somewhat more efficient than our approach, see for example the timings for. Second, as shown in Figure 12, the advantage of the accept-reject vanishes as we increase the number of shocks. Notice that the Gibbs sampler algorithm is faster than the baseline specification in Chan, Matthes, and Yu (2025), which features 8 structural shocks. Third, as shown in Figure 13, as we continue to increase the number of sign restrictions by adding shocks 9 and 10, respectively, the accept-reject algorithm reproduces the explosive patterns in Figures 1-14. In contrast, the efficiency of the Gibbs sampler algorithm remains nearly unchanged.

(a) Demand shock

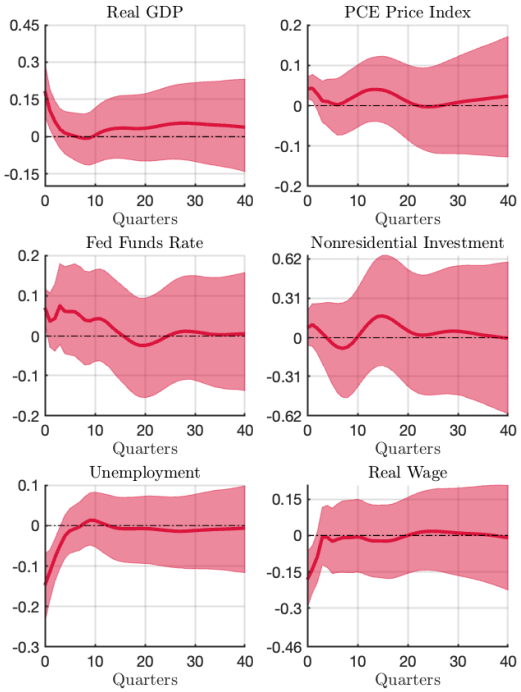

(b) Investment shock

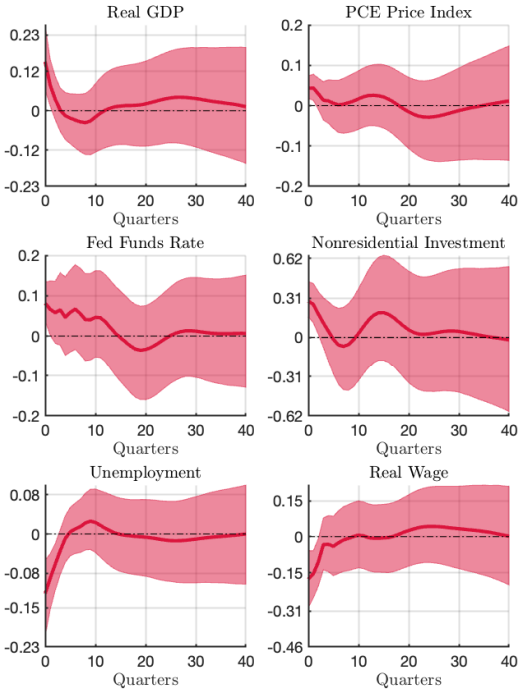

(c) Financial shock

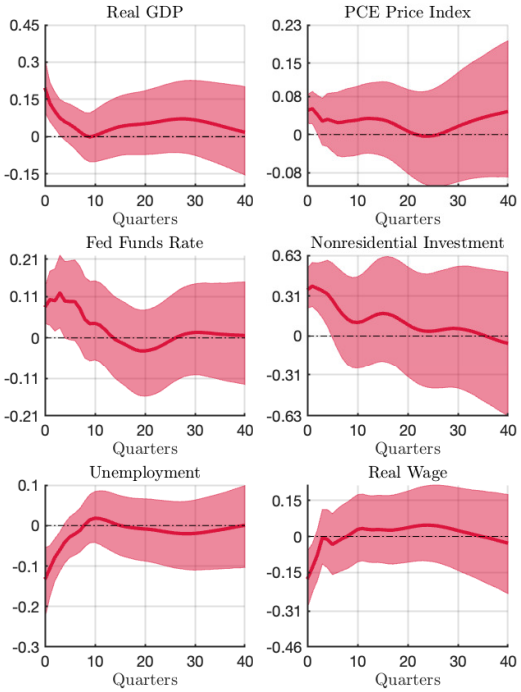

(d) Monetary policy shock

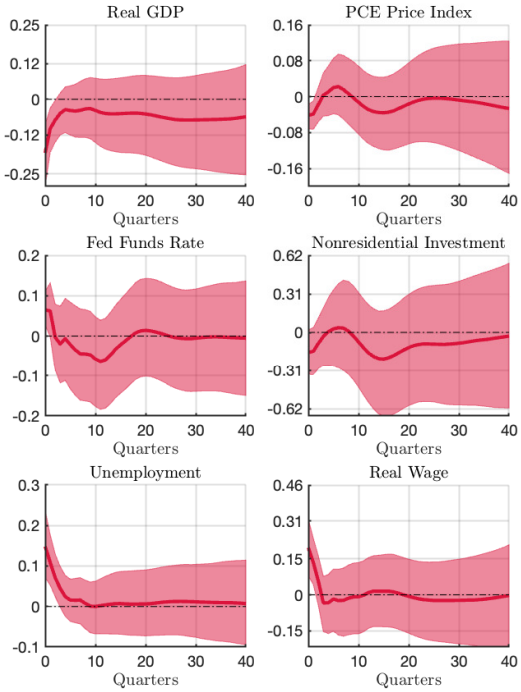

Figure 10: Impulse Responses

(a) Government spending shock

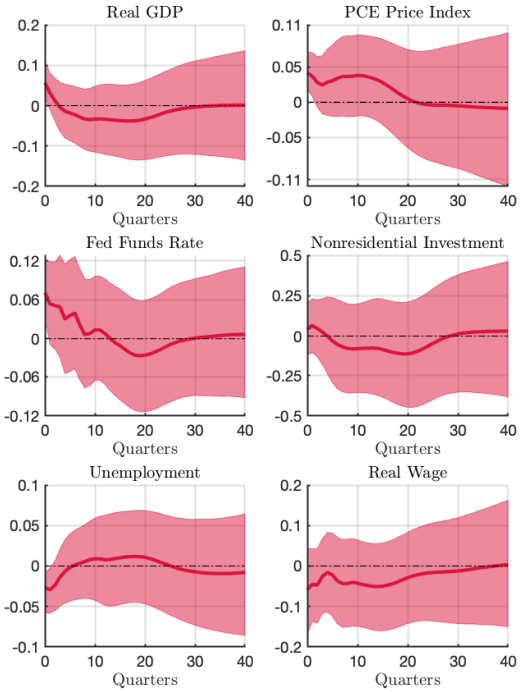

(b) Technology shock

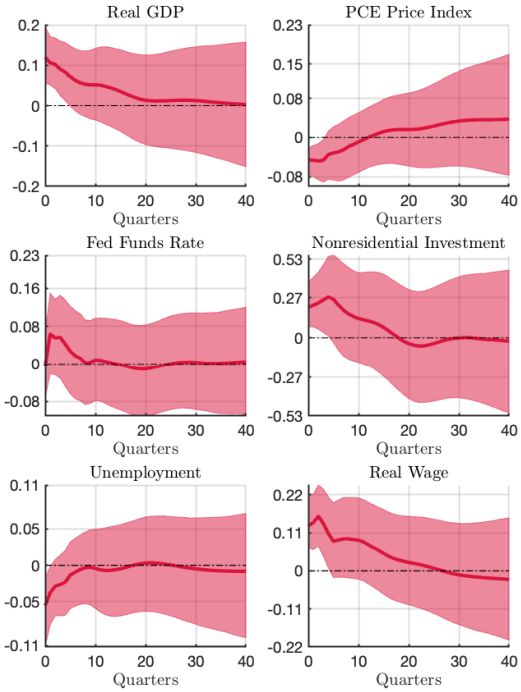

(c) Labor supply shock

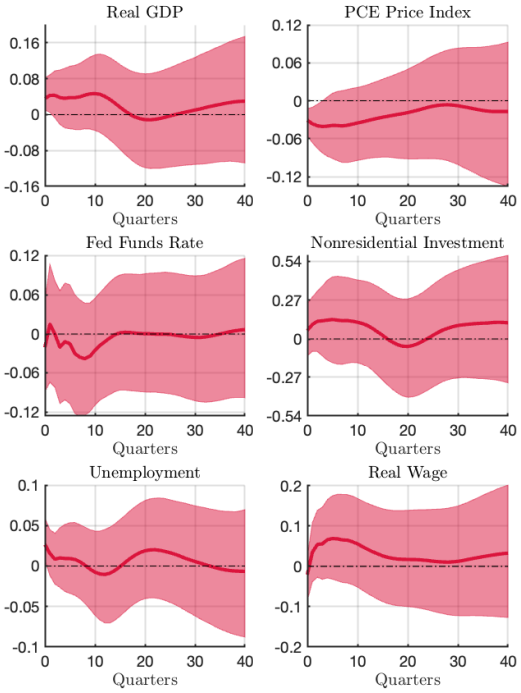

(d) Wage bargaining shock

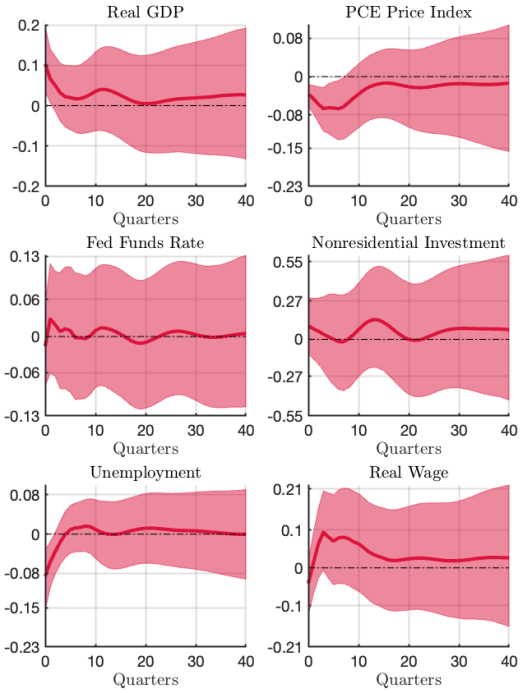

Figure 11: Impulse Responses

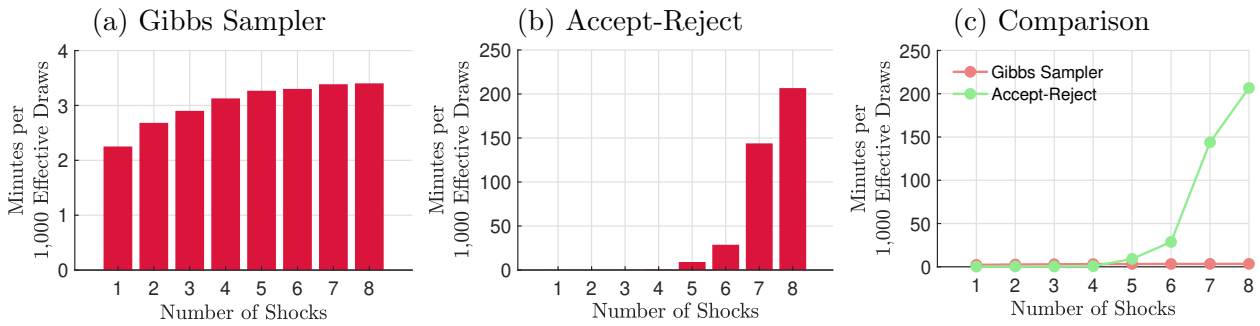

Figure 12: Time Per 1,000 Effective Draws

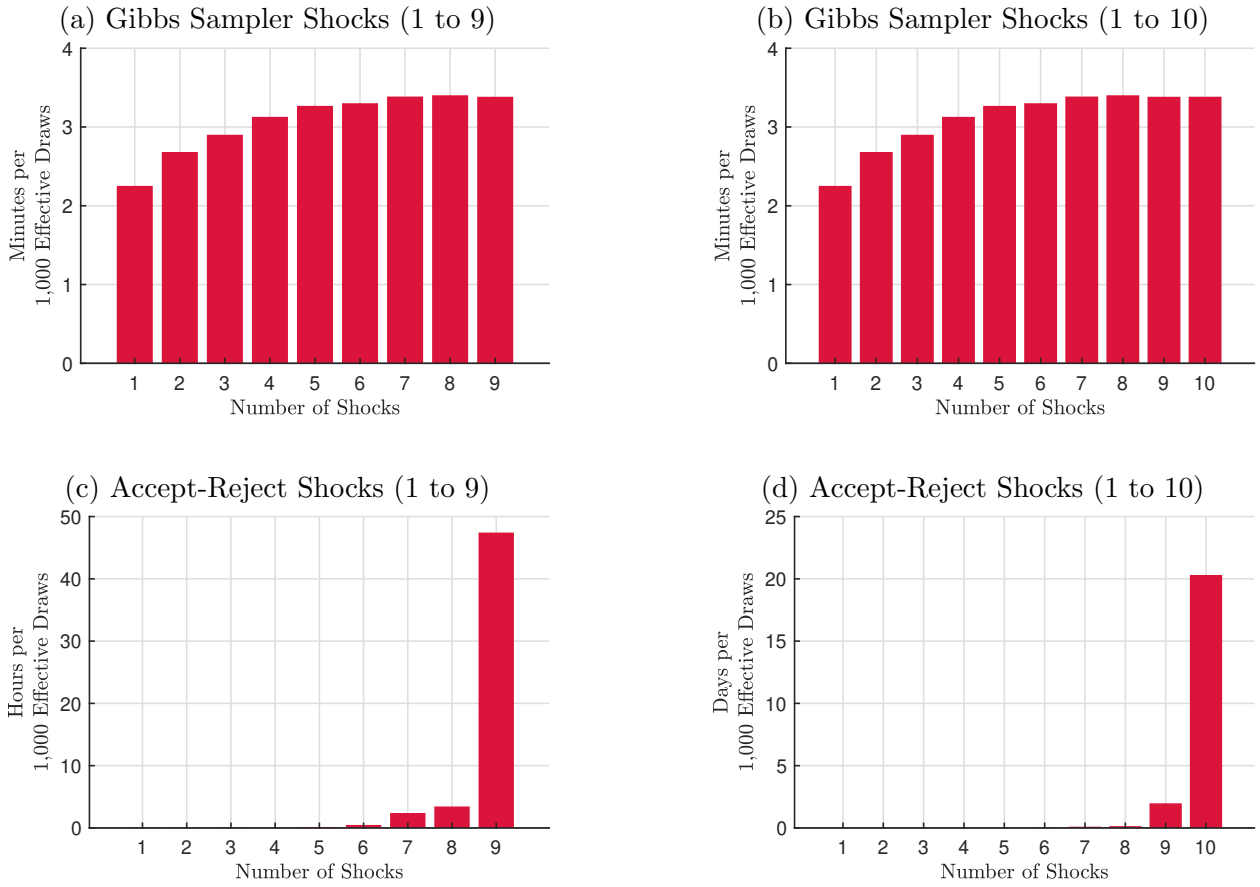

Figure 13: Gibbs Sampler vs. Accept-Reject

Note: The time of the accept-reject algorithm for shocks 9 and 10 is extrapolated based on 10 draws.

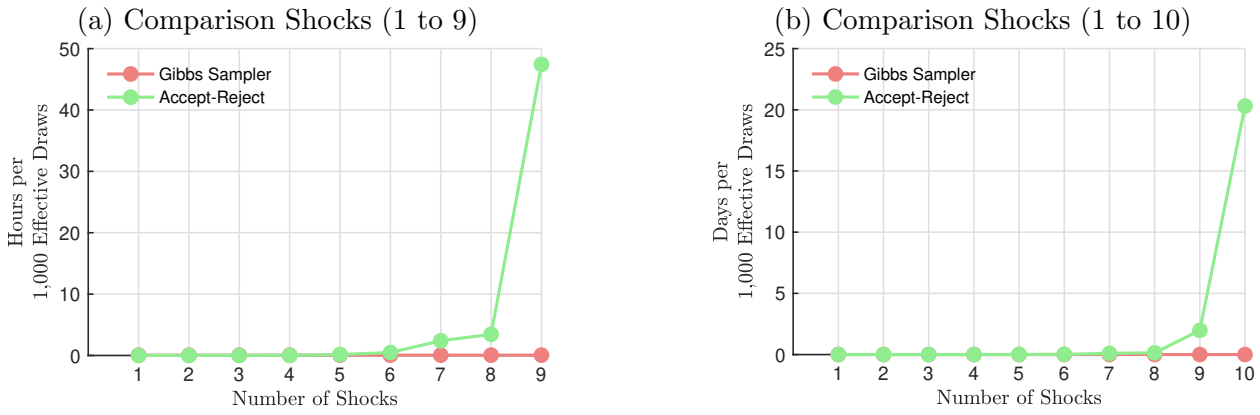

Figure 14: Gibbs Sampler vs. Accept-Reject

Note: The time of the accept-reject algorithm for shocks 9 and 10 is extrapolated based on 10 draws.