

Comparing Stochastic Volatility Specifications for Large Bayesian VARs

Joshua C.C. Chan*

Purdue University

November 2022

Abstract

Large Bayesian vector autoregressions with various forms of stochastic volatility have become increasingly popular in empirical macroeconomics. One main difficulty for practitioners is to choose the most suitable stochastic volatility specification for their particular application. We develop Bayesian model comparison methods—based on marginal likelihood estimators that combine conditional Monte Carlo and adaptive importance sampling—to choose among a variety of stochastic volatility specifications. The proposed methods can also be used to select an appropriate shrinkage prior on the VAR coefficients, which is a critical component for avoiding over-fitting in high-dimensional settings. Using US quarterly data of different dimensions, we find that both the Cholesky stochastic volatility and factor stochastic volatility outperform the common stochastic volatility specification. Their superior performance, however, can mostly be attributed to the more flexible priors that accommodate cross-variable shrinkage.

Keywords: large vector autoregression, marginal likelihood, Bayesian model comparison, stochastic volatility, shrinkage prior

*I would like to thank Roberto Casarin, Todd Clark, William McCausland, Editor Serena Ng, and an anonymous Associate Editor and reviewer for their many insightful comments and constructive suggestions. This paper has also benefited from helpful discussions with conference and seminar participants at the 11th European Central Bank Conference on Forecasting Techniques, Ca' Foscari University of Venice, the University of Montreal and the Federal Reserve Bank of Kansas City. All remaining errors are, of course, my own.

1 Introduction

Large Bayesian vector autoregressions (VARs) are now widely used for empirical macroeconomic analysis and forecasting thanks to the seminal work of Banbura, Giannone, and Reichlin (2010).¹ Since it is well established that time-varying volatility is vitally important for small VARs, it is expected to be even more so for large systems.² Consequently, there has been a lot of recent research devoted to designing stochastic volatility specifications suitable for large systems. Prominent examples include the common stochastic volatility models (Carriero, Clark, and Marcellino, 2016; Chan, 2020), the Cholesky stochastic volatility models (Cogley and Sargent, 2005; Carriero, Clark, and Marcellino, 2019) and the factor stochastic volatility models (Pitt and Shephard, 1999b; Chib, Nardari, and Shephard, 2006; Kastner, 2019). Since these stochastic volatility models are widely different and the choice among these alternatives involves important trade-offs—e.g., flexibility versus speed of estimation—one main issue facing practitioners is the lack of tools to compare these high-dimensional, non-linear and non-nested models.

Of course, the natural Bayesian model comparison criterion is the marginal likelihood (or the marginal data density), and in principle it can be used to select among these stochastic volatility models. In practice, however, computing the marginal likelihood for high-dimensional VARs with stochastic volatility is hardly trivial due to the large number of VAR coefficients and the latent state variables (e.g., stochastic volatility and latent factors). We tackle this obstacle by developing new methods to estimate the marginal likelihood of large Bayesian VARs with a variety of stochastic volatility specifications.

More specifically, we combine two popular variance reduction techniques, namely, conditional Monte Carlo and adaptive importance sampling, to construct our marginal likelihood estimators. We first analytically integrate out the large number of VAR

¹The literature on using large Bayesian VARs for structural analysis and forecasting is rapidly expanding. Early applications include Carriero, Kapetanios, and Marcellino (2009), Koop (2013), Koop and Korobilis (2013), Banbura, Giannone, Modugno, and Reichlin (2013) and Carriero, Clark, and Marcellino (2015).

²For small VARs with only a few variables, papers such as Clark (2011), D’Agostino, Gambetti, and Giannone (2013), Clark and Ravazzolo (2014), Cross and Poon (2016) and Chan and Eisenstat (2018) highlight the importance of time-varying volatility for model-fit and forecasting. For large VARs with more than a dozen variables, Carriero, Clark, and Marcellino (2016) present evidence that the data overwhelmingly favors a VAR with stochastic volatility; Koop and Korobilis (2013), Carriero, Clark, and Marcellino (2019) and Chan, Eisenstat, Hou, and Koop (2020) show that VARs with stochastic volatility forecast better than their homoskedastic counterparts, especially for density forecasts.

coefficients—i.e., we derive an analytical expression of the likelihood unconditional on the VAR coefficients that can be evaluated quickly. In the second step, we construct an adaptive importance sampling estimator—obtained by minimizing the Kullback-Leibler divergence to the ideal zero-variance importance sampling density—to integrate out the log-volatility via Monte Carlo. By carefully combining these two ways of integration (analytical and Monte Carlo integration), we are able to efficiently evaluate the marginal likelihood of a variety of popular stochastic volatility models for large Bayesian VARs.

Compared to earlier marginal likelihood estimators for Bayesian VARs with stochastic volatility, such as Chan and Eisenstat (2018) and Chan (2020), the new method offers two main advantages. First, it analytically integrates out the large number of VAR coefficients. As such, it reduces the variance of the estimator by eliminating the portion contributed by the VAR coefficients. This reduction is expected to be substantial in large VARs. Second, the new method is based on a *global* approximation of the joint distribution of the log-volatility instead of *local* approximations used in earlier marginal likelihood estimators. More specifically, earlier estimators are based on a second-order Taylor expansion of the log target density around the mode. This approximation is therefore entirely based on the properties at the point of expansion (i.e., the functional value, the gradient and the Hessian at the mode). It is guaranteed to approximate the target density well in the neighborhood of the point of expansion, but its accuracy typically deteriorates rapidly away from the approximation point. In contrast, the new approximation is constructed by solving a functional minimization problem: within a class of multivariate Gaussian distributions, it obtains the member that minimizes the Kullback-Leibler divergence to the target density. Hence, the minimizer is a global approximation in the sense that it incorporates information from the entire support of the target density.

In addition to comparing different stochastic volatility specifications, the proposed method can also be used to select an appropriate shrinkage prior on the VAR coefficients. Since even small VARs have a large number of parameters, shrinkage priors are essential to avoid over-fitting in high-dimensional settings. The most prominent example of these shrinkage priors is the Minnesota prior first introduced by Doan, Litterman, and Sims (1984) and Litterman (1986), not long after the seminal work on VARs by Sims (1980). There are now a wide range of more flexible variants (see, e.g., Kadiyala and Karlsson, 1993, 1997; Giannone, Lenza, and Primiceri, 2015; Chan, 2021), and choosing among them for a particular application has become a practical issue for empirical economists.

In particular, we focus on Minnesota priors with two potentially useful features: 1) allowing the overall shrinkage hyperparameters to be estimated from the data rather than fixing them at some subjectively elicited values; and 2) cross-variable shrinkage, i.e., the idea that coefficients on ‘other’ lags should be shrunk to zero more aggressively than those on ‘own’ lags. The proposed marginal likelihood estimators can provide a way to compare shrinkage priors with and without these features.

Through a series of Monte Carlo experiments, we demonstrate that the proposed estimators work well in practice. In particular, we show that one can correctly distinguish the three different stochastic volatility specifications: the common stochastic volatility (VAR-CSV), the Cholesky stochastic volatility (VAR-SV) and the factor stochastic volatility (VAR-FSV). In addition, the proposed marginal likelihood estimators can also be used to identify the correct number of factors in factor stochastic volatility models.

In an empirical application using US quarterly data, we compare the three stochastic volatility specifications in fitting datasets of different dimensions (7, 15 and 30 variables). The model comparison results show that the data overwhelmingly prefers VAR-SV and VAR-FSV over the more restrictive VAR-CSV for all model dimensions. We also find strong evidence in favor of the two aforementioned features of some Minnesota priors: both cross-variable shrinkage and a data-based approach to determine the overall shrinkage strength are empirically important. In fact, when we turn off the cross-variable shrinkage in the priors of VAR-SV and VAR-FSV, they perform similarly as VAR-CSV, suggesting that the superior performance of the former two models can mostly be attributed to the more flexible priors that accommodate cross-variable shrinkage. These results thus illustrate that in high-dimensional settings, choosing a flexible shrinkage prior is as important as selecting a flexible stochastic volatility specification. Lastly, we also demonstrate how the proposed methodology can be used to improve other marginal likelihood estimators and model selection criteria such as the deviance information criterion.

The rest of the paper is organized as follows. We first outline in Section 2 the three stochastic volatility specifications designed for large BayBunesian VARs, followed by an overview of various data-driven Minnesota priors. Then, Section 3 describes the two components of the proposed marginal likelihood estimators: adaptive importance sampling and conditional Monte Carlo. The methodology is illustrated via a concrete example of estimating the marginal likelihood of the common stochastic volatility model.

We then conduct a series of Monte Carlo experiments in Section 4 to assess how the proposed marginal likelihood estimators perform in selecting the correct data generating process. It is followed by an empirical application in Section 5, where we compare the three stochastic volatility specifications in the context of Bayesian VARs of different dimensions. Section 6 discusses various extensions of using the proposed methods to improve other marginal likelihood estimators and alternative model comparison methods. There we also consider a stochastic volatility model with an outlier component. Lastly, Section 7 concludes and briefly discusses some future research directions.

2 Stochastic Volatility Models for Large VARs

In this section we first describe a few recently proposed stochastic volatility specifications designed for large Bayesian VARs. We then outline a few data-driven Minnesota priors that are particularly useful for these high-dimensional models.

2.1 Common Stochastic Volatility

Let $\mathbf{y}_t = (y_{1,t}, \dots, y_{n,t})'$ be an $n \times 1$ vector of variables that is observed over the periods $t = 1, \dots, T$. The first specification is the common stochastic volatility model introduced in Carriero, Clark, and Marcellino (2016). The conditional mean equation is the standard reduced-form VAR with p lags:

$$\mathbf{y}_t = \mathbf{a}_0 + \mathbf{A}_1 \mathbf{y}_{t-1} + \dots + \mathbf{A}_p \mathbf{y}_{t-p} + \boldsymbol{\varepsilon}_t, \quad (1)$$

where \mathbf{a}_0 is an $n \times 1$ vector of intercepts and $\mathbf{A}_1, \dots, \mathbf{A}_p$ are all $n \times n$ coefficient matrices. To allow for heteroskedastic errors, the covariance matrix of the innovation $\boldsymbol{\varepsilon}_t$ is scaled by a common, time-varying factor that can be interpreted as the overall macroeconomic volatility:

$$\boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, e^{h_t} \boldsymbol{\Sigma}). \quad (2)$$

The log-volatility h_t in turn follows a stationary AR(1) process:

$$h_t = \phi h_{t-1} + u_t^h, \quad u_t^h \sim \mathcal{N}(0, \sigma^2), \quad (3)$$

for $t = 2, \dots, T$, where $|\phi| < 1$ and the initial condition is specified as $h_1 \sim \mathcal{N}(0, \sigma^2/(1 - \phi^2))$. Note that the unconditional mean of the AR(1) process is assumed to be zero for identification. We refer to this common stochastic volatility model as VAR-CSV.

One drawback of this volatility specification is that it appears to be restrictive. For example, all variances are scaled by a single factor and, consequently, they are always proportional to each other. Nevertheless, there is empirical evidence that the error variances of macroeconomic variables tend to move closely together (see, e.g., Carriero, Clark, and Marcellino, 2016; Chan, Eisenstat, and Strachan, 2020), and a common stochastic volatility is a parsimonious way to model this empirical feature.

Estimating large VARs is in general computationally intensive because of the large number of VAR coefficients $\mathbf{A} = (\mathbf{a}_0, \mathbf{A}_1, \dots, \mathbf{A}_p)'$. One main advantage of the common stochastic volatility specification is that—if a natural conjugate prior on $(\mathbf{A}, \boldsymbol{\Sigma})$ is used—it leads to many useful analytical results that make estimation fast. In particular, there are efficient algorithms to generate the large number of VAR coefficients. Moreover, as demonstrated in Chan (2020), similar computational gains can be achieved for a much wider class of VARs with non-Gaussian, heteroskedastic and serially dependent innovations. Recent empirical applications using this common stochastic volatility and its variants include Mumtaz (2016), Mumtaz and Theodoridis (2017), Götz and Hauzenberger (2018), Poon (2018), Louzis (2019), LeSage and Hendrikz (2019), Zhang and Nguyen (2020) and Hartwig (2021).

2.2 Cholesky Stochastic Volatility

A more flexible way to model multivariate heteroscedasticity and time-varying covariances is to incorporate multiple stochastic volatility processes, as first proposed in Cogley and Sargent (2005). Specifically, consider the same reduced-form VAR in (1), but the innovation $\boldsymbol{\varepsilon}_t$ has a more flexible time-varying covariance matrix modeled as follows:

$$\boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_t), \quad \boldsymbol{\Sigma}_t^{-1} = \mathbf{B}'_0 \mathbf{D}_t^{-1} \mathbf{B}_0, \quad (4)$$

where \mathbf{B}_0 is an $n \times n$ lower triangular matrix with ones on the diagonal and $\mathbf{D}_t = \text{diag}(e^{h_{1,t}}, \dots, e^{h_{n,t}})$. The law of motion for each element of $\mathbf{h}_t = (h_{1,t}, \dots, h_{n,t})'$ is specified

as an independent autoregressive process:

$$h_{i,t} = \mu_i + \phi_i(h_{i,t-1} - \mu_i) + u_{i,t}^h, \quad u_{i,t}^h \sim \mathcal{N}(0, \sigma_i^2) \quad (5)$$

for $t = 2, \dots, T$, where the initial condition is specified as $h_{i,1} \sim \mathcal{N}(\mu_i, \sigma_i^2/(1 - \phi_i^2))$. This stochastic volatility specification is sometimes called the Cholesky stochastic volatility. We refer to this stochastic volatility model as VAR-SV.

This model can be estimated using the algorithm described in Del Negro and Primiceri (2015). However, when the number of variables is large, the conventional way of drawing all VAR coefficients jointly becomes excessively computationally intensive. To tackle this computational problem, Carriero, Clark, and Marcellino (2019) introduce a blocking scheme that makes it possible to estimate the reduced-form VAR equation by equation that drastically reduces the computational time.³

In contrast to the common stochastic volatility model, VAR-SV is more flexible in that it contains n stochastic volatility processes, which can accommodate more complex co-volatility patterns. But this comes at a cost of more intensive posterior computations: the complexity of estimating VAR-SV is $\mathcal{O}(n^4)$ compared to $\mathcal{O}(n^3)$ for VAR-CSV (when a natural conjugate prior is used). Recent empirical applications using this Cholesky stochastic volatility in the context of large Bayesian VARs include Banbura and van Vlodrop (2018), Bianchi, Guidolin, and Ravazzolo (2018), Huber and Feldkircher (2019), Cross, Hou, and Poon (2019), Baumeister, Korobilis, and Lee (2020), Koop, McIntyre, Mitchell, and Poon (2020), Tallman and Zaman (2020), Zens, Böck, and Zörner (2020) and Chan (2021).

2.3 Factor Stochastic Volatility

The third stochastic volatility specification that is suitable for large systems belongs to the class of factor stochastic volatility models (Pitt and Shephard, 1999a; Chib, Nardari, and Shephard, 2006; Kastner, 2019). More specifically, consider the same reduced-form

³As pointed out in Bognanni (2022), the algorithm in Carriero, Clark, and Marcellino (2019) to sample the VAR coefficients equation by equation can only be viewed as an approximation, as it ignores certain integrating constants of the distributions of VAR coefficients in other equations. Carriero, Chan, Clark, and Marcellino (2022) provide an alternative algorithm to draw the VAR coefficients equation by equation that has the same computational complexity of $\mathcal{O}(n^4)$.

VAR in (1), but the innovation is instead decomposed as:

$$\boldsymbol{\varepsilon}_t = \mathbf{L}\mathbf{f}_t + \mathbf{u}_t,$$

where $\mathbf{f}_t = (f_{1,t}, \dots, f_{r,t})'$ is a $r \times 1$ vector of latent factors and \mathbf{L} is the associated $n \times r$ factor loading matrix. For identification purposes we assume that \mathbf{L} is a lower triangular matrix with ones on the main diagonal.⁴ Furthermore, to ensure one can separately identify the common and the idiosyncratic components, we adopt a sufficient condition in Anderson and Rubin (1956) that requires $r \leq (n - 1)/2$.⁵

The disturbances \mathbf{u}_t and the latent factors \mathbf{f}_t are assumed to be independent at all leads and lags. Furthermore, they are jointly Gaussian:

$$\begin{pmatrix} \mathbf{u}_t \\ \mathbf{f}_t \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \mathbf{0} \\ \mathbf{0} \end{pmatrix}, \begin{pmatrix} \mathbf{D}_t & \mathbf{0} \\ \mathbf{0} & \mathbf{G}_t \end{pmatrix} \right), \quad (6)$$

where $\mathbf{D}_t = \text{diag}(e^{h_{1,t}}, \dots, e^{h_{n,t}})$ and $\mathbf{G}_t = \text{diag}(e^{h_{n+1,t}}, \dots, e^{h_{n+r,t}})$ are diagonal matrices. Here the correlations among the elements of the innovation $\boldsymbol{\varepsilon}_t$ are induced by the latent factors. In typical applications, a small number of factors would be sufficient to capture the time-varying covariance structure even when n is large.

Next, for each $i = 1, \dots, n + r$, the evolution of the log-volatility is modeled as:

$$h_{i,t} = \mu_i + \phi_i(h_{i,t-1} - \mu_i) + u_{i,t}^h, \quad u_{i,t}^h \sim \mathcal{N}(0, \sigma_i^2) \quad (7)$$

for $t = 2, \dots, T$. The initial state is assumed to follow the stationary distribution: $h_{i,1} \sim \mathcal{N}(\mu_i, \sigma_i^2/(1 - \phi_i^2))$. We refer to this factor stochastic volatility model as VAR-FSV.

⁴One main reason to impose these identification restrictions on \mathbf{L} is to avoid the so-called label switching problem, which makes it a bit more difficult to construct appropriate importance sampling densities for the log-volatility. One way to handle this issue is to first postprocess the posterior draws to sort them into the correct categories using, e.g., the approach in Kaufmann and Schumacher (2019). Then, one can use these sorted MCMC draws to construct the importance sampling densities discussed below.

⁵In this paper we follow the common approach of fixing the number of factors r in the estimation and selecting r via the marginal likelihood. An alternative approach is to allow the inclusion of potentially infinitely many factors, but to use a shrinkage prior on the factor loadings so that they are increasingly shrunk to 0 as the column index increases (Bhattacharya and Dunson, 2011; Kastner, 2019). This latter approach circumvents the need to select the number of factors, but it makes the interpretation of the factors more difficult.

In a sense VAR-FSV is even more flexible than VAR-SV as the former contains $n + r$ stochastic volatility processes compared to n in the latter. In terms of estimation, both are more computationally intensive to estimate relative to VAR-CSV, but they can still be fitted reasonably quickly even when n is large. In particular, given the latent factors VAR-FSV becomes n unrelated regressions. Consequently, the n -equation system can be estimated equation by equation. While various factor stochastic volatility specifications are widely applied in financial applications (see, e.g., Aguilar and West, 2000; Jin, Maheu, and Yang, 2019; Hautsch and Voigt, 2019; Li and Scharth, 2020; McCausland, Miller, and Pelletier, 2020), they are not yet commonly used in conjunction with a large VAR (with the notable exception of Kastner and Huber, 2020).

2.4 Hierarchical Minnesota Priors

Next, we briefly describe the prior specifications for the above three stochastic volatility models; the details of the priors are given in Appendix A. In general, we assume exactly the same priors on the common parameters across models. When this is not possible, we use similar priors so that they are comparable across models. Below we focus on the shrinkage priors on the VAR coefficients as they are critical in the current context.

More specifically, we consider Minnesota-type priors on the VAR coefficients. This family of priors was first developed by Doan, Litterman, and Sims (1984) and Litterman (1986), and a number of more flexible variants have been developed since. We adopt a version in which the overall shrinkage hyperparameters are treated as parameters to be estimated from the data, in the spirit of Giannone, Lenza, and Primiceri (2015). Compared to the conventional version in which the hyperparameters are fixed at some subjective values, this data-based variant is often found to fit the data substantially better and have better out-of-sample forecast performance (see, e.g., Carriero, Clark, and Marcellino, 2015; Chan, Jacobi, and Zhu, 2020).

First, for the VAR-CSV we consider the natural conjugate prior on $(\mathbf{A}, \mathbf{\Sigma})$ that depends on a shrinkage hyperparameter κ :

$$\mathbf{\Sigma} \sim \mathcal{IW}(\nu_0, \mathbf{S}_0), \quad (\text{vec}(\mathbf{A}) \mid \mathbf{\Sigma}, \kappa) \sim \mathcal{N}(\text{vec}(\mathbf{A}_0), \mathbf{\Sigma} \otimes \mathbf{V}_{\mathbf{A}}),$$

where the prior hyperparameters $\text{vec}(\mathbf{A}_0)$ and $\mathbf{V}_{\mathbf{A}}$ are chosen in the spirit of the Minnesota

prior. More specifically, since in the empirical applications we will be using data in growth rates, we set $\text{vec}(\mathbf{A}_0) = \mathbf{0}$ to shrink all the coefficients to 0. For the covariance matrix $\mathbf{V}_\mathbf{A}$, it is assumed to be diagonal and it depends on a single hyperparameter κ that controls the overall shrinkage strength (see Appendix A for details). As mentioned above, we treat κ as an unknown parameter to be estimated. This natural conjugate prior substantially speeds up posterior simulations due to the availability of various analytical results. One major drawback of this natural conjugate prior, however, is that it cannot accommodate cross-variable shrinkage—i.e., the prior belief that coefficients on ‘own’ lags are on average larger than ‘other’ lags.

Next, for the VAR-SV, we assume that the intercept and the VAR coefficients are *a priori* independent across equations, and each vector $\boldsymbol{\alpha}_i = (a_{i,0}, \mathbf{A}_{i,1}, \dots, \mathbf{A}_{i,p})'$, for $i = 1, \dots, n$, where $\mathbf{A}_{i,j}$ denotes the i -th row of \mathbf{A}_j , has a hierarchical normal prior: $\boldsymbol{\alpha}_i \sim \mathcal{N}(\boldsymbol{\alpha}_{0,i}, \mathbf{V}_{\boldsymbol{\alpha}_i})$. Here the prior mean $\boldsymbol{\alpha}_{0,i}$ is set to be zero to shrink the VAR coefficients to zero. The prior covariance matrix $\mathbf{V}_{\boldsymbol{\alpha}_i}$ is assumed to be diagonal and it depends on two hyperparameters: κ_1 and κ_2 . The hyperparameter κ_1 controls the overall shrinkage strength for VAR coefficients on own lags, whereas κ_2 controls those on lags of other variables. Allowing the two shrinkage parameters to be different across the two groups of VAR coefficients is empirically important. For example, Carriero, Clark, and Marcellino (2015) and Chan (2021) find empirical evidence in support of cross-variable shrinkage. In addition, following recent works such as Giannone, Lenza, and Primiceri (2015) and Amir-Ahmadi, Matthes, and Wang (2020) that demonstrate the benefits of tuning these hyperparameters in a data-based manner relative to fixing them at specific values, we treat both κ_1 and κ_2 to be parameters to be estimated.

Let $\boldsymbol{\beta}_i$ denote the free elements in the i -th row of the impact matrix \mathbf{B}_0 for $i = 2, \dots, n$. That is, $\boldsymbol{\beta}_i = (B_{0,i,1}, \dots, B_{0,i,i-1})'$, where $B_{0,i,j}$ is the (i, j) element of \mathbf{B}_0 . We assume that $\boldsymbol{\beta}_i$ has a hierarchical normal prior: $\boldsymbol{\beta}_i \sim \mathcal{N}(\boldsymbol{\beta}_{0,i}, \mathbf{V}_{\boldsymbol{\beta}_i})$. We set the the prior mean $\boldsymbol{\beta}_{0,i}$ to be zero so that the free elements in the impact matrix are shrunk to zero. The prior covariance matrix $\mathbf{V}_{\boldsymbol{\beta}_i}$ is set to be diagonal, and it depends on a hyperparameter κ_3 , which controls the overall shrinkage strength. We again treat κ_3 as an unknown parameter to be estimated. Allowing this shrinkage parameter to be determined by the data is found to substantially improve forecasts (see, e.g., Chan, 2021).

Finally, for the VAR-FSV we assume a similar normal prior on $\boldsymbol{\alpha}_i = (a_{i,0}, \mathbf{A}_{i,1}, \dots, \mathbf{A}_{i,p})'$,

the intercept and the VAR coefficients in the i -th equation: $\boldsymbol{\alpha}_i \sim \mathcal{N}(\boldsymbol{\alpha}_{0,i}, \mathbf{V}_{\boldsymbol{\alpha}_i})$. Again we set $\boldsymbol{\alpha}_{0,i} = \mathbf{0}$ and assume $\mathbf{V}_{\boldsymbol{\alpha}_i}$ to be diagonal and depends on two hyperparameters: one controls the overall shrinkage strength on coefficients on own lags and the other on coefficients on other lags. Details of these Minnesota priors and priors on other parameters are given in Appendix A.

3 Marginal Likelihood Estimation: Adaptive Importance Sampling and Conditional Monte Carlo

In this section we first give an overview of the marginal likelihood and its role in Bayesian model comparison. We then describe the two components of the proposed marginal likelihood estimators: adaptive importance sampling and conditional Monte Carlo. Specifically, conditional Monte Carlo is used to integrate out the large number of VAR coefficients analytically, and adaptive importance sampling is used to bias the distributions of the remaining parameters to oversample the ‘important’ regions of the parameter space that contribute most to the integral (and this bias is later corrected in the estimation). Finally, to illustrate the methodology, we provide the details of estimating the marginal likelihood of the VAR-CSV by combining adaptive importance sampling and conditional Monte Carlo.

3.1 Overview of the Marginal Likelihood

Suppose we want to compare a collection of models $\{M_1, \dots, M_K\}$, where each model M_k is formally defined by a likelihood function $p(\mathbf{y} | \boldsymbol{\theta}_k, M_k)$ and a prior on the model-specific parameter vector $\boldsymbol{\theta}_k$ denoted by $p(\boldsymbol{\theta}_k | M_k)$. A natural Bayesian model comparison criterion is the Bayes factor in favor of M_i against M_j , defined as

$$\text{BF}_{ij} = \frac{p(\mathbf{y} | M_i)}{p(\mathbf{y} | M_j)},$$

where

$$p(\mathbf{y} | M_k) = \int p(\mathbf{y} | \boldsymbol{\theta}_k, M_k) p(\boldsymbol{\theta}_k | M_k) d\boldsymbol{\theta}_k$$

is the *marginal likelihood* under model M_k , $k = i, j$. One can show that if the data is generated from model M_i , the Bayes factor would on average pick the correct model over some distinct model, say, M_j . To see that, we compute the expected log Bayes factor in favor of model M_i with respect to the distribution $p(\mathbf{y} | M_i)$:

$$\mathbb{E}_{\mathbf{y}} \left[\log \frac{p(\mathbf{y} | M_i)}{p(\mathbf{y} | M_j)} \right] = \int p(\mathbf{y} | M_i) \log \frac{p(\mathbf{y} | M_i)}{p(\mathbf{y} | M_j)} d\mathbf{y}.$$

This expression is the Kullback-Leibler divergence from $p(\mathbf{y} | M_i)$ to $p(\mathbf{y} | M_j)$, and it is strictly positive unless $p(\mathbf{y} | M_i) = p(\mathbf{y} | M_j)$, in which case it is zero. In other words, the marginal likelihood of M_i is larger than that of M_j on average.

Furthermore, the Bayes factor is related to the posterior odds ratio between the two models as follows:

$$\frac{\mathbb{P}(M_i | \mathbf{y})}{\mathbb{P}(M_j | \mathbf{y})} = \frac{\mathbb{P}(M_i)}{\mathbb{P}(M_j)} \times \text{BF}_{ij},$$

where $\mathbb{P}(M_i)/\mathbb{P}(M_j)$ is the prior odds ratio. It follows that if both models are equally probable *a priori*, i.e., $p(M_i) = p(M_j)$, the posterior odds ratio between the two models is then equal to the Bayes factor or ratio of marginal likelihoods. In that case, if, say, $\text{BF}_{ij} = 100$, then model M_i is 100 times more likely than model M_j given the data. For a more detailed discussion of the Bayes factor and its role in Bayesian model comparison, see Koop (2003) or Chan, Koop, Poirier, and Tobias (2019). From here onwards we suppress the model indicator; for example we denote the marginal likelihood and the likelihood by $p(\mathbf{y})$ and $p(\mathbf{y} | \boldsymbol{\theta})$, respectively.

3.2 Adaptive Importance Sampling

Next, we give an overview of an adaptive importance sampling approach called the improved cross-entropy method for estimating the marginal likelihood. It is based on the idea of biasing the sampling distribution in such a way that more ‘important values’ are generated in the simulation. The sample is then weighted to correct for the use of a different distribution to give an unbiased estimator. In our context of estimating the

marginal likelihood $p(\mathbf{y})$, consider the following importance sampling estimator:

$$\widehat{p(\mathbf{y})}_{\text{IS}} = \frac{1}{R} \sum_{r=1}^R \frac{p(\mathbf{y} | \boldsymbol{\theta}^{(r)})p(\boldsymbol{\theta}^{(r)})}{g(\boldsymbol{\theta}^{(r)})}, \quad (8)$$

where $\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(R)}$ are independent draws obtained from the importance sampling density $g(\cdot)$ that dominates $p(\mathbf{y} | \cdot)p(\cdot)$ —i.e., $g(\boldsymbol{\theta}) = 0 \Rightarrow p(\mathbf{y} | \boldsymbol{\theta})p(\boldsymbol{\theta}) = 0$.

The importance sampling estimator in (8) is unbiased and simulation consistent for any g that dominates $p(\mathbf{y} | \cdot)p(\cdot)$. But its performance depends critically on the choice of g . Below we follow Chan and Eisenstat (2015) to use an improved version of the classic cross-entropy method to construct g optimally. The original cross-entropy method was developed for rare-event simulation by Rubinstein (1997, 1999) using a multi-level procedure to obtain the optimal importance sampling density (see also Rubinstein and Kroese, 2004, for a book-length treatment). Chan and Eisenstat (2015) show that this optimal importance sampling density can be obtained more accurately in one step using MCMC methods. This adaptive importance sampling estimator is then used to estimate the marginal likelihood.

To motivate the cross-entropy method, first note that there exists an importance sampling density that gives a zero-variance estimator of the marginal likelihood. In particular, it is easy to verify that if we use the posterior distribution as the importance sampling density, i.e., $g^*(\boldsymbol{\theta}) = p(\boldsymbol{\theta} | \mathbf{y}) = p(\mathbf{y} | \boldsymbol{\theta})p(\boldsymbol{\theta})/p(\mathbf{y})$, then the associated importance sampling estimator (8) has zero variance. Of course in practice g^* cannot be used as its normalizing constant is exactly the marginal likelihood, the unknown quantity we wish to estimate. Nevertheless, this provides a signpost to obtain an optimal importance sampling density. Intuitively, if we choose an importance sampling density g that is ‘close enough’ to g^* so that both behave similarly, the resulting importance sampling estimator should have reasonable accuracy. Hence, our goal is to locate a convenient density that is in a well-defined sense ‘close’ to g^* .

To that end, consider a parametric family $\mathcal{F} = \{f(\boldsymbol{\theta}; \mathbf{v})\}$ indexed by the parameter vector \mathbf{v} within which we locate the optimal importance sampling density. We find the density $f(\boldsymbol{\theta}; \mathbf{v}^*) \in \mathcal{F}$ such that it is the ‘closest’ to g^* . One convenient measure of closeness between densities is the *Kullback-Leibler divergence* or the *cross-entropy distance*. Specifically, let f_1 and f_2 be two probability density functions. Then, the cross-entropy

distance from f_1 to f_2 is defined as:

$$\mathcal{D}(f_1, f_2) = \int f_1(\mathbf{x}) \log \frac{f_1(\mathbf{x})}{f_2(\mathbf{x})} d\mathbf{x}.$$

Given this measure, we can then locate the density $f(\cdot; \mathbf{v}) \in \mathcal{F}$ such that $\mathcal{D}(g^*, f(\cdot; \mathbf{v}))$ is minimized: $\mathbf{v}_{\text{ce}}^* = \operatorname{argmin}_{\mathbf{v}} \mathcal{D}(g^*, f(\cdot; \mathbf{v}))$. It can be shown that solving the CE minimization problem is equivalent to finding

$$\mathbf{v}_{\text{ce}}^* = \operatorname{argmax}_{\mathbf{v}} \int p(\mathbf{y} | \boldsymbol{\theta}) p(\boldsymbol{\theta}) \log f(\boldsymbol{\theta}; \mathbf{v}) d\boldsymbol{\theta}.$$

This optimization problem is often difficult to solve analytically. Instead, we consider its stochastic counterpart:

$$\widehat{\mathbf{v}}_{\text{ce}}^* = \operatorname{argmax}_{\mathbf{v}} \frac{1}{M} \sum_{m=1}^M \log f(\boldsymbol{\theta}_m; \mathbf{v}), \quad (9)$$

where $\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(M)}$ are draws from the posterior distribution $p(\boldsymbol{\theta} | \mathbf{y}) \propto p(\mathbf{y} | \boldsymbol{\theta}) p(\boldsymbol{\theta})$.

In other words, $\widehat{\mathbf{v}}_{\text{ce}}^*$ is exactly the maximum likelihood estimate for \mathbf{v} if we treat $f(\boldsymbol{\theta}; \mathbf{v})$ as the likelihood function with parameter vector \mathbf{v} and $\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(M)}$ an observed sample. Since finding the maximum likelihood estimator is a standard problem, solving (9) is typically easy. In particular, analytical solutions to (9) can be found explicitly for the exponential family (e.g., Rubinstein and Kroese, 2004, p. 70).

The parametric family \mathcal{F} is often chosen so that each member $f(\boldsymbol{\theta}; \mathbf{v})$ is a product of densities, e.g., $f(\boldsymbol{\theta}; \mathbf{v}) = f(\boldsymbol{\theta}_1; \mathbf{v}_1) \times \dots \times f(\boldsymbol{\theta}_B; \mathbf{v}_B)$, where $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_B)$ and $\mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_B)$. In that case, one can reduce the possibly high-dimensional maximization problem (9) into B low-dimensional problems, which can then be readily solved. Once the optimal density is obtained, we then set $g(\cdot) = f(\cdot; \widehat{\mathbf{v}}_{\text{ce}}^*)$ and use it to construct the importance sampling estimator in (8).

3.3 Conditional Monte Carlo

Recall that the marginal likelihood can be written as $p(\mathbf{y}) = \mathbb{E}p(\mathbf{y} | \boldsymbol{\theta})$, where \mathbf{y} is fixed and the expectation is taken with respect to the prior distribution $p(\boldsymbol{\theta})$. Suppose that there is a random vector $\boldsymbol{\psi}$ with density function $p(\boldsymbol{\psi})$ such that the conditional expectation $\mathbb{E}[p(\mathbf{y} | \boldsymbol{\theta}) | \boldsymbol{\psi}]$ can be computed analytically. By the law of iterated expectation, we have

$$p(\mathbf{y}) = \mathbb{E}[\mathbb{E}[p(\mathbf{y} | \boldsymbol{\theta}) | \boldsymbol{\psi}]],$$

where the outer expectation is taken with respect to $p(\boldsymbol{\psi})$ and the inner expectation is taken with respect to the conditional distribution $p(\boldsymbol{\theta} | \boldsymbol{\psi})$. Hence, $\mathbb{E}[p(\mathbf{y} | \boldsymbol{\theta}) | \boldsymbol{\psi}]$ is an unbiased estimator of $p(\mathbf{y})$. In addition, by the variance decomposition formula, we obtain

$$\text{Var}(p(\mathbf{y} | \boldsymbol{\theta})) = \mathbb{E}[\text{Var}(p(\mathbf{y} | \boldsymbol{\theta}) | \boldsymbol{\psi})] + \text{Var}(\mathbb{E}[p(\mathbf{y} | \boldsymbol{\theta}) | \boldsymbol{\psi}]).$$

Since $\mathbb{E}[\text{Var}(p(\mathbf{y} | \boldsymbol{\theta}) | \boldsymbol{\psi})] > 0$ if the density $p(\boldsymbol{\theta} | \boldsymbol{\psi})$ is non-degenerate, we have

$$\text{Var}(\mathbb{E}[p(\mathbf{y} | \boldsymbol{\theta}) | \boldsymbol{\psi}]) < \text{Var}(p(\mathbf{y} | \boldsymbol{\theta})).$$

Therefore, the conditional Monte Carlo estimator $\mathbb{E}[p(\mathbf{y} | \boldsymbol{\theta}) | \boldsymbol{\psi}]$ *always* provides variance reduction. The degree of variance reduction, of course, depends on the dimension and variability of $\boldsymbol{\psi}$. Loosely speaking, to maximize variance reduction, one feasible approach is to find the ‘smallest’ subset $\boldsymbol{\psi}$ of $\boldsymbol{\theta}$ such that $\mathbb{E}[p(\mathbf{y} | \boldsymbol{\theta}) | \boldsymbol{\psi}]$ is computable.

In our VARs the parameter vector $\boldsymbol{\theta}$ contains the VAR coefficients and other latent variables such as the stochastic volatility and dynamic factors. To have the ‘smallest’ $\boldsymbol{\psi}$ is equivalent to integrating out as many elements in $\boldsymbol{\theta}$ as possible. In particular, since there are a large number of VAR coefficients in our setting, ideally we would like to remove their contribution to the variability in the marginal likelihood estimation. Fortunately, for all the VARs we consider, we are able to integrate them out analytically and obtain an analytical expression of $\mathbb{E}[p(\mathbf{y} | \boldsymbol{\theta}) | \boldsymbol{\psi}]$. We will give an explicit example in the following section.

3.4 An Example: Marginal Likelihood of VAR-CSV

As an illustration, we now provide the details of estimating the marginal likelihood of the VAR-CSV using the proposed method by combining adaptive importance sampling and conditional Monte Carlo. Other VARs with stochastic volatility can be handled similarly and the technical details are given in Appendix A.

Let $\mathbf{x}'_t = (1, \mathbf{y}'_{t-1}, \dots, \mathbf{y}'_{t-p})$ be a $1 \times k$ vector of an intercept and lags with $k = 1 + np$. Then, stacking the observations over $t = 1, \dots, T$, we can rewrite (1)–(2) as:

$$\mathbf{Y} = \mathbf{X}\mathbf{A} + \boldsymbol{\varepsilon}, \quad \text{vec}(\boldsymbol{\varepsilon}) \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma} \otimes \mathbf{D}_{\mathbf{h}}),$$

where $\mathbf{D}_{\mathbf{h}} = \text{diag}(e^{h_1}, \dots, e^{h_T})$, $\mathbf{A} = (\mathbf{a}_0, \mathbf{A}_1, \dots, \mathbf{A}_p)'$ is $k \times n$, and the matrices \mathbf{Y} , \mathbf{X} and $\boldsymbol{\varepsilon}$ are respectively of dimensions $T \times n$, $T \times k$ and $T \times n$. In addition, the natural conjugate prior on $(\mathbf{A}, \boldsymbol{\Sigma})$ depends on a single hyperparameter κ that controls the overall shrinkage strength, which is treated as an unknown parameter. Finally, the log-volatility h_t follows the AR(1) process in (3) with AR coefficient ϕ and error variance σ^2 .

Next, we define an appropriate conditional Monte Carlo estimator, which will then be combined with adaptive importance sampling. Using the notations in Section 3.3, the set of model parameters and latent variables is $\boldsymbol{\theta} = \{\mathbf{A}, \boldsymbol{\Sigma}, \phi, \sigma^2, \kappa, \mathbf{h}\}$, and the conditional likelihood (i.e., conditional on the latent variables \mathbf{h}) is given by $p(\mathbf{y} | \boldsymbol{\theta}) = p(\mathbf{Y} | \mathbf{h}, \mathbf{A}, \boldsymbol{\Sigma}, \phi, \sigma^2, \kappa) = p(\mathbf{Y} | \mathbf{h}, \mathbf{A}, \boldsymbol{\Sigma})$. When n is large, the dimensions of both \mathbf{A} and $\boldsymbol{\Sigma}$ are large, and we expect them to contribute the most to the variability of $p(\mathbf{Y} | \mathbf{h}, \mathbf{A}, \boldsymbol{\Sigma})$. Hence, we aim to integrate them out. To that end, we let $\boldsymbol{\psi} = \{\mathbf{h}, \phi, \kappa\}$ and consider the conditional expectation $\mathbb{E}[p(\mathbf{y} | \boldsymbol{\theta}) | \boldsymbol{\psi}] = p(\mathbf{Y} | \mathbf{h}, \kappa)$, which can be computed analytically:

$$\begin{aligned} p(\mathbf{Y} | \mathbf{h}, \kappa) &= \int p(\mathbf{Y} | \mathbf{h}, \mathbf{A}, \boldsymbol{\Sigma}) p(\mathbf{A}, \boldsymbol{\Sigma} | \kappa) d(\mathbf{A}, \boldsymbol{\Sigma}) \\ &= \pi^{-\frac{Tn}{2}} e^{-\frac{n}{2} \mathbf{1}'_T \mathbf{h}} |\mathbf{V}_{\mathbf{A}}|^{-\frac{n}{2}} |\mathbf{K}_{\mathbf{A}}|^{-\frac{n}{2}} \frac{\Gamma_n\left(\frac{\nu_0 + T}{2}\right) |\mathbf{S}_0|^{\frac{\nu_0}{2}}}{\Gamma_n\left(\frac{\nu_0}{2}\right) |\widehat{\mathbf{S}}|^{\frac{\nu_0 + T}{2}}}, \end{aligned} \quad (10)$$

where $\Gamma_n(\cdot)$ is the multivariate gamma function and

$$\begin{aligned} \mathbf{K}_{\mathbf{A}} &= \mathbf{V}_{\mathbf{A}}^{-1} + \mathbf{X}'\mathbf{D}_{\mathbf{h}}^{-1}\mathbf{X}, \quad \widehat{\mathbf{A}} = \mathbf{K}_{\mathbf{A}}^{-1}(\mathbf{V}_{\mathbf{A}}^{-1}\mathbf{A}_0 + \mathbf{X}'\mathbf{D}_{\mathbf{h}}^{-1}\mathbf{Y}), \\ \widehat{\mathbf{S}} &= \mathbf{S}_0 + \mathbf{A}'_0\mathbf{V}_{\mathbf{A}}^{-1}\mathbf{A}_0 + \mathbf{Y}'\mathbf{D}_{\mathbf{h}}^{-1}\mathbf{Y} - \widehat{\mathbf{A}}'\mathbf{K}_{\mathbf{A}}\widehat{\mathbf{A}}. \end{aligned}$$

Note that the overall shrinkage parameter κ appears in the prior covariance matrix $\mathbf{V}_\mathbf{A}$. The details of the derivations are given in Appendix A. We note that (10) can be calculated quickly without inverting any large matrices; see, e.g., Chan (2020) for various computational shortcuts.

Now, the marginal likelihood can be written in terms of $p(\mathbf{Y} | \mathbf{h}, \kappa)$ as follows:

$$p(\mathbf{y}) = \int p(\mathbf{Y} | \mathbf{h}, \kappa) p(\mathbf{h} | \phi, \sigma^2) p(\phi, \sigma^2, \kappa) d(\mathbf{h}, \phi, \sigma^2, \kappa),$$

where $p(\mathbf{h} | \phi, \sigma^2)$ is a Gaussian density implied by the state equation (3) and $p(\phi, \sigma^2, \kappa)$ is the prior for (ϕ, σ^2, κ) . Hence, the corresponding conditional Monte Carlo estimator is given by

$$\widehat{p(\mathbf{y})}_{\text{CMC}} = \frac{1}{M} \sum_{m=1}^M p(\mathbf{Y} | \mathbf{h}^{(m)}, \kappa^{(m)}),$$

where $\mathbf{h}^{(1)}, \dots, \mathbf{h}^{(M)}$ are drawn from the prior $p(\mathbf{h}) = \int p(\mathbf{h} | \phi, \sigma^2) p(\phi, \sigma^2) d(\phi, \sigma^2)$ and $\kappa^{(1)}, \dots, \kappa^{(M)}$ are from the prior $p(\kappa)$.

The conditional Monte Carlo estimator $\widehat{p(\mathbf{y})}_{\text{CMC}}$ alone does not work well in practice, because most draws from the prior distributions $p(\mathbf{h})$ and $p(\kappa)$ do not coincide with the high-density region of $p(\mathbf{Y} | \mathbf{h}, \kappa)$ (as a function of \mathbf{h} and κ). To tackle this problem, next we combine the conditional Monte Carlo approach with importance sampling, where the importance sampling density is obtained using the improved cross-entropy method. Chan and Eisenstat (2015) first propose using the improved cross-entropy method to estimate the marginal likelihood, but the models considered there are low dimensional. In contrast, for our settings with a large number of latent variables, additional consideration on parsimonious parameterization is required for obtaining appropriate importance sampling densities.

More specifically, we consider the parametric family

$$\mathcal{F} = \{f_{\mathcal{N}}(\mathbf{h}; \widehat{\mathbf{h}}, \widehat{\mathbf{K}}_{\mathbf{h}}^{-1}) f_{\mathcal{N}}(\phi; \widehat{\phi}, \widehat{K}_{\phi}^{-1}) f_{\mathcal{G}}(\kappa; \widehat{\nu}_{\kappa}, \widehat{S}_{\kappa})\},$$

where $f_{\mathcal{N}}$ and $f_{\mathcal{G}}$ denote Gaussian and gamma densities, respectively. Moreover, $\widehat{\mathbf{h}}$ and $\widehat{\mathbf{K}}_{\mathbf{h}}$ are respectively the mean vector and precision matrix—i.e., inverse covariance matrix—of the T -variate Gaussian density of \mathbf{h} ; $\widehat{\phi}$ and \widehat{K}_{ϕ} are the mean and precision of the

univariate Gaussian density of ϕ ; and $\widehat{\nu}_\kappa$ and \widehat{S}_κ are respectively the shape and rate of the gamma importance density for κ .⁶ Now, we aim to choose these parameters so that the associated member in \mathcal{F} is the closest in cross-entropy distance to the theoretical zero-variance importance sampling density $p(\mathbf{h}, \phi, \kappa | \mathbf{y}) \propto p(\mathbf{Y} | \mathbf{h}, \kappa)p(\mathbf{h} | \phi)p(\phi, \kappa)$, which is simply the marginal posterior density. Draws from this marginal distribution can be obtained using the posterior sampler described in Appendix A—i.e., we obtain posterior draws from the full posterior $p(\mathbf{A}, \mathbf{\Sigma}, \mathbf{h}, \phi, \sigma^2, \kappa | \mathbf{y})$ and keep only the draws of \mathbf{h} , ϕ , and κ .

Now, given the posterior draws $(\mathbf{h}^{(1)}, \phi^{(1)}, \kappa^{(1)}), \dots, (\mathbf{h}^{(M)}, \phi^{(M)}, \kappa^{(M)})$ and the parametric family \mathcal{F} , the optimization problem in (9) can be divided into 3 lower-dimensional problems: 1) obtain $\widehat{\mathbf{h}}$ and $\widehat{\mathbf{K}}_{\mathbf{h}}$ given the posterior draws of \mathbf{h} ; 2) obtain $\widehat{\phi}$ and \widehat{K}_ϕ given the posterior draws of ϕ ; and 3) obtain $\widehat{\nu}_\kappa$ and \widehat{S}_κ given the posterior draws of κ . The latter two steps are easy as both densities are univariate, and they can be solved using similar methods as described in Chan and Eisenstat (2015). Below we focus on the first step.

If we do not impose any restrictions on the Gaussian density $f_{\mathcal{N}}(\mathbf{h}; \widehat{\mathbf{h}}, \widehat{\mathbf{K}}_{\mathbf{h}}^{-1})$, in principle we can obtain $\widehat{\mathbf{h}}$ and $\widehat{\mathbf{K}}_{\mathbf{h}}$ analytically given the posterior draws $\mathbf{h}^{(1)}, \dots, \mathbf{h}^{(M)}$. However, there are two related issues with this approach, First, if unrestricted, the precision matrix $\widehat{\mathbf{K}}_{\mathbf{h}}$ is a full, $T \times T$ matrix. Evaluating and sampling from this Gaussian density would be time-consuming. Second, since $\widehat{\mathbf{K}}_{\mathbf{h}}$ is a symmetric but otherwise unrestricted matrix, there are $T(T + 1)/2$ parameters to be estimated. Consequently, one would require a large number of posterior draws to ensure that $\widehat{\mathbf{h}}$ and $\widehat{\mathbf{K}}_{\mathbf{h}}$ are accurately estimated.

In view of these potential difficulties, we consider instead a restricted family of Gaussian densities that exploits the time-series structure. Specifically, this family is parameterized by the parameters ρ , $\mathbf{a} = (a_1, \dots, a_T)'$ and $\mathbf{b} = (b_1, \dots, b_T)'$ as follows. First, h_1 has the marginal Gaussian distribution $h_1 \sim \mathcal{N}(a_1, b_1)$. For $t = 2, \dots, T$,

$$h_t = a_t + \rho h_{t-1} + \eta_t, \quad \eta_t \sim \mathcal{N}(0, b_t). \quad (11)$$

In other words, the joint distribution of \mathbf{h} is implied by an AR(1) process with time-varying intercepts and variances. It is easy to see that this parametric family includes the prior density of \mathbf{h} implied by the state equation (3) as a member.⁷

⁶Note that σ^2 only appears in the prior of \mathbf{h} , and one can integrate it out analytically: $p(\mathbf{h}) = \int p(\mathbf{h} | \phi, \sigma^2)p(\phi, \sigma^2)d(\phi, \sigma^2) = \int p(\mathbf{h} | \phi)p(\phi)d\phi$, where $p(\mathbf{h} | \phi)$ has an analytical expression, which is given in Appendix A. Hence, there is no need to simulate σ^2 .

⁷We also investigate other Gaussian families such as the joint distributions implied by the AR(2),

To facilitate computations, we vectorize the AR(1) process and write

$$\mathbf{H}_\rho \mathbf{h} = \mathbf{a} + \boldsymbol{\eta}, \quad \boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, \mathbf{B}),$$

where $\mathbf{B} = \text{diag}(b_1, \dots, b_T)$ and

$$\mathbf{H}_\rho = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ -\rho & 1 & \cdots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & -\rho & 1 \end{pmatrix}.$$

Since $|\mathbf{H}_\rho| = 1$ for any ρ , \mathbf{H}_ρ is invertible. Then, the Gaussian distribution implied by the AR(1) process in (11) has the form $\mathbf{h} \sim \mathcal{N}(\mathbf{H}_\rho^{-1} \mathbf{a}, (\mathbf{H}_\rho' \mathbf{B}^{-1} \mathbf{H}_\rho)^{-1})$. Note that the number of parameters here is only $2T + 1$ instead of unrestricted case of $T + T(T + 1)/2$.

Next, given this parametric family and posterior draws $\mathbf{h}^{(1)}, \dots, \mathbf{h}^{(M)}$, we can solve the maximization problem (9) in two steps. First, note that given ρ , we can obtain the maximizer $\hat{\mathbf{a}} = (\hat{a}_1, \dots, \hat{a}_T)'$ and $\hat{\mathbf{b}} = (\hat{b}_1, \dots, \hat{b}_T)'$ analytically. More specifically, by maximizing the log-likelihood

$$\ell(\rho, \mathbf{a}, \mathbf{b}) = -\frac{TM}{2} \log(2\pi) - \frac{M}{2} \log |\mathbf{B}| - \frac{1}{2} \sum_{m=1}^M (\mathbf{H}_\rho \mathbf{h}^{(m)} - \mathbf{a})' \mathbf{B}^{-1} (\mathbf{H}_\rho \mathbf{h}^{(m)} - \mathbf{a})$$

with respect to \mathbf{a} and \mathbf{b} , we obtain the maximizer $\hat{\mathbf{a}} = \frac{1}{M} \sum_{m=1}^M \mathbf{H}_\rho \mathbf{h}^{(m)}$ and $\hat{b}_t = \frac{1}{M} \sum_{m=1}^M (h_t^{(m)} - \hat{a}_t)^2, t = 1, \dots, T$.

Second, given the analytical solution $\hat{\mathbf{a}}$ and $\hat{\mathbf{b}}$ —which are functions of ρ —we can readily evaluate the one-dimensional concentrated log-likelihood $\ell(\rho, \hat{\mathbf{a}}, \hat{\mathbf{b}})$. Then, $\hat{\rho}$ can be obtained numerically by maximizing $\ell(\rho, \hat{\mathbf{a}}, \hat{\mathbf{b}})$ with respect to ρ . Finally, we use $f_{\mathcal{N}}(\mathbf{h}; \mathbf{H}_{\hat{\rho}}^{-1} \hat{\mathbf{a}}, (\mathbf{H}_{\hat{\rho}}' \hat{\mathbf{B}}^{-1} \mathbf{H}_{\hat{\rho}})^{-1})$, where $\hat{\mathbf{B}} = \text{diag}(\hat{b}_1, \dots, \hat{b}_T)$, as the importance sampling density for \mathbf{h} .

Chan and Eisenstat (2018) also consider a Gaussian importance sampling density for the log-volatility \mathbf{h} . However, there the approximation is based on a second-order Taylor expansion at the mode. Hence, it is a local approximation that might not be close to

MA(1) and ARMA(1,1) processes. None of them perform noticeably better than the simple AR(1) process in (11).

the target density at points away from the mode. In contrast, our proposed Gaussian importance sampling density is a global approximation that takes into account of the whole support of the distribution, and is therefore expected to behave more similarly to the target distribution $p(\mathbf{h} | \mathbf{y})$ over its entire support.

4 Monte Carlo Experiments

In this section we conduct a series of Monte Carlo experiments to assess how the proposed marginal likelihood estimators perform in selecting the correct data generating process. More specifically, we focus on the following three questions. First, can we distinguish the three stochastic volatility models described in Section 2? Second, can we discriminate between models with time-varying volatility against homoskedastic errors? Finally, for factor stochastic volatility models, can we identify the correct number of factors?

4.1 Can We Distinguish the Stochastic Volatility Specifications?

In the first part of this Monte Carlo exercise, we address the question of whether one can distinguish the three stochastic volatility models—namely, the common stochastic volatility model (VAR-CSV), the Cholesky stochastic volatility model (VAR-SV) and the factor stochastic volatility model (VAR-FSV)—in the context of large Bayesian VARs. To that end, we generate 100 datasets from each of the three models. Each dataset consists of $n = 10$ variables, $T = 400$ observations and $p = 2$ lags. Given each dataset, we compute the log marginal likelihoods of the same three models using the proposed method described in Section 3.

For VAR-CSV, we generate the intercepts from $\mathcal{U}(-10, 10)$. The diagonal elements of the first VAR coefficient matrix are iid $\mathcal{U}(-0.2, 0.4)$ and the off-diagonal elements are $\mathcal{U}(-0.2, 0.2)$; all elements of the second VAR coefficient matrix are iid $\mathcal{N}(0, 0.05^2)$. The error covariance matrix Σ is generated from the inverse-Wishart distribution $\mathcal{IW}(n + 5, 0.7 \times \mathbf{I}_n + 0.3 \times \mathbf{1}_n \mathbf{1}_n')$, where $\mathbf{1}_n$ is an $n \times 1$ column of ones. For the log-volatility process, we set $\phi = 0.98$ and $\sigma^2 = 0.1$. For VAR-SV, the VAR coefficients are generated as before, and the free elements of the impact matrix are generated independently from the $\mathcal{N}(0, 0.5^2)$ distribution. For the log-volatility process, we set $\mu_i = -1$, $\phi_i = 0.98$

and $\sigma_i^2 = 0.1$ for $i = 1, \dots, n$. Finally, for VAR-FSV, we set the number of factor r to be 3. The VAR coefficients are generated the same way as other models. For the log-volatility process, we set $\mu_i = -1$, $\phi_i = 0.98$ and $\sigma_i^2 = 0.1$ for $i = 1, \dots, n$, and $\mu_{n+j} = 0$, $\phi_{n+j} = 0.98$ and $\sigma_{n+j}^2 = 0.1$ for $j = 1, \dots, r$. That is, the log stochastic volatility processes associated with the factors have a larger mean, but are otherwise the same as the idiosyncratic stochastic volatility processes.⁸ In the first experiment, we generate 100 datasets from VAR-CSV as described above. For each dataset, we then compute the log marginal likelihoods of VAR-SV and VAR-FSV relative to that of the true model VAR-CSV. Specifically, we subtract the latter log marginal likelihood from the log marginal likelihoods of VAR-SV and VAR-FSV. The results are reported in Figure 1.

Since a model is preferred by the data if it has a larger log marginal likelihood value, a difference that is negative indicates that the correct model is favored. It is clear from the histograms that for all the datasets the correct model VAR-CSV compares favorably to the other two stochastic volatility specifications, often by a large margin.

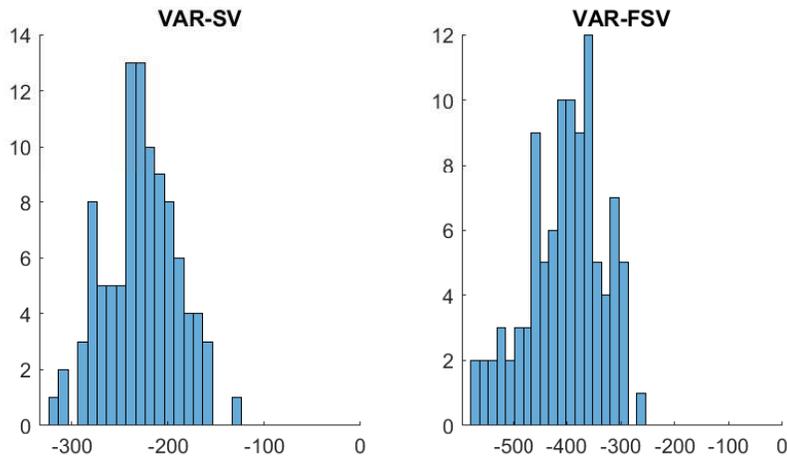


Figure 1: Histograms of log marginal likelihoods under VAR-SV (left panel) and VAR-FSV (right panel) relative to the true model (VAR-CSV). A negative value indicates that the correct model is favored.

Next, we generate 100 datasets from VAR-SV. For each dataset, we then compute the

⁸Since the main goal of the Monte Carlo experiments is to assess if one can distinguish the stochastic volatility specifications, we aim to use similar priors across the models so as to minimize their impacts. In particular, for the hierarchical Minnesota priors we fix all shrinkage hyperparameters to be $\kappa = \kappa_1 = \kappa_2 = 0.2^2$. That is, we turn off the cross-variable shrinkage feature of the Minnesota priors under VAR-SV and VAR-FSV, so that they are comparable to the symmetric Minnesota prior under VAR-CSV.

log marginal likelihoods of VAR-CSV and VAR-FSV relative to that of the true model. The results are reported in Figure 2. Again, the model comparison result shows that, for all datasets, the correct model VAR-SV is overwhelmingly favored compared to the other two stochastic volatility specifications.

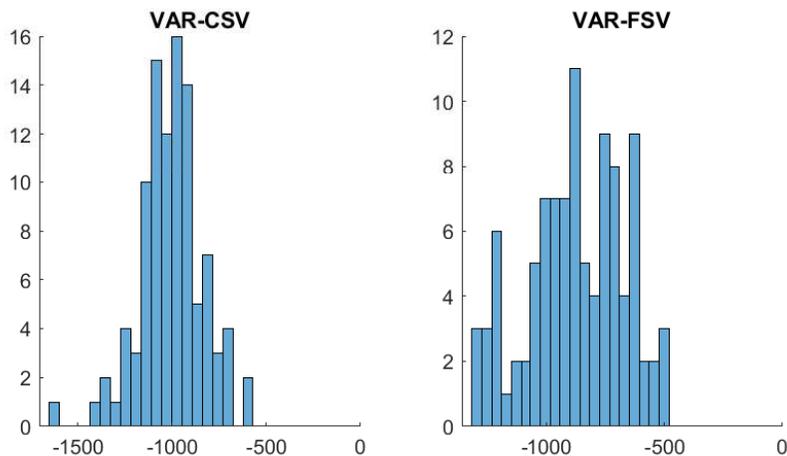


Figure 2: Histograms of log marginal likelihoods under VAR-CSV (left panel) and VAR-FSV (right panel) relative to the true model (VAR-SV). A negative value indicates that the correct model is favored.

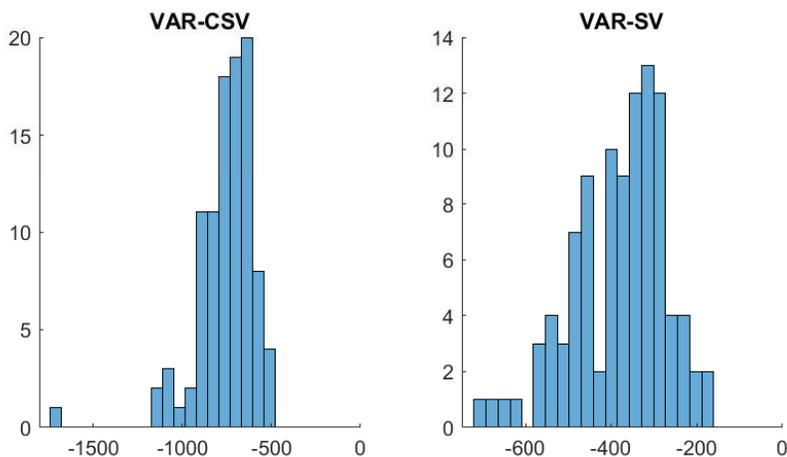


Figure 3: Histograms of log marginal likelihoods under VAR-CSV (left panel) and VAR-SV (right panel) relative to the true model (VAR-FSV). A negative value indicates that the correct model is favored.

Lastly, we generate 100 datasets from VAR-FSV, and for each dataset we compute the log

marginal likelihoods of VAR-CSV and VAR-SV relative to that of VAR-FSV. The results are reported in Figure 3. Once again, these results show that the proposed marginal likelihood estimators can select the correct model for all the simulated datasets.

Appendix C provides additional Monte Carlo results to further assess the performance of the proposed marginal likelihood estimators in higher dimensions ($n = 20$) and noisier VAR equations (magnitudes of the error covariance matrices increase 10-fold). For these additional settings, the correct model is again overwhelmingly favored compared to the other two stochastic volatility models. All in all, this series of Monte Carlo experiments show that using the proposed marginal likelihood estimators one can clearly distinguish the three stochastic volatility specifications, even for a moderate number of variables and sample size.

4.2 Can We Discriminate between Models with Time-Varying Volatility against Homoskedastic Errors?

Since VARs with stochastic volatility are by design more flexible than conventional VARs with homoskedastic errors, one concern of using these more flexible models is that they might overfit the data. In this Monte Carlo experiment we investigate if the proposed marginal likelihood estimators can distinguish models with and without stochastic volatility. More specifically, we generate 100 datasets from a standard VAR with homoskedastic errors. For each dataset, we then compute the log marginal likelihoods of the three stochastic volatility models: VAR-CSV, VAR-SV and VAR-FSV.⁹ The differences in log marginal likelihoods relative to the homoskedastic VAR are reported in Figure 4.

Recall that a model is preferred by the data if it has a larger log marginal likelihood value. Hence, a negative value indicates that the correct homoskedastic VAR is selected. Overall, the histograms show that for almost all datasets the correct homoskedastic model compares favorably to the three stochastic volatility models. This Monte Carlo experiment highlights that the marginal likelihood does not always favor the model with the best model-fit; it also has a built-in penalty for model complexity. And it is only when

⁹We also generate data from the stochastic volatility models and compare their marginal likelihoods with a homoskedastic VAR. The results overwhelmingly favor the stochastic volatility models as the homoskedastic VAR does not fit the data well at all. For space constraint, however, we do not report these model comparison results.

the gain in model-fit outweighs the additional model complexity does the more complex model have a larger marginal likelihood value.

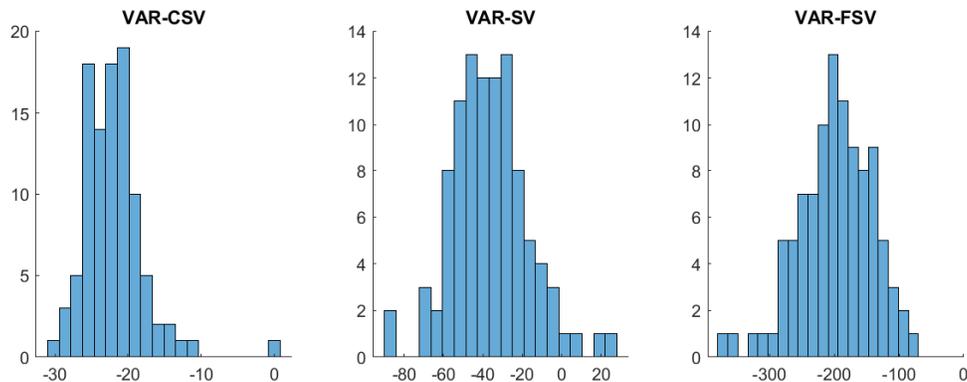


Figure 4: Histograms of log marginal likelihoods under VAR-CSV (left panel), VAR-SV (middle panel) and VAR-FSV (right panel) relative to the true model (homoskedastic VAR). A negative value indicates that the correct model is favored.

It is also interesting to note that VAR-CSV generally performs better than the other two stochastic volatility models. This is perhaps not surprising, as the VAR-CSV is the most parsimonious extension of the homoskedastic VAR—it has only one stochastic volatility process h_t , and when h_t is identically zero it replicates the homoskedastic VAR. Nevertheless, even for this closely related extension, the marginal likelihood criterion clearly indicates that the stochastic volatility process is spurious and it tends to favor the homoskedastic model.

4.3 Can We Identify the Correct Number of Factors?

One important specification choice for factor stochastic volatility models is to select the number of factors. This represents the trade-off between a more parsimonious model with fewer stochastic volatility factors versus a more complex model with more factors but better model-fit. In this Monte Carlo experiment, we investigate if the proposed marginal likelihood estimator for the VAR-FSV can pick the correct number of factors.

To that end, we generate 100 datasets from VAR-FSV with $r = 3$ factors. Then, for each dataset, we compute the log marginal likelihood of VAR-FSV models with $r = 2, 3$ and 4

factors. This set of number of factors is chosen to shed light on the effects of under-fitting versus over-fitting. We report the log marginal likelihoods of the 2- and 4-factor models relative to the true 3-factor model in Figure 5. The results clearly show that the proposed method is able to identify the correct number of factors. In particular, for all datasets the 3-factor model outperforms the more parsimonious 2-factor model and the more flexible 4-factor model.

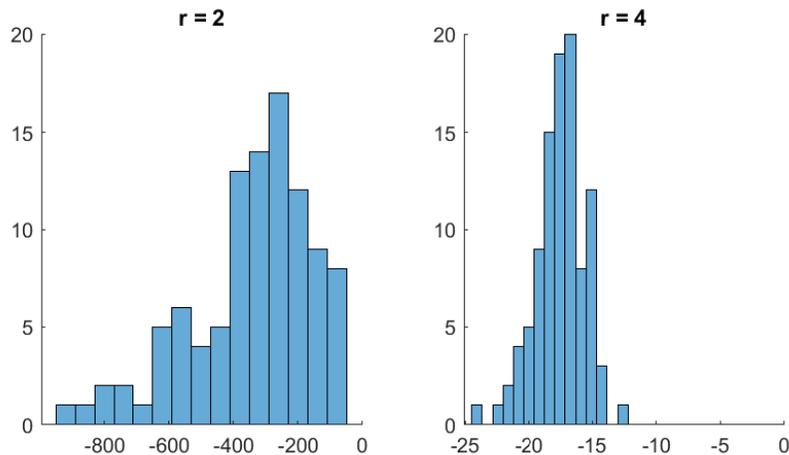


Figure 5: Histograms of log marginal likelihoods under the 2-factor model (left panel) and the 4-factor model (right panel) relative to the true 3-factor model. A negative value indicates that the correct model is favored.

The 2-factor model under-fits the data and performs much worse than the correct 3-factor model. On the other hand, while the 4-factor model is able to replicate features of the 3-factor model, it also includes spurious features that tend to over-fit the data. Consequently, the 3-factor model is strongly preferred by the data over the two alternatives. However, the impacts of under- and over-fitting are not symmetric. In particular, under-fitting receives a much heavier penalty than over-fitting, as illustrated by the much larger differences (in magnitude) between the 2- and 3-factor models than those between the 3- and 4-factor models. Regardless of these differences, we conclude that the proposed method is able to select the correct number of factors.

5 Empirical Application

In this empirical application we compare the three stochastic volatility specifications—i.e., the common stochastic volatility model (VAR-CSV), the Cholesky stochastic volatility model (VAR-SV) and the factor stochastic volatility model (VAR-FSV)—in the context of Bayesian VARs of different dimensions. We use a dataset that consists of 30 US quarterly variables with a sample period from 1959Q1 to 2019Q4. It is constructed from the FRED-QD database at the Federal Reserve Bank of St. Louis as described in McCracken and Ng (2020). The dataset contains a range of widely used macroeconomic and financial variables, including Real GDP and its components, various measures of inflation, labor market variables and interest rates. They are transformed to stationarity, typically to annualized growth rates. We consider VARs that are small ($n = 7$), medium ($n = 15$) and large ($n = 30$). The complete list of variables for each dimension and how they are transformed is given in Appendix B.

In contrast to the Monte Carlo study where the shrinkage hyperparameters in the hierarchical Minnesota priors are fixed at certain subjective values, here we treat them as unknown parameters to be estimated. This is motivated by a few recent papers that find this data-based approach outperforms subjective prior elicitation in both in-sample model fit and out-of-sample forecast performance (see, e.g., Carriero, Clark, and Marcellino, 2015; Giannone, Lenza, and Primiceri, 2015; Amir-Ahmadi, Matthes, and Wang, 2020). For easy comparison, we set the lag length to be $p = 4$ for all VARs.¹⁰ We compute the log marginal likelihoods of the VARs using the proposed hybrid algorithm described in Section 3. The importance sampling density for each model is constructed using the cross-entropy method with 20,000 posterior draws after a burn-in period of 1,000. Then, the log marginal likelihood estimate is computed using an importance sample of size 10,000.

¹⁰The lag length can of course be chosen by comparing the marginal likelihood. In a preliminary analysis, we find that a lag length of $p = 4$ is generally sufficient. In addition, with our flexible hierarchical Minnesota priors on the VAR coefficients, adding more lags than necessary does not substantially affect model performance. For instance, for all of the three stochastic volatility specifications, the best models for $n = 7$ have $p = 3$ lags. Adding one more lag reduces the log marginal likelihood by only about 1-2 for all models.

5.1 Comparing Stochastic Volatility Specifications

We first report the main model comparison results on comparing the three stochastic volatility specifications: VAR-CSV, VAR-SV and VAR-FSV. As a benchmark, we also include the standard homoskedastic VAR with the natural conjugate prior (its marginal likelihood is available analytically). In addition to the widely different likelihoods implied by these stochastic volatility specifications, they also differ in the flexibility of the shrinkage priors employed. More specifically, both the standard VAR and VAR-CSV use the natural conjugate prior, which has only one hyperparameter that controls the shrinkage strength of all VAR coefficients. In contrast, both the priors under VAR-SV and VAR-FSV can accommodate cross-variable shrinkage. That is, there are two separate shrinkage hyperparameters, one controls the shrinkage strength on own lags, whereas the other controls that of lags of other variables. In what follows we focus on the overall ranking of the models; we will investigate the role of the priors in the next section.

Table 1 reports the log marginal likelihood estimates of the four VARs across the three model dimensions ($n = 7, 15, 30$). First, it is immediately apparent that all three stochastic volatility models perform substantially better than the standard homoskedastic VAR. For example, the log marginal likelihood difference between VAR-CSV and VAR is about 760 for $n = 30$, highlighting overwhelming evidence in favor of the heteroskedastic model. This finding is in line with the growing body of evidence that shows the importance of time-varying volatility in modeling both small and large macroeconomic datasets.

Table 1: Log marginal likelihood estimates (numerical standard errors) of a standard homoskedastic VAR, VAR-CSV, VAR-SV and VAR-FSV.

	VAR	VAR-CSV	VAR-SV	VAR-FSV		
				$k = 1$	$k = 2$	$k = 3$
$n = 7$	-2,583.4	-2,410	-2,315	-2,318	-2,320	-2,331
	-	(0.1)	(0.3)	(0.3)	(0.7)	(0.7)
				$k = 3$	$k = 4$	$k = 5$
$n = 15$	-6,918.8	-6,618	-6,443	-6,468	-6,454	-6,485
	-	(0.1)	(0.4)	(0.7)	(0.8)	(1.3)
				$k = 9$	$k = 10$	$k = 11$
$n = 30$	-12,783.4	-12,024	-11,382	-11,571	-11,567	-11,608
	-	(0.1)	(0.6)	(1.2)	(1.8)	(1.8)

Second, among the three stochastic volatility specifications, the data overwhelmingly prefers VAR-SV and VAR-FSV over the more restrictive VAR-CSV for all model dimensions, due to a combination of the more flexible likelihoods and priors. For example, the log marginal likelihood differences between VAR-SV and VAR-CSV are 95 for $n = 7$, 175 for $n = 15$ and 642 for $n = 30$. Third, for all three model dimensions, VAR-SV is the most favored stochastic volatility specification, though VAR-FSV comes in close second. Finally, we note that the optimal number of factors changes across the model dimension. It is perhaps not surprising that more factors are needed to model the more complex error covariance structure as the model dimension increases. For instance, for $n = 7$ the 1-factor model performs the best, whereas one needs 10 factors when $n = 30$.

To corroborate these model comparison results, we next perform a recursive out-of-sample forecasting exercise using an evaluation period from 1975Q1 to the end of the sample. More specifically, Table 2 reports the joint density forecasting results (as measured by the sum of log predictive likelihoods) for the four VARs across the three model dimensions (for VAR-FSV, we fix $k = 4$). It is clear from the table that one can draw mostly similar conclusions from these forecasting results. In particular, all three stochastic volatility models perform substantially better than the homoskedastic VAR for both 1- and 4-step-ahead forecast horizons and for all model dimensions. This again highlights the importance of allowing for some form of time-varying volatility. In addition, for 1-step-ahead, both VAR-SV and VAR-FSV forecast better than the more restrictive VAR-CSV for all model dimensions. However, when the forecast horizon increases, VAR-CSV tends to perform better for higher model dimensions.

Table 2: Sum of log predictive likelihoods of a standard homoskedastic VAR, VAR-CSV, VAR-SV and VAR-FSV. A larger value indicates better forecasting performance.

	1-step-ahead				4-step-ahead			
	VAR	VAR-CSV	VAR-SV	VAR-FSV	VAR	VAR-CSV	VAR-SV	VAR-FSV
$n = 7$	-1,858	-1,719	-1,625	-1,613	-2,573	-2,321	-2,311	-2,312
$n = 15$	-5,070	-4,764	-4,612	-4,604	-6,243	-5,631	-5,821	-5,772
$n = 30$	-11,271	-8,579	-8,035	-8,424	-15,715	-10,030	-11,739	-11,348

5.2 Comparing Shrinkage Priors

In this section we compare different types of Minnesota priors for each of the three stochastic volatility specifications. In particular, we investigate the potential benefits of allowing for cross-variable shrinkage and selecting the overall shrinkage hyperparameters in a data-driven manner. To that end, we consider two useful benchmarks. First, for VAR-SV and VAR-FSV we consider the special case where $\kappa_1 = \kappa_2$ —i.e., we turn off the cross-variable shrinkage and require the shrinkage hyperparameters on own and other lags to be the same. We refer to this version as the symmetric prior. The second benchmark is a set of subjectively chosen hyperparameter values that apply cross-variable shrinkage. In particular, we follow Carriero, Clark, and Marcellino (2015) and use the values $\kappa_1 = 0.04$ and $\kappa_2 = 0.0016$. This second benchmark is referred to as the subjective prior. Finally, our baseline prior, where κ_1 and κ_2 are estimated from the data and could potentially be different, is referred to as the asymmetric prior.

To fix ideas, we focus on VARs with $n = 15$ variables. Table 3 reports the log marginal likelihood estimates of the three stochastic volatility specifications with different shrinkage priors. First, for both VAR-SV and VAR-FSV, the asymmetric prior significantly outperforms the symmetric version that requires $\kappa_1 = \kappa_2$. This result suggests that it is beneficial to shrink the coefficients on own lags differently than those on lags of other variables. This makes intuitive sense as one would expect that, on average, a variable’s own lags would be more important for its future evolution than lags of other variables. By relaxing the restriction that $\kappa_1 = \kappa_2$, the log marginal likelihood values of VAR-SV and VAR-FSV increase by 177 and 204, respectively.

Table 3: Log marginal likelihood estimates (numerical standard errors) of the stochastic volatility specifications with different shrinkage priors for $n = 15$.

	VAR-CSV	VAR-SV	VAR-FSV ($k = 4$)
Subjective prior	−6,702 (0.1)	−6,474 (0.6)	−6,491 (0.9)
Symmetric prior	−6,618 (0.1)	−6,620 (0.4)	−6,658 (1.1)
Asymmetric prior	-	−6,443 (0.7)	−6,454 (0.8)

In addition, there are also substantial benefits of allowing the shrinkage hyperparameters

to be estimated instead of fixing them subjectively. For example, the log marginal likelihood value of VAR-CSV with the symmetric prior is 84 larger than that of the subjective prior; the log marginal likelihood value of VAR-SV with the asymmetric prior is 31 larger than that of the subjective prior. These results suggest that while those widely-used subjective hyperparameter values seem to work well for some datasets, they might not be suitable for others that contain different variables and span different sample periods.

We reported in last section that the data overwhelmingly preferred the more flexible VAR-SV and VAR-FSV over VAR-CSV. But it was unclear whether it was due to the more flexible likelihoods or priors (recall that VAR-CSV can only accommodate the symmetric prior). The result here suggests that the superior performance of VAR-SV and VAR-FSV can be attributed to the more flexible priors rather than the more flexible stochastic volatility specifications. In particular, both VAR-SV and VAR-FSV with the symmetric prior actually perform slightly worse than VAR-CSV (respectively, $-6,620$, $-6,658$ and $-6,618$ for the three models). Only when the asymmetric prior is used do VAR-SV and VAR-FSV outperform VAR-CSV. This stark result shows that choosing a suitable shrinkage prior in high-dimensional settings is as important as, if not more important than, selecting a suitable stochastic volatility specification.

Next, Table 4 reports the posterior estimates of the shrinkage hyperparameters under the three stochastic volatility specifications. First, under VAR-CSV, the posterior mean of κ varies across the model dimensions, from 0.32 for $n = 7$ to 0.22 for $n = 15$ and 0.1 for $n = 30$. This finding highlights the fact that any particular subjectively elicited value—e.g., the widely-used value of 0.04 for the natural conjugate prior; see [Carriero, Clark, and Marcellino \(2016\)](#) and [Chan \(2020\)](#)—is unlikely to be suitable for all datasets of different dimensions. In addition, the estimates also suggest that when the model dimension increases, more aggressive shrinkage is needed, confirming the findings in, e.g., [Banbura, Giannone, and Reichlin \(2010\)](#).

Second, it is clear that the estimates of κ_1 are orders of magnitude larger than those of κ_2 under VAR-SV and VAR-FSV, the two models that accommodate cross-variable shrinkage. For example, for VAR-FSV with $n = 15$, the estimates of κ_1 and κ_2 are, respectively, 0.17 and 0.0018, a difference of 2 orders of magnitude. These estimates confirm the model comparison results above and suggest that the data strongly prefers shrinking the coefficients on lags of other variables much more aggressively to zero than

those on own lags. Again, this is intuitive as one would expect that own lags are more informative, on average, than lags of other variables in forecasting a variable’s future evolution. In addition, these estimates are also rather different from those widely-used subjectively chosen values of $\kappa_1 = 0.04$ and $\kappa_2 = 0.0016$, reinforcing the conclusion that any particular set of fixed hyperparameter values is unlikely to be suitable for all datasets of different dimensions.

Table 4: Posterior means (standard deviations) of the shrinkage hyperparameters in VAR-CSV, VAR-SV and VAR-FSV.

	VAR-CSV	VAR-SV		VAR-FSV	
	κ	κ_1	κ_2	κ_1	κ_2
$n = 7$	0.32 (0.065)	0.23 (0.051)	0.0032 (0.0014)	0.24 (0.054)	0.0037 (0.0015)
$n = 15$	0.22 (0.027)	0.16 (0.030)	0.0022 (0.0006)	0.17 (0.032)	0.0018 (0.0004)
$n = 30$	0.10 (0.008)	0.15 (0.020)	0.00002 (0.00001)	0.15 (0.020)	0.00003 (0.00001)

All in all, our results confirm the substantial benefits of allowing for cross-variable shrinkage and selecting the shrinkage hyperparameters using a data-based approach. They also highlight that in high-dimensional settings, choosing a flexible shrinkage prior is as important as selecting a flexible stochastic volatility specification.

6 Extensions

The proposed marginal likelihood estimators are based on two key ingredients: a conditional Monte Carlo component to integrate out the large number of VAR coefficients analytically and an adaptive importance sampling method to integrate out the log-volatility via Monte Carlo simulation. These two ideas can be more widely applied to other marginal likelihood estimators or alternative model comparison methods to make them viable in high-dimensional settings, as we demonstrate below. In addition, we further illustrate how the proposed methods can be used to compute the marginal likelihood of a stochastic volatility model with an outlier component.

6.1 Improving the Modified Harmonic Mean Estimator

We first illustrate how one can use the conditional Monte Carlo and adaptive importance sampling methods to improve alternative estimators of the marginal likelihood. To fix ideas, we focus on the modified harmonic mean estimator proposed by Gelfand and Dey (1994); similar improvements can be made for other estimators such as Chib’s method (Chib, 1995; Chib and Jeliazkov, 2001) and sequential Monte Carlo (e.g., Bognanni and Zito, 2020).

The theoretical justification of the modified harmonic mean estimator is based on the following observation: for any density function f with support contained in the support of the posterior density, we have

$$\mathbb{E} \left(\frac{f(\boldsymbol{\theta})}{p(\boldsymbol{\theta})p(\mathbf{y}|\boldsymbol{\theta})} \mid \mathbf{y} \right) = \int \frac{f(\boldsymbol{\theta})}{p(\boldsymbol{\theta})p(\mathbf{y}|\boldsymbol{\theta})} \frac{p(\boldsymbol{\theta})p(\mathbf{y}|\boldsymbol{\theta})}{p(\mathbf{y})} d\boldsymbol{\theta} = p(\mathbf{y})^{-1}, \quad (12)$$

where the expectation is taken with respect to the posterior distribution $p(\boldsymbol{\theta}|\mathbf{y}) = p(\boldsymbol{\theta})p(\mathbf{y}|\boldsymbol{\theta})/p(\mathbf{y})$. It follows that one can estimate the marginal likelihood $p(\mathbf{y})$ using the following estimator:

$$\text{GD} = \left\{ \frac{1}{R} \sum_{r=1}^R \frac{f(\boldsymbol{\theta}^{(r)})}{p(\boldsymbol{\theta}^{(r)})p(\mathbf{y}|\boldsymbol{\theta}^{(r)})} \right\}^{-1}, \quad (13)$$

where $\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(R)}$ are posterior draws. This modified harmonic mean is simulation-consistent—it converges to $p(\mathbf{y})$ in probability as R tends to infinity—but it is not unbiased in general.

Even though the estimator in (13) is simulation-consistent for any density f with support contained in the support of $p(\boldsymbol{\theta}|\mathbf{y})$, its finite-sample properties depend critically on the choice of f . Geweke (1999) shows that if f has tails lighter than those of $p(\boldsymbol{\theta}|\mathbf{y})$, the estimator has a finite variance. He further recommends a Gaussian approximation of $p(\boldsymbol{\theta}|\mathbf{y})$ with tail truncations determined by asymptotic arguments. More specifically, let $\hat{\boldsymbol{\theta}}$ and $\mathbf{Q}_{\boldsymbol{\theta}}$ denote the posterior mean and covariance matrix respectively. Then, f is set to be the $\mathcal{N}(\hat{\boldsymbol{\theta}}, \mathbf{Q}_{\boldsymbol{\theta}})$ density truncated within the region

$$\{\boldsymbol{\theta} \in \mathbb{R}^m : (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}})' \mathbf{Q}_{\boldsymbol{\theta}}^{-1} (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}) < \chi_{\alpha, m}^2\},$$

where $\chi_{\alpha, m}^2$ is the $(1 - \alpha)$ quantile of the χ_m^2 distribution and m is the dimension of $\boldsymbol{\theta}$. The GD estimator in (13) with the truncated Gaussian tuning function is commonly-used in economics; examples include Sims, Waggoner, and Zha (2008) and Justiniano and Primiceri (2008).

One potential difficulty of implementing this approach in high-dimensional setting is that the prescribed Gaussian density has a large number of parameters to be estimated. As an example, for the VAR-SV model with $T = 300$, $n = 20$, and $p = 4$, the Gaussian density is 2,133-dimensional and has over 2.2 million parameters. As such, it is impractical to estimate them with any reasonable accuracy using MCMC output.

Below we outline a few ways to make the GD estimator viable in high-dimensional settings using the proposed adaptive importance sampling and conditional Monte Carlo methods. First, we consider a tuning function of the form $f(\boldsymbol{\theta}) = f_{\mathbf{h}}(\mathbf{h})f_{\boldsymbol{\theta}_{-\mathbf{h}}}(\boldsymbol{\theta}_{-\mathbf{h}})$, where $\boldsymbol{\theta}_{-\mathbf{h}}$ denotes the set of model parameters and latent variables excluding the log-volatility \mathbf{h} . Then, we use the adaptive importance sampling method described in Section 3.2 to obtain $f_{\mathbf{h}}$, whereas $f_{\boldsymbol{\theta}_{-\mathbf{h}}}$ is a truncated Gaussian distribution constructed as described above. This approach reduces the number of parameters of f in two ways: 1) by factoring f into two lower-dimensional Gaussian densities, it reduces the number of covariance parameters; 2) as described in Section 3.2, the way $f_{\mathbf{h}}$ is parameterized reduces the number of parameters from $(T^2 + 3T)/2$ to $2T + 1$.¹¹ However, this approach alone is unlikely to work well as there remains a large number of parameters to be estimated. Using the same example of the VAR-SV with $T = 300$, $n = 20$, and $p = 4$, the tuning function has still over 1.6 million parameters.

Next, we incorporate the proposed conditional Monte Carlo estimator to improve the GD estimator. Using the notations in Section 3.3, partition $\boldsymbol{\theta}$ into two blocks, $\boldsymbol{\gamma}$ and $\boldsymbol{\psi}$, and suppose we can compute $g(\mathbf{y} | \boldsymbol{\psi}) \equiv \mathbb{E}[p(\mathbf{y} | \boldsymbol{\theta}) | \boldsymbol{\psi}] = \int p(\mathbf{y} | \boldsymbol{\gamma}, \boldsymbol{\psi})p(\boldsymbol{\gamma} | \boldsymbol{\psi})d\boldsymbol{\gamma}$ analytically, where $p(\boldsymbol{\gamma} | \boldsymbol{\psi})$ is the prior of $\boldsymbol{\gamma}$ that can potentially depend on $\boldsymbol{\psi}$. Since the marginal likelihood can be expressed as $p(\mathbf{y}) = \int g(\mathbf{y} | \boldsymbol{\psi})p(\boldsymbol{\psi})d\boldsymbol{\psi}$, a similar derivation as (12)

¹¹As discussed in Section 3.4, the choice of the restricted family of Gaussian densities exploits the time-series structure of the problem. In particular, the log-volatility is assumed to follow an AR(1) process, which implies the precision or inverse covariance matrix of the joint distribution is tridiagonal. This strikes the balance between flexibility and parsimony.

suggests that we can estimate $p(\mathbf{y})$ using the following estimator:

$$\text{GD}_\psi = \left\{ \frac{1}{R} \sum_{r=1}^R \frac{f_\psi(\boldsymbol{\psi}^{(r)})}{p(\boldsymbol{\psi}^{(r)})g(\mathbf{y} | \boldsymbol{\psi}^{(r)})} \right\}^{-1}, \quad (14)$$

where $\boldsymbol{\psi}^{(1)}, \dots, \boldsymbol{\psi}^{(R)}$ are posterior draws from $p(\boldsymbol{\psi} | \mathbf{y}) = g(\mathbf{y} | \boldsymbol{\psi})p(\boldsymbol{\psi})/p(\mathbf{y})$. These draws can be obtained using a standard MCMC algorithm: first obtain posterior draws of the pairs $(\boldsymbol{\gamma}^{(1)}, \boldsymbol{\psi}^{(1)}), \dots, (\boldsymbol{\gamma}^{(R)}, \boldsymbol{\psi}^{(R)})$ and then discard the first components. One can use the truncated Gaussian density described above as the tuning function f_ψ . Using the same VAR-SV example, the number of parameters of the tuning function here is drastically reduced to about 46,000. Even so, estimating these parameters with sufficient accuracy using MCMC output is still computationally intensive.

Finally, one can combine both approaches to further improve the performance of the estimator. Specifically, consider the estimator in (14) and a tuning function of the form $f_\psi(\boldsymbol{\psi}) = f_{\mathbf{h}}(\mathbf{h})f_{\boldsymbol{\psi}_{-\mathbf{h}}}(\boldsymbol{\psi}_{-\mathbf{h}})$, where $f_{\mathbf{h}}$ is obtained using the adaptive importance sampling method and $f_{\boldsymbol{\psi}_{-\mathbf{h}}}$ is a truncated Gaussian density. For the same VAR-SV example, the number of parameters for this combined approach is only 610, and they can be estimated with adequate accuracy using MCMC output.

To illustrate how these improved estimators perform empirically, we revisit the model comparison exercise in Section 5 and compute the marginal likelihood of VAR-CSV with $n = 15$ using two versions of the GD estimator. Recall that the set of model parameters and latent variables for VAR-CSV is $\boldsymbol{\theta} = \{\mathbf{A}, \boldsymbol{\Sigma}, \phi, \sigma^2, \kappa, \mathbf{h}\}$. We first consider the estimator in (14), where we integrate out $\boldsymbol{\gamma} = \{\mathbf{A}, \boldsymbol{\Sigma}, \sigma^2\}$ analytically and set f_ψ to be a truncated Gaussian density with $\boldsymbol{\psi} = \{\phi, \kappa, \mathbf{h}\}$. Second, we use the same estimator in (14), but with a different tuning function: we set $f_\psi(\boldsymbol{\psi}) = f_{\mathbf{h}}(\mathbf{h})f_{\boldsymbol{\psi}_{-\mathbf{h}}}(\boldsymbol{\psi}_{-\mathbf{h}})$, where $f_{\mathbf{h}}$ is obtained using the adaptive importance sampling method and $f_{\boldsymbol{\psi}_{-\mathbf{h}}}$ is a truncated bivariate Gaussian density with $\boldsymbol{\psi}_{-\mathbf{h}} = \{\phi, \kappa\}$. For each estimator, we obtain 100 marginal likelihood estimates, where each estimate is based on 50,000 posterior draws. To summarize the results we report the boxplots of the log marginal likelihood estimates in Figure 6. The middle line of each box denotes the median, while the lower and upper lines represent, respectively, the 25- and the 75-percentiles. The whiskers extend to the maximum and minimum. For comparison we also include the estimates from the proposed method.¹²

¹²The original GD estimator in (13) and the version with only the adaptive importance sampling

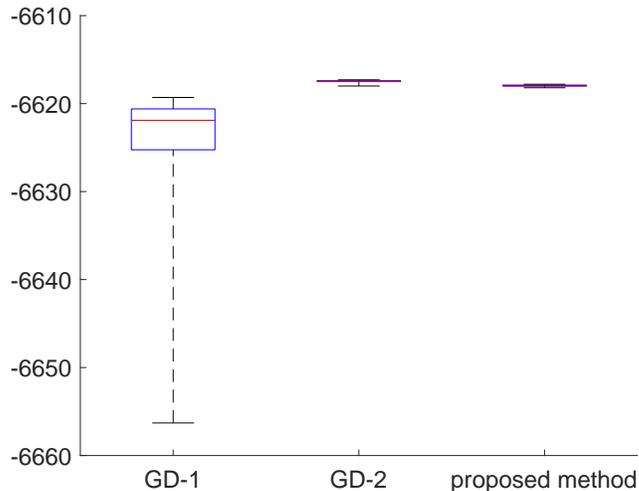


Figure 6: Boxplots of the log marginal likelihoods for VAR-CSV with $n = 15$ estimated using three estimators: the GD estimator in which $\gamma = \{\mathbf{A}, \mathbf{\Sigma}, \sigma^2\}$ is integrated out analytically (GD-1), the GD estimator in which γ is integrated out analytically and the adaptive importance sampling method is used to obtain $f_{\mathbf{h}}$ (GD-2) and the proposed method.

While the original GD estimator in (13) has too many parameters to provide stable estimates, the version in which $\gamma = \{\mathbf{A}, \mathbf{\Sigma}, \sigma^2\}$ is integrated out analytically (GD-1) performs reasonably well. In particular, the variance of the estimator is relatively small and most of the marginal likelihood estimates are between $-6,630$ and $-6,620$. Nevertheless, the distribution of the estimator has a heavier left tail and it seems to underestimate the value of the marginal likelihood. In contrast, when both the conditional Monte Carlo and adaptive importance sampling methods are used to improve the GD estimator (GD-2), its variance is substantially reduced and its performance is virtually the same as the proposed estimator that is based on the cross-entropy method. In particular, all estimates of both estimators are around $-6,618$.

6.2 The Observe-Data Deviance Information Criterion

Another commonly-used model comparison criterion is the deviance information criterion (DIC) introduced in the seminal paper by Spiegelhalter, Best, Carlin, and van der

method have too many parameters, and they are unstable in this high-dimensional setting.

Linde (2002). Below we discuss how one can use the proposed conditional Monte Carlo and adaptive importance sampling methods to efficiently compute the DIC for high-dimensional stochastic volatility models.

The DIC is based on the concept of deviance, which is defined as

$$D(\boldsymbol{\theta}) = -2 \log p(\mathbf{y} | \boldsymbol{\theta}),$$

where $p(\mathbf{y} | \boldsymbol{\theta})$ is the (observed-data) likelihood.¹³ Next, the effective number of parameters, p_D , of the model is defined to be

$$p_D = \overline{D(\boldsymbol{\theta})} - D(\tilde{\boldsymbol{\theta}}),$$

where $\overline{D(\boldsymbol{\theta})} = -2\mathbb{E}_{\boldsymbol{\theta}}[\log p(\mathbf{y} | \boldsymbol{\theta}) | \mathbf{y}]$ is the posterior mean deviance and $\tilde{\boldsymbol{\theta}}$ is an estimate of $\boldsymbol{\theta}$, which is typically taken as the posterior mean or mode. Finally, the deviance information criterion is the sum of the posterior mean deviance, which can be viewed as a Bayesian measure of model fit, and the effective number of parameters that measures model complexity:

$$\text{DIC} = \overline{D(\boldsymbol{\theta})} + p_D.$$

The DIC may therefore be interpreted as a trade-off between model fit and model complexity. Given a set of competing models, the preferred model is the one with the minimum DIC value.

The main difficulty in computing the DIC for high-dimensional stochastic volatility models is that the likelihood is not available in closed-form and it is typically computationally intensive to evaluate. We show how one can use the conditional Monte Carlo and adaptive importance sampling methods developed in Section 3 to overcome this difficulty. To fix ideas, we focus on the likelihood evaluation of the VAR-FSV model; likelihood evaluation for other stochastic volatility models can be done similarly.

The latent variables in the VAR-FSV model are the factors and the log-volatility. We aim

¹³Here we distinguish between the conditional likelihood, i.e., the data density conditional on the latent variables, and the observed-data likelihood (data density marginal of the latent variables). While the DIC can be defined in terms of the conditional likelihood (Celeux, Forbes, Robert, and Titterton, 2006), a few recent papers have pointed out that this version is problematic for theoretical and practical reasons (see, e.g., Chan and Grant, 2016; Li, Yu, and Zeng, 2020). We therefore focus on the original DIC based on the observed-data likelihood.

to integrate out the factors analytically and the log-volatility using importance sampling. To that end, we stack the observations $\mathbf{y} = (\mathbf{y}'_1, \dots, \mathbf{y}'_T)'$ over $t = 1, \dots, T$ and write the VAR-FSV model as

$$\mathbf{y} = \mathbf{X}\boldsymbol{\alpha} + (\mathbf{I}_T \otimes \mathbf{L})\mathbf{f}, \quad \mathbf{u}^y \sim \mathcal{N}(\mathbf{0}, \mathbf{D}), \quad \mathbf{f} \sim \mathcal{N}(\mathbf{0}, \mathbf{G}),$$

where \mathbf{X} is an appropriately defined design matrix consisting of intercepts and lagged values, $\mathbf{D} = \text{diag}(\mathbf{D}_1, \dots, \mathbf{D}_T)$ and $\mathbf{G} = \text{diag}(\mathbf{G}_1, \dots, \mathbf{G}_T)$ with $\mathbf{D}_t = \text{diag}(e^{h_{1,t}}, \dots, e^{h_{n,t}})$ and $\mathbf{G}_t = \text{diag}(e^{h_{n+1,t}}, \dots, e^{h_{n+r,t}})$. Hence, the distribution of \mathbf{y} marginal of the factors \mathbf{f} is $(\mathbf{y} | \boldsymbol{\alpha}, \mathbf{l}, \mathbf{h}) \sim \mathcal{N}(\mathbf{X}\boldsymbol{\alpha}, \mathbf{S}_y)$ with density function

$$p(\mathbf{y} | \boldsymbol{\alpha}, \mathbf{l}, \mathbf{h}) = (2\pi)^{-\frac{Tn}{2}} |\mathbf{S}_y|^{-\frac{1}{2}} e^{-\frac{1}{2}(\mathbf{y}-\mathbf{X}\boldsymbol{\alpha})'\mathbf{S}_y^{-1}(\mathbf{y}-\mathbf{X}\boldsymbol{\alpha})},$$

where $\mathbf{S}_y = \mathbf{D} + (\mathbf{I}_T \otimes \mathbf{L})\mathbf{G}(\mathbf{I}_T \otimes \mathbf{L}')$. Finally, the likelihood function (marginal of both \mathbf{f} and \mathbf{h} but conditional on the time-invariant model parameters) can be evaluated via:

$$p(\mathbf{y} | \boldsymbol{\alpha}, \mathbf{l}, \boldsymbol{\mu}, \boldsymbol{\phi}, \boldsymbol{\sigma}^2) = \int p(\mathbf{y} | \boldsymbol{\alpha}, \mathbf{l}, \mathbf{h}) \prod_{j=1}^{n+r} p(\mathbf{h}_{j,\cdot} | \mu_j, \phi_j, \sigma_j^2) d\mathbf{h},$$

where $\mathbf{h}_{j,\cdot} = (h_{j,1}, \dots, h_{j,T})'$ and $p(\mathbf{h}_{j,\cdot} | \mu_j, \phi_j, \sigma_j^2)$ is a Gaussian density that is available analytically. To implement the importance sampling approach described in Section 3.2, we consider the parametric family

$$\mathcal{F} = \left\{ \prod_{j=1}^{n+r} f_{\mathcal{N}}(\mathbf{h}_{j,\cdot}; \hat{\mathbf{h}}_{j,\cdot}, \hat{\mathbf{K}}_{\mathbf{h}_{j,\cdot}}^{-1}) \right\}.$$

The parameters of the importance sampling densities are chosen by solving the maximization problem in (9). In particular, all the T -variate Gaussian densities are obtained using the procedure described in Section 3.4.

To illustrate the methods, we compute the observed-data DIC of the 3 stochastic volatility models with $n = 7$, and the results are reported in Table 5. In contrast to the marginal likelihood results, VAR-CSV is the most favored model according to the observed-data DIC. This is due to a combination of better model-fit (as shown by a lower posterior mean deviance value) and lower model complexity (in terms of a lower effective number of parameters). Despite being commonly used in the literature, rigorous justification of

the DIC as a model selection criterion has only been recently investigated (Li, Yu, and Zeng, 2021). More research to better understand the theoretical and empirical differences between the DIC and the marginal likelihood is therefore needed.

Table 5: The observed-data DIC, the posterior mean deviance $\overline{D(\boldsymbol{\theta})}$ and the effective number of parameters p_D of VAR-CSV, VAR-SV and VAR-FSV for $n = 7$.

	VAR-CSV	VAR-SV	VAR-FSV
DIC	3,367	3,440	3,517
$\overline{D(\boldsymbol{\theta})}$	3,181	3,195	3,239
p_D	186	245	278

6.3 A Stochastic Volatility Model with an Outlier Component

There has been a lot of recent work to develop VARs that can handle COVID-19 outliers; examples include Hartwig (2021), Ng (2021), Lenza and Primiceri (2022) and Carriero, Clark, Marcellino, and Mertens (2022). In particular, Carriero, Clark, Marcellino, and Mertens (2022) propose various stochastic volatility specifications that have an explicit component to model infrequent volatility outliers. Below we consider one of these stochastic volatility models. It extends the Cholesky stochastic volatility by incorporating a discrete mixture representation proposed in Stock and Watson (2016) for handling outliers in unobserved component models. More specifically, with the same reduced-form VAR in (1), the innovation $\boldsymbol{\varepsilon}_t$ is now distributed as:

$$\boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_t), \quad \boldsymbol{\Sigma}_t^{-1} = \mathbf{B}'_0 \mathbf{D}_t^{-1} \mathbf{B}_0 o_t^{-2}. \quad (15)$$

The new addition is the outlier parameter o_t , which is modeled to have a mixture distribution that distinguishes between regular observations with $o_t = 1$ and outliers for which $o_t > 2$. More precisely, with probability $1 - p_o$, $o_t = 1$; and with probability p_o , o_t follows a uniform distribution on $(2, 20)$: $o_t \sim \mathcal{U}(2, 20)$. The outlier probability p_o is assumed to have a beta prior, where its hyperparameters are calibrated so that the mean outlier frequency is once every 4 years in quarterly data. We refer to this model as VAR-SVO.

Following Stock and Watson (2016) and Carriero, Clark, Marcellino, and Mertens (2022), the distribution of o_t is discretized using a grid. Consequently, o_t follows a discrete distribution that can be easily handled. We refer the readers to Carriero, Clark, Marcellino, and

Mertens (2022) for estimation details. We fit VAR-SVO using the datasets ($n = 7, 15, 30$) in Section 5 and compute the corresponding marginal likelihood values. The results are reported in Table 6. For convenience, we also reproduce the marginal likelihood values of other models (the value of the best VAR-FSV for each dimension is reported).

Table 6: Log marginal likelihood estimates (numerical standard errors) of a standard homoskedastic VAR, VAR-CSV, VAR-SV, VAR-FSV and VAR-SVO.

	VAR	VAR-CSV	VAR-SV	VAR-FSV	VAR-SVO
$n = 7$	-2,583.4	-2,410	-2,315	-2,318	-2,306
	–	(0.1)	(0.3)	(0.4)	(0.6)
$n = 15$	-6,918.8	-6,618	-6,443	-6,454	-6,386
	–	(0.1)	(0.7)	(0.8)	(0.8)
$n = 30$	-12,783.4	-12,024	-11,382	-11,567	-11,255
	–	(0.1)	(1.2)	(1.8)	(1.9)

For all model dimensions, VAR-SVO is the most favored specification by the data. These results suggest that there are infrequent, extreme observations even before the COVID-19 pandemic. Moreover, they also highlight that an explicit outlier component is useful for handling these apparent outliers, even in the presence of time-varying volatility.

7 Concluding Remarks and Future Research

As large Bayesian VARs are now widely used in empirical applications, choosing the most suitable stochastic volatility specification and shrinkage priors for a particular dataset has become an increasingly pertinent problem. We took a first step to address this issue by developing Bayesian model comparison methods to select among a variety of stochastic volatility specifications and shrinkage priors in the context of large VARs. We demonstrated via a series of Monte Carlo experiments that the proposed method worked well—particularly it could be used to discriminate VARs with different stochastic volatility specifications.

Using US datasets of different dimensions, we showed that the data strongly preferred the Cholesky stochastic volatility, whereas the factor stochastic volatility was also competitive. This finding thus suggests that more future research on factor stochastic volatility

models would be fruitful, given that they are not as commonly used in empirical macroeconomics. Our results also confirmed the vital role of flexible shrinkage priors: both cross-variable shrinkage and a data-based approach to determine the overall shrinkage strength were empirically important.

In future work, it would be useful to extend the proposed estimators to compare large time-varying parameter VARs. Existing evidence seems to suggest that in a large VAR, only a few of the coefficients are time-varying. Such a model comparison method would be helpful for comparing different dynamic shrinkage priors, and thus provide useful guidelines for empirical researchers.

Appendix A: Estimation Details

In this appendix we provide the details of the priors and the estimation of the Bayesian VARs with three different stochastic volatility specifications. We also discuss the details of the adaptive importance sampling approach for estimating the marginal likelihoods for these models.

Common Stochastic Volatility

We first outline the estimation of the common stochastic volatility model given in (1)–(3). To that end, let $\mathbf{x}'_t = (1, \mathbf{y}'_{t-1}, \dots, \mathbf{y}'_{t-p})$ be a $1 \times k$ vector of an intercept and lags with $k = 1 + np$. Then, stacking the observations over $t = 1, \dots, T$, we have

$$\mathbf{Y} = \mathbf{X}\mathbf{A} + \boldsymbol{\varepsilon},$$

where $\mathbf{A} = (\mathbf{a}_0, \mathbf{A}_1, \dots, \mathbf{A}_p)'$ is $k \times n$, and the matrices \mathbf{Y} , \mathbf{X} and $\boldsymbol{\varepsilon}$ are respectively of dimensions $T \times n$, $T \times k$ and $T \times n$. Under the common stochastic volatility model, the innovations are distributed as $\text{vec}(\boldsymbol{\varepsilon}) \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma} \otimes \mathbf{D}_h)$, where $\mathbf{D}_h = \text{diag}(e^{h_1}, \dots, e^{h_T})$. The log-volatility h_t evolves as a stationary AR(1) process given in (3), which for convenience we reproduce below:

$$h_t = \phi h_{t-1} + u_t^h, \quad u_t^h \sim \mathcal{N}(0, \sigma^2),$$

for $t = 2, \dots, T$, where the process is initialized as $h_1 \sim \mathcal{N}(0, \sigma^2/(1 - \phi^2))$.

Next, we describe the priors on the model parameters \mathbf{A} , $\boldsymbol{\Sigma}$, κ , ϕ and σ^2 . First, we consider a natural conjugate prior on $(\mathbf{A}, \boldsymbol{\Sigma} \mid \kappa)$:

$$\boldsymbol{\Sigma} \sim \mathcal{IW}(\nu_0, \mathbf{S}_0), \quad (\text{vec}(\mathbf{A}) \mid \boldsymbol{\Sigma}, \kappa) \sim \mathcal{N}(\text{vec}(\mathbf{A}_0), \boldsymbol{\Sigma} \otimes \mathbf{V}_A),$$

where the prior hyperparameters $\text{vec}(\mathbf{A}_0)$ and \mathbf{V}_A are chosen in the spirit of the Minnesota prior. More specifically, we set $\mathbf{A}_0 = \mathbf{0}$ and the covariance matrix \mathbf{V}_A is assumed to be diagonal with diagonal elements $v_{A,ii} = \kappa/(l^2 \widehat{s}_r)$ for a coefficient associated to the l -th lag of variable r and $v_{A,ii} = 100$ for an intercept, where \widehat{s}_r is the sample variance of an AR(4) model for the variable r . Note that here a single hyperparameter κ controls the overall shrinkage strength and this prior does not distinguish ‘own’ versus ‘other’ lags. Here we

treat κ to be an unknown parameter with a hierarchical gamma prior: $\kappa \sim \mathcal{G}(c_1, c_2)$. Finally, for the prior distributions of ϕ and σ^2 , they are respectively truncated normal and inverse-gamma: $\phi \sim \mathcal{N}(\phi_0, V_\phi)1(|\phi| < 1)$ and $\sigma^2 \sim \mathcal{IG}(\nu_{\sigma^2}, S_{\sigma^2})$, where $1(\cdot)$ denotes the indicator function.

With the priors on the model parameters specified above, one can obtain posterior draws by sequentially sampling from:

1. $p(\mathbf{A}, \mathbf{\Sigma} | \mathbf{Y}, \mathbf{h}, \phi, \sigma^2, \kappa) = p(\mathbf{A}, \mathbf{\Sigma} | \mathbf{Y}, \mathbf{h}, \kappa)$;
2. $p(\mathbf{h} | \mathbf{Y}, \mathbf{A}, \mathbf{\Sigma}, \phi, \sigma^2, \kappa) = p(\mathbf{h} | \mathbf{Y}, \mathbf{A}, \mathbf{\Sigma}, \phi, \sigma^2)$;
3. $p(\phi | \mathbf{Y}, \mathbf{A}, \mathbf{\Sigma}, \mathbf{h}, \sigma^2, \kappa) = p(\phi | \mathbf{h}, \sigma^2)$;
4. $p(\sigma^2 | \mathbf{Y}, \mathbf{A}, \mathbf{\Sigma}, \mathbf{h}, \phi, \kappa) = p(\sigma^2 | \mathbf{h}, \phi)$;
5. $p(\kappa | \mathbf{Y}, \mathbf{A}, \mathbf{\Sigma}, \mathbf{h}, \phi, \sigma^2) = p(\kappa | \mathbf{A}, \mathbf{\Sigma})$.

Step 1: we use the results proved in Chan (2020) that $(\mathbf{A}, \mathbf{\Sigma} | \mathbf{Y}, \mathbf{h}, \kappa)$ has a normal-inverse-Wishart distribution. More specifically, let

$$\begin{aligned} \mathbf{K}_A &= \mathbf{V}_A^{-1} + \mathbf{X}'\mathbf{D}_h^{-1}\mathbf{X}, \\ \widehat{\mathbf{A}} &= \mathbf{K}_A^{-1}(\mathbf{V}_A^{-1}\mathbf{A}_0 + \mathbf{X}'\mathbf{D}_h^{-1}\mathbf{Y}), \\ \widehat{\mathbf{S}} &= \mathbf{S}_0 + \mathbf{A}_0'\mathbf{V}_A^{-1}\mathbf{A}_0 + \mathbf{Y}'\mathbf{D}_h^{-1}\mathbf{Y} - \widehat{\mathbf{A}}'\mathbf{K}_A\widehat{\mathbf{A}}. \end{aligned}$$

Then $(\mathbf{A}, \mathbf{\Sigma} | \mathbf{Y}, \mathbf{h}, \kappa)$ has a normal-inverse-Wishart distribution with parameters $\nu_0 + T$, $\widehat{\mathbf{S}}$, $\widehat{\mathbf{A}}$ and \mathbf{K}_A^{-1} . We can sample $(\mathbf{A}, \mathbf{\Sigma} | \mathbf{Y}, \mathbf{h}, \kappa)$ in two steps. First, sample $\mathbf{\Sigma}$ marginally from the inverse-Wishart distribution $(\mathbf{\Sigma} | \mathbf{Y}, \mathbf{h}) \sim \mathcal{IW}(\nu_0 + T, \widehat{\mathbf{S}})$. Then, given the sampled $\mathbf{\Sigma}$, obtain \mathbf{A} from the normal distribution:

$$(\text{vec}(\mathbf{A}) | \mathbf{Y}, \mathbf{\Sigma}, \mathbf{h}) \sim \mathcal{N}(\text{vec}(\widehat{\mathbf{A}}), \mathbf{\Sigma} \otimes \mathbf{K}_A^{-1}).$$

We note that one can sample from this normal distribution efficiently without explicitly computing the inverse \mathbf{K}_A^{-1} ; we refer the readers to Chan (2020) for computational details.

Step 2: note that

$$p(\mathbf{h} | \mathbf{Y}, \mathbf{A}, \boldsymbol{\Sigma}, \phi, \sigma^2) \propto p(\mathbf{h} | \phi, \sigma^2) \prod_{t=1}^T p(\mathbf{y}_t | \mathbf{A}, \boldsymbol{\Sigma}, h_t),$$

where $p(\mathbf{h} | \phi, \sigma^2)$ is a Gaussian density implied by the state equation (3). The log conditional likelihood $p(\mathbf{y}_t | \mathbf{A}, \boldsymbol{\Sigma}, h_t)$ has the following explicit expression:

$$\log p(\mathbf{y}_t | \mathbf{A}, \boldsymbol{\Sigma}, h_t) = c_t - \frac{n}{2} h_t - \frac{1}{2} e^{-h_t} \boldsymbol{\varepsilon}_t' \boldsymbol{\Sigma}^{-1} \boldsymbol{\varepsilon}_t,$$

where c_t is a constant independent of h_t . It is easy to check that

$$\begin{aligned} \frac{\partial}{\partial h_t} \log p(\mathbf{y}_t | \mathbf{A}, \boldsymbol{\Sigma}, h_t) &= -\frac{n}{2} + \frac{1}{2} e^{-h_t} \boldsymbol{\varepsilon}_t' \boldsymbol{\Sigma}^{-1} \boldsymbol{\varepsilon}_t, \\ \frac{\partial^2}{\partial h_t^2} \log p(\mathbf{y}_t | \mathbf{A}, \boldsymbol{\Sigma}, h_t) &= -\frac{1}{2} e^{-h_t} \boldsymbol{\varepsilon}_t' \boldsymbol{\Sigma}^{-1} \boldsymbol{\varepsilon}_t < 0. \end{aligned}$$

Given the above first and second derivatives of the log conditional likelihood, one can use the Newton-Raphson algorithm to obtain the mode of $\log p(\mathbf{h} | \mathbf{Y}, \mathbf{A}, \boldsymbol{\Sigma}, \phi, \sigma^2)$ and compute the negative Hessian evaluated at the mode, which are denoted as $\hat{\mathbf{h}}$ and $\mathbf{K}_{\mathbf{h}}$, respectively. Since the Hessian is negative definite everywhere, $\mathbf{K}_{\mathbf{h}}$ is positive definite. Next, using $\mathcal{N}(\hat{\mathbf{h}}, \mathbf{K}_{\mathbf{h}}^{-1})$ as a proposal distribution, one can sample \mathbf{h} directly using an acceptance-rejection Metropolis-Hastings step. We refer the readers to Chan (2020) for details.

Step 3: the conditional density of ϕ is of the form

$$p(\phi | \mathbf{h}, \sigma^2) \propto p(\phi) g(\phi) e^{-\frac{1}{2\sigma^2} \sum_{t=2}^T (h_t - \phi h_{t-1})^2},$$

where $g(\phi) = (1 - \phi^2)^{\frac{1}{2}} e^{-\frac{1}{2\sigma^2} (1 - \phi^2) h_1^2}$ and $p(\phi)$ is the truncated normal prior. This density is nonstandard, but we can sample ϕ via an independence-chain Metropolis-Hastings step using the proposal distribution $\mathcal{N}(\hat{\phi}, K_{\phi}^{-1}) 1(|\phi| < 1)$, where

$$K_{\phi} = V_{\phi}^{-1} + \frac{1}{\sigma^2} \sum_{t=2}^T h_{t-1}^2 \quad \hat{\phi} = K_{\phi}^{-1} \left(V_{\phi}^{-1} \phi_0 + \frac{1}{\sigma^2} \sum_{t=2}^T h_t h_{t-1} \right).$$

Then, given the current draw ϕ , a proposal ϕ^* is accepted with probability $\min(1, g(\phi^*)/g(\phi))$;

otherwise the Markov chain stays at the current state ϕ .

Step 4: σ^2 can be sampled easily as it has the following inverse-gamma distribution:

$$(\sigma^2 | \mathbf{h}, \phi) \sim \mathcal{IG}(\nu_{\sigma^2} + T/2, \tilde{S}_{\sigma^2}),$$

where $\tilde{S}_{\sigma^2} = S_{\sigma^2} + \left[(1 - \phi^2)h_1^2 + \sum_{t=2}^T (h_t - \phi h_{t-1})^2 \right] / 2$.

Step 5: first note that κ only appears in its gamma prior $\kappa \sim \mathcal{G}(c_1, c_2)$ and $\mathbf{V}_{\mathbf{A}}$, the prior covariance matrix of \mathbf{A} . Recall that $\mathbf{V}_{\mathbf{A}}$ is a $k \times k$ diagonal matrix in which the first element—corresponding to the prior variance of the intercept—does not involve κ . More explicitly, for $i = 2, \dots, k$, write the i -th diagonal element of $\mathbf{V}_{\mathbf{A}}$ as $v_{\mathbf{A},ii} = \kappa C_i$ for some constant C_i . Then, we have

$$\begin{aligned} p(\kappa | \mathbf{A}, \boldsymbol{\Sigma}) &\propto \kappa^{c_1-1} e^{-c_2 \kappa} \times |\mathbf{V}_{\mathbf{A}}|^{-\frac{n}{2}} e^{-\frac{1}{2} \text{tr}(\boldsymbol{\Sigma}^{-1}(\mathbf{A}-\mathbf{A}_0)' \mathbf{V}_{\mathbf{A}}^{-1}(\mathbf{A}-\mathbf{A}_0))} \\ &\propto \kappa^{c_1 - \frac{(k-1)n}{2} - 1} e^{-c_2 \kappa} e^{-\frac{1}{2} \text{tr}(\mathbf{V}_{\mathbf{A}}^{-1}(\mathbf{A}-\mathbf{A}_0) \boldsymbol{\Sigma}^{-1}(\mathbf{A}-\mathbf{A}_0)')} \\ &\propto \kappa^{c_1 - \frac{(k-1)n}{2} - 1} e^{-\frac{1}{2} (2c_2 \kappa + \kappa^{-1} \sum_{i=2}^k Q_i / C_i)}, \end{aligned}$$

where Q_i is the i -th diagonal element of $\mathbf{Q} = (\mathbf{A} - \mathbf{A}_0) \boldsymbol{\Sigma}^{-1} (\mathbf{A} - \mathbf{A}_0)'$. Note that this is the kernel of the generalized inverse Gaussian distribution $\mathcal{GIG} \left(c_1 - \frac{(k-1)n}{2}, 2c_2, \sum_{i=2}^k Q_i / C_i \right)$. Draws from the generalized inverse Gaussian distribution can be obtained using the algorithm in Devroye (2014).

Next, we derive the expression of $p(\mathbf{Y} | \mathbf{h}, \kappa)$ given in (10). First let k_1 denote the normalizing constant of the normal-inverse-Wishart prior: $k_1 = (2\pi)^{-\frac{nk}{2}} 2^{-\frac{n\nu_0}{2}} |\mathbf{V}_{\mathbf{A}}|^{-\frac{n}{2}} \Gamma_n(\nu_0/2)^{-1} |\mathbf{S}_0|^{\frac{\nu_0}{2}}$. Then, by direct computation:

$$\begin{aligned} p(\mathbf{Y} | \mathbf{h}, \kappa) &= \int p(\mathbf{Y} | \mathbf{A}, \boldsymbol{\Sigma}, \mathbf{h}) p(\mathbf{A}, \boldsymbol{\Sigma} | \kappa) d(\mathbf{A}, \boldsymbol{\Sigma}) \\ &= \int (2\pi)^{-\frac{Tn}{2}} |\boldsymbol{\Sigma}|^{-\frac{T}{2}} e^{-\frac{n}{2} \mathbf{1}'_T \mathbf{h}} e^{-\frac{1}{2} \text{tr}(\boldsymbol{\Sigma}^{-1}(\mathbf{Y}-\mathbf{X}\mathbf{A})' \mathbf{D}_{\mathbf{h}}^{-1}(\mathbf{Y}-\mathbf{X}\mathbf{A}))} \\ &\quad \times k_1 |\boldsymbol{\Sigma}|^{-\frac{\nu_0+n+k+1}{2}} e^{-\frac{1}{2} \text{tr}(\boldsymbol{\Sigma}^{-1} \mathbf{S}_0)} e^{-\frac{1}{2} \text{tr}(\boldsymbol{\Sigma}^{-1}(\mathbf{A}-\mathbf{A}_0)' \mathbf{V}_{\mathbf{A}}^{-1}(\mathbf{A}-\mathbf{A}_0))} d(\mathbf{A}, \boldsymbol{\Sigma}) \\ &= k_1 (2\pi)^{-\frac{Tn}{2}} e^{-\frac{n}{2} \mathbf{1}'_T \mathbf{h}} \int |\boldsymbol{\Sigma}|^{-\frac{\nu_0+T+n+k+1}{2}} e^{-\frac{1}{2} \text{tr}(\boldsymbol{\Sigma}^{-1} \hat{\mathbf{S}})} e^{-\frac{1}{2} \text{tr}(\boldsymbol{\Sigma}^{-1}(\mathbf{A}-\hat{\mathbf{A}})' \mathbf{K}_{\mathbf{A}}^{-1}(\mathbf{A}-\hat{\mathbf{A}}))} d(\mathbf{A}, \boldsymbol{\Sigma}) \\ &= \pi^{-\frac{Tn}{2}} e^{-\frac{n}{2} \mathbf{1}'_T \mathbf{h}} |\mathbf{V}_{\mathbf{A}}|^{-\frac{n}{2}} |\mathbf{K}_{\mathbf{A}}|^{-\frac{n}{2}} \frac{\Gamma_n \left(\frac{\nu_0+T}{2} \right) |\mathbf{S}_0|^{\frac{\nu_0}{2}}}{\Gamma_n \left(\frac{\nu_0}{2} \right) |\hat{\mathbf{S}}|^{\frac{\nu_0+T}{2}}}, \end{aligned}$$

where the shrinkage hyperparameter κ appears in $\mathbf{V}_{\mathbf{A}}$, and the last equality holds because

$$\begin{aligned} & \int |\boldsymbol{\Sigma}|^{-\frac{\nu_0+T+n+k+1}{2}} e^{-\frac{1}{2}\text{tr}(\boldsymbol{\Sigma}^{-1}\widehat{\mathbf{S}})} e^{-\frac{1}{2}\text{tr}(\boldsymbol{\Sigma}^{-1}(\mathbf{A}-\widehat{\mathbf{A}})'\mathbf{K}_{\mathbf{A}}^{-1}(\mathbf{A}-\widehat{\mathbf{A}}))} d(\mathbf{A}, \boldsymbol{\Sigma}) \\ &= (2\pi)^{\frac{nk}{2}} 2^{\frac{n(\nu_0+T)}{2}} |\mathbf{K}_{\mathbf{A}}^{-1}|^{\frac{n}{2}} \Gamma_n\left(\frac{\nu_0+T}{2}\right) |\widehat{\mathbf{S}}|^{-\frac{\nu_0+T}{2}}. \end{aligned}$$

Finally, the marginal density $p(\mathbf{h} | \phi)$ has the following analytical expression:

$$\begin{aligned} p(\mathbf{h} | \phi) &= \int p(\mathbf{h} | \phi, \sigma^2) p(\sigma^2) d\sigma^2 \\ &= (2\pi)^{-\frac{T}{2}} |1 - \phi^2|^{\frac{1}{2}} S_{\sigma^2}^{\nu_{\sigma^2}} \Gamma(\nu_{\sigma^2})^{-1} \int (\sigma^2)^{-(\nu_{\sigma^2} + \frac{T}{2} + 1)} e^{-\frac{1}{\sigma^2} \widetilde{S}_{\sigma^2}} d\sigma^2 \\ &= (2\pi)^{-\frac{T}{2}} |1 - \phi^2|^{\frac{1}{2}} \frac{\Gamma(\nu_{\sigma^2} + \frac{T}{2}) S_{\sigma^2}^{\nu_{\sigma^2}}}{\Gamma(\nu_{\sigma^2}) \widetilde{S}_{\sigma^2}^{\nu_{\sigma^2} + \frac{T}{2}}}, \end{aligned}$$

where $\Gamma(\cdot)$ is the gamma function and \widetilde{S}_{σ^2} is defined in Step 4 of the MCMC algorithm.

Cholesky Stochastic Volatility

Next, we outline the estimation of the VAR-SV model:

$$\mathbf{y}_t = \mathbf{a}_0 + \mathbf{A}_1 \mathbf{y}_{t-1} + \cdots + \mathbf{A}_p \mathbf{y}_{t-p} + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_t), \quad (16)$$

where $\boldsymbol{\Sigma}_t^{-1} = \mathbf{B}'_0 \mathbf{D}_t^{-1} \mathbf{B}_0$, $\mathbf{D}_t = \text{diag}(e^{h_{1,t}}, \dots, e^{h_{n,t}})$ and \mathbf{B}_0 is an $n \times n$ lower triangular matrix with ones on the diagonal. Each element of $\mathbf{h}_t = (h_{1,t}, \dots, h_{n,t})'$ follows an autoregressive process:

$$h_{i,t} = \mu_i + \phi_i (h_{i,t-1} - \mu_i) + u_{i,t}^h, \quad u_{i,t}^h \sim \mathcal{N}(0, \sigma_i^2)$$

for $t = 2, \dots, T$, and the initial condition is specified as $h_{i,1} \sim \mathcal{N}(\mu_i, \sigma_i^2 / (1 - \phi_i^2))$.

We aim to estimate the system equation by equation, which substantially speeds up the computation time. Recall that $\boldsymbol{\alpha}_i$ denotes the $k \times 1$ vector that consists of the intercept and VAR coefficients in the i -th equation, and $\boldsymbol{\beta}_i$ represents the $(i-1) \times 1$ vector of free elements in the i -th row of the impact matrix \mathbf{B}_0 . Then, the parameters for the i -th equation are $\boldsymbol{\alpha}_i, \boldsymbol{\beta}_i, \mu_i, \phi_i$, and σ_i^2 . We assume the following independent prior

distributions on the parameters:

$$\begin{aligned}
(\boldsymbol{\alpha}_i | \boldsymbol{\kappa}) &\sim \mathcal{N}(\boldsymbol{\alpha}_{0,i}, \mathbf{V}_{\boldsymbol{\alpha}_i}), \quad (\boldsymbol{\beta}_i | \boldsymbol{\kappa}) \sim \mathcal{N}(\boldsymbol{\beta}_{0,i}, \mathbf{V}_{\boldsymbol{\beta}_i}), \\
\mu_i &\sim \mathcal{N}(\mu_{0,i}, V_{\mu_i}), \quad \phi_i \sim \mathcal{N}(\phi_{0,i}, V_{\phi_i})1(|\phi_i| < 1), \quad \sigma_i^2 \sim \mathcal{IG}(\nu_i, S_i),
\end{aligned} \tag{17}$$

where $\boldsymbol{\kappa}$ is a vector of hyperparameters that is described in more detail below. The prior mean vector $\boldsymbol{\alpha}_{0,i}$ and the prior covariance matrix $\mathbf{V}_{\boldsymbol{\alpha}_i}$ are selected to mimic the Minnesota prior. More specifically, we set $\boldsymbol{\alpha}_{0,i} = \mathbf{0}$ to shrink the VAR coefficients to zero. For $\mathbf{V}_{\boldsymbol{\alpha}_i}$, we assume it to be diagonal with the k -th diagonal element $V_{\boldsymbol{\alpha}_i, kk}$ set to be:

$$V_{\boldsymbol{\alpha}_i, kk} = \begin{cases} \frac{\kappa_1}{l^2}, & \text{for the coefficient on the } l\text{-th lag of variable } i, \\ \frac{\kappa_2 s_i^2}{l^2 s_j^2}, & \text{for the coefficient on the } l\text{-th lag of variable } j, j \neq i, \\ 100s_i^2, & \text{for the intercept,} \end{cases}$$

where s_r^2 denotes the sample variance of the residuals from an AR(4) model for the variable r for $r = 1, \dots, n$. In addition, we set the prior mean vector $\boldsymbol{\beta}_{0,i}$ to be zero to shrink the impact matrix to the identity matrix. In addition, the prior covariance matrix $\mathbf{V}_{\boldsymbol{\beta}_i}$ is assumed to be diagonal, where the j -th diagonal element is set to be $\kappa_3 s_i^2 / s_j^2$. Finally, we treat the shrinkage hyperparameters $\boldsymbol{\kappa} = (\kappa_1, \kappa_2, \kappa_3)'$ as unknown parameters to be estimated with hierarchical gamma priors $\kappa_i \sim \mathcal{G}(c_{j,1}, c_{j,2}), j = 1, 2, 3$.

Let $\mathbf{y}_{i,\cdot} = (y_{i,1}, \dots, y_{i,T})'$ denote the vector of observed values for the i -th variable for $i = 1, \dots, n$. We similarly define $\mathbf{h}_{i,\cdot} = (h_{i,1}, \dots, h_{i,T})'$. Next, stack $\mathbf{y} = (\mathbf{y}'_{1,\cdot}, \dots, \mathbf{y}'_{n,\cdot})'$, $\mathbf{h} = (\mathbf{h}'_{1,\cdot}, \dots, \mathbf{h}'_{n,\cdot})'$ and $\boldsymbol{\alpha} = (\boldsymbol{\alpha}'_1, \dots, \boldsymbol{\alpha}'_n)'$; similarly define $\boldsymbol{\beta}, \boldsymbol{\mu}, \boldsymbol{\phi}$ and $\boldsymbol{\sigma}^2$. Then, posterior draws can be obtained by sampling sequentially from:

1. $p(\boldsymbol{\alpha} | \mathbf{y}, \boldsymbol{\beta}, \mathbf{h}, \boldsymbol{\mu}, \boldsymbol{\phi}, \boldsymbol{\sigma}^2, \boldsymbol{\kappa})$;
2. $p(\boldsymbol{\beta} | \mathbf{y}, \boldsymbol{\alpha}, \mathbf{h}, \boldsymbol{\mu}, \boldsymbol{\phi}, \boldsymbol{\sigma}^2, \boldsymbol{\kappa}) = \prod_{i=2}^n p(\boldsymbol{\beta}_i | \mathbf{y}, \boldsymbol{\alpha}, \mathbf{h}_{i,\cdot}, \boldsymbol{\kappa})$;
3. $p(\mathbf{h} | \mathbf{y}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\mu}, \boldsymbol{\phi}, \boldsymbol{\sigma}^2, \boldsymbol{\kappa}) = \prod_{i=1}^n p(\mathbf{h}_{i,\cdot} | \mathbf{y}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mu_i, \phi_i, \sigma_i^2)$;
4. $p(\boldsymbol{\mu} | \mathbf{y}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{h}, \boldsymbol{\phi}, \boldsymbol{\sigma}^2, \boldsymbol{\kappa}) = \prod_{i=1}^n p(\mu_i | \mathbf{h}_{i,\cdot}, \phi_i, \sigma_i^2)$;
5. $p(\boldsymbol{\phi} | \mathbf{y}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{h}, \boldsymbol{\mu}, \boldsymbol{\sigma}^2, \boldsymbol{\kappa}) = \prod_{i=1}^n p(\phi_i | \mathbf{h}_{i,\cdot}, \mu_i, \sigma_i^2)$;
6. $p(\boldsymbol{\sigma}^2 | \mathbf{y}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{h}, \boldsymbol{\mu}, \boldsymbol{\phi}, \boldsymbol{\kappa}) = \prod_{i=1}^n p(\sigma_i^2 | \mathbf{h}_{i,\cdot}, \mu_i, \phi_i)$;

$$7. p(\boldsymbol{\kappa} | \mathbf{y}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{h}, \boldsymbol{\mu}, \boldsymbol{\phi}, \boldsymbol{\sigma}^2) = p(\boldsymbol{\kappa} | \boldsymbol{\alpha}, \boldsymbol{\beta}).$$

Step 1: we sample the reduced-form VAR coefficients equation by equation using the triangular algorithm proposed in Carriero, Chan, Clark, and Marcellino (2022). To that end, define $\mathbf{A}_{i=0}$ to be a $k \times n$ matrix that has exactly the same elements as $\mathbf{A} = (\mathbf{a}_0, \mathbf{A}_1, \dots, \mathbf{A}_p)'$ except for the i -th column, which is set to be zero, i.e., $\mathbf{A}_{i=0} = (\boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_{i-1}, \mathbf{0}, \boldsymbol{\alpha}_{i+1}, \dots, \boldsymbol{\alpha}_n)$. Then, we can rewrite (16) as

$$\mathbf{B}_0(\mathbf{y}_t - \mathbf{A}'_{i=0}\mathbf{x}_t) = \mathbf{B}_{0,1:n,i}\boldsymbol{\alpha}'_i\mathbf{x}_t + \tilde{\boldsymbol{\varepsilon}}_t, \quad \tilde{\boldsymbol{\varepsilon}}_t \equiv \mathbf{B}_0\boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{D}_t),$$

where $\mathbf{x}_t = (1, \mathbf{y}'_{t-1}, \dots, \mathbf{y}'_{t-p})'$ and $\mathbf{B}_{0,1:n,i}$ is the i -th column of \mathbf{B}_0 . Transposing and stacking the above equation over $t = 1, \dots, T$, we can write the system more compactly as

$$(\mathbf{Y} - \mathbf{X}\mathbf{A}_{i=0})\mathbf{B}'_0 = \mathbf{X}\boldsymbol{\alpha}_i\mathbf{B}'_{0,1:n,i} + \tilde{\boldsymbol{\varepsilon}},$$

where the matrices \mathbf{Y} , \mathbf{X} and $\tilde{\boldsymbol{\varepsilon}}$ are respectively of dimensions $T \times n$, $T \times k$ and $T \times n$. Vectoring the above system, we obtain

$$\mathbf{z}^i = \mathbf{W}^i\boldsymbol{\alpha}_i + \text{vec}(\tilde{\boldsymbol{\varepsilon}}), \quad \text{vec}(\tilde{\boldsymbol{\varepsilon}}) \sim \mathcal{N}(\mathbf{0}, \mathbf{D}),$$

where $\mathbf{z}^i = \text{vec}((\mathbf{Y} - \mathbf{X}\mathbf{A}_{i=0})\mathbf{B}'_0)$, $\mathbf{W}^i = \mathbf{B}_{0,1:n,i} \otimes \mathbf{X}$ and $\mathbf{D} = \text{diag}(\mathbf{D}_{1,\cdot}, \dots, \mathbf{D}_{n,\cdot})$ with $\mathbf{D}_{i,\cdot} = \text{diag}(e^{h_{i,1}}, \dots, e^{h_{i,T}})$. Given the conditionally Gaussian prior on $\boldsymbol{\alpha}_i$ in (17) (conditional on $\boldsymbol{\kappa}$), it follows from standard linear regression results that

$$(\boldsymbol{\alpha}_i | \mathbf{y}, \boldsymbol{\beta}, \boldsymbol{\alpha}_{-i}, \mathbf{h}, \boldsymbol{\mu}, \boldsymbol{\phi}, \boldsymbol{\sigma}^2, \boldsymbol{\kappa}) \sim \mathcal{N}(\hat{\boldsymbol{\alpha}}_i, \mathbf{K}_{\boldsymbol{\alpha}_i}^{-1}),$$

where $\boldsymbol{\alpha}_{-i} = (\boldsymbol{\alpha}'_1, \dots, \boldsymbol{\alpha}'_{i-1}, \boldsymbol{\alpha}'_{i+1}, \dots, \boldsymbol{\alpha}'_n)'$,

$$\mathbf{K}_{\boldsymbol{\alpha}_i} = \mathbf{V}_{\boldsymbol{\alpha}_i}^{-1} + \mathbf{W}^{i'}\mathbf{D}^{-1}\mathbf{W}^i, \quad \hat{\boldsymbol{\alpha}}_i = \mathbf{K}_{\boldsymbol{\alpha}_i}^{-1} \left(\mathbf{V}_{\boldsymbol{\alpha}_i}^{-1}\boldsymbol{\alpha}_{0,i} + \mathbf{W}^{i'}\mathbf{D}^{-1}\mathbf{z}^i \right).$$

Some of the above computations can be simplified by noting that \mathbf{B}_0 is a lower triangular matrix. In addition, random draws from the high-dimensional $\mathcal{N}(\hat{\boldsymbol{\alpha}}_i, \mathbf{K}_{\boldsymbol{\alpha}_i}^{-1})$ distribution can be obtained efficiently without inverting any large matrices; see, e.g., Chan (2021) for computational details.

Step 2: to sample $\boldsymbol{\beta}$, the lower triangular elements of \mathbf{B}_0 , first note that given \mathbf{y} and $\boldsymbol{\alpha}$, the innovations $\boldsymbol{\varepsilon}_1, \dots, \boldsymbol{\varepsilon}_T$ can be computed using (16). Next, we rewrite the model

as a system of $(n - 1)$ regressions in which $\varepsilon_{i,t}$ is regressed on the negative values of $\varepsilon_{1,t}, \dots, \varepsilon_{i-1,t}$ for $i = 2, \dots, n$, and $B_{0,i,1}, \dots, B_{0,i,i-1}$ are the corresponding regression coefficients. More specifically, note that

$$\begin{aligned} \mathbf{B}_0 \boldsymbol{\varepsilon}_t &= \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ B_{0,2,1} & 1 & 0 & \cdots & 0 \\ B_{0,3,1} & B_{0,3,2} & 1 & \cdots & \vdots \\ \vdots & \vdots & & \ddots & \vdots \\ B_{0,n,1} & B_{0,n,2} & B_{0,n,3} & \cdots & 1 \end{pmatrix} \begin{pmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \\ \varepsilon_{3,t} \\ \vdots \\ \varepsilon_{n,t} \end{pmatrix} = \begin{pmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} + B_{0,2,1} \varepsilon_{1,t} \\ \varepsilon_{3,t} + B_{0,3,1} \varepsilon_{1,t} + B_{0,3,2} \varepsilon_{2,t} \\ \vdots \\ \varepsilon_{n,t} + \sum_{j=1}^{n-1} B_{0,n,j} \varepsilon_{j,t} \end{pmatrix} \\ &= \begin{pmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \\ \varepsilon_{3,t} \\ \vdots \\ \varepsilon_{n,t} \end{pmatrix} - \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & \cdots & \cdots & 0 \\ -\varepsilon_{1,t} & 0 & 0 & 0 & 0 & \cdots & & \vdots \\ 0 & -\varepsilon_{1,t} & -\varepsilon_{2,t} & 0 & 0 & \cdots & & 0 \\ \vdots & \vdots & & \ddots & \ddots & & & 0 \\ 0 & \cdots & 0 & \cdots & 0 & -\varepsilon_{1,t} & \cdots & -\varepsilon_{t,n-1} \end{pmatrix} \begin{pmatrix} B_{0,2,1} \\ B_{0,3,1} \\ B_{0,3,2} \\ \vdots \\ B_{0,n,n-1} \end{pmatrix}. \end{aligned}$$

Or more succinctly,

$$\mathbf{B}_0 \boldsymbol{\varepsilon}_t = \boldsymbol{\varepsilon}_t - \mathbf{E}_t \boldsymbol{\beta}.$$

Since $|\boldsymbol{\Sigma}_t| = |\mathbf{D}_t|$, we can rewrite the likelihood implied by (16) as

$$\begin{aligned} p(\mathbf{y} | \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{h}) &\propto \left(\prod_{t=1}^T |\mathbf{D}_t|^{-\frac{1}{2}} \right) \exp \left(-\frac{1}{2} \sum_{t=1}^T \boldsymbol{\varepsilon}'_t (\mathbf{B}'_0 \mathbf{D}_t^{-1} \mathbf{B}_0) \boldsymbol{\varepsilon}_t \right) \\ &= \left(\prod_{t=1}^T |\mathbf{D}_t|^{-\frac{1}{2}} \right) \exp \left(-\frac{1}{2} \sum_{t=1}^T (\boldsymbol{\varepsilon}_t - \mathbf{E}_t \boldsymbol{\beta})' \mathbf{D}_t^{-1} (\boldsymbol{\varepsilon}_t - \mathbf{E}_t \boldsymbol{\beta}) \right). \end{aligned}$$

That is, the likelihood is the same as that implied by the regression

$$\boldsymbol{\varepsilon}_t = \mathbf{E}_t \boldsymbol{\beta} + \boldsymbol{\eta}_t, \tag{18}$$

where $\boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{D}_t)$. In addition, since the covariance matrix \mathbf{D}_t is diagonal and \mathbf{E}_t is block-diagonal, we can sample $\boldsymbol{\beta}_2, \dots, \boldsymbol{\beta}_n$ equation by equation to improve computational efficiency.

To that end, we stack the i -th equation, $i = 2, \dots, n$, over $t = 1, \dots, T$ and obtain

$$\boldsymbol{\varepsilon}_{i,\cdot} = \mathbf{E}_{i,\cdot} \boldsymbol{\beta}_i + \boldsymbol{\eta}_{i,\cdot},$$

where $\boldsymbol{\varepsilon}_{i,\cdot} = (\varepsilon_{i,1}, \dots, \varepsilon_{i,T})'$, $\mathbf{E}_{i,\cdot} = (-\varepsilon_{1,\cdot}, \dots, -\varepsilon_{i-1,\cdot})$ and $\boldsymbol{\eta}_{i,\cdot} \sim \mathcal{N}(\mathbf{0}, \mathbf{D}_{i,\cdot})$ with $\mathbf{D}_{i,\cdot} = \text{diag}(e^{h_{i,1}}, \dots, e^{h_{i,T}})$. Finally, given the hierarchical Gaussian prior $(\boldsymbol{\beta}_i | \boldsymbol{\kappa}) \sim \mathcal{N}(\boldsymbol{\beta}_{0,i}, \mathbf{V}_{\boldsymbol{\beta}_i})$, it follows that

$$(\boldsymbol{\beta}_i | \mathbf{y}, \boldsymbol{\alpha}, \mathbf{h}_{i,\cdot}, \boldsymbol{\kappa}) \sim \mathcal{N}(\widehat{\boldsymbol{\beta}}_i, \mathbf{K}_{\boldsymbol{\beta}_i}^{-1}),$$

where

$$\mathbf{K}_{\boldsymbol{\beta}_i} = \mathbf{V}_{\boldsymbol{\beta}_i}^{-1} + \mathbf{E}'_{i,\cdot} \mathbf{D}_{i,\cdot}^{-1} \mathbf{E}_{i,\cdot}, \quad \widehat{\boldsymbol{\beta}}_i = \mathbf{K}_{\boldsymbol{\beta}_i}^{-1} \left(\mathbf{V}_{\boldsymbol{\beta}_i}^{-1} \boldsymbol{\beta}_{0,i} + \mathbf{E}'_{i,\cdot} \mathbf{D}_{i,\cdot}^{-1} \boldsymbol{\varepsilon}_{i,\cdot} \right).$$

Step 3: to sample \mathbf{h} , we first compute the ‘‘orthogonalized’’ innovations: $\tilde{\boldsymbol{\varepsilon}}_t = \mathbf{B}_0(\mathbf{y}_t - \mathbf{a}_0 - \mathbf{A}_1 \mathbf{y}_{t-1} - \dots - \mathbf{A}_p \mathbf{y}_{t-p})$ for $t = 1, \dots, T$. Then, we can sample each vector $\mathbf{h}_{i,\cdot}$ separately for $i = 1, \dots, n$. More specifically, using the time series $\tilde{\varepsilon}_{i,1}, \dots, \tilde{\varepsilon}_{i,T}$ as raw data, we directly apply the auxiliary mixture sampler of Kim, Shephard, and Chib (1998) in conjunction with the precision sampler of Chan and Jeliaskov (2009) to sample $(\mathbf{h}_{i,\cdot} | \mathbf{y}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mu_i, \phi_i, \sigma_i^2)$ for $i = 1, \dots, n$. For a textbook treatment, see, e.g., Chapter 19 in Chan, Koop, Poirier, and Tobias (2019).

Step 4: this step can be done easily, as μ_1, \dots, μ_n are conditionally independent given \mathbf{h} and other parameters, and each follows a normal distribution:

$$(\mu_i | \mathbf{h}_{i,\cdot}, \phi_i, \sigma_i^2) \sim \mathcal{N}(\widehat{\mu}_i, K_{\mu_i}^{-1}),$$

where

$$K_{\mu_i} = V_{\mu_i}^{-1} + \frac{1}{\sigma_i^2} [1 - \phi_i^2 + (T-1)(1 - \phi_i)^2]$$

$$\widehat{\mu}_i = K_{\mu_i}^{-1} \left[V_{\mu_i}^{-1} \mu_{0,i} + \frac{1}{\sigma_i^2} \left((1 - \phi_i^2) h_{i,1} + (1 - \phi_i) \sum_{t=2}^T (h_{i,t} - \phi_i h_{i,t-1}) \right) \right].$$

Step 5: to sample $\boldsymbol{\phi}$, first note that

$$p(\boldsymbol{\phi}_i | \mathbf{h}_{i,\cdot}, \mu_i, \sigma_i^2) \propto p(\phi_i) g(\phi_i) e^{-\frac{1}{2\sigma_i^2} \sum_{t=2}^T (h_{i,t} - \mu_i - \phi_i (h_{i,t-1} - \mu_i))^2},$$

where $g(\phi_i) = (1 - \phi_i^2)^{1/2} e^{-\frac{1}{2\sigma_i^2} (1 - \phi_i^2) (h_{i,1} - \mu_i)^2}$ and $p(\phi_i)$ is the truncated normal prior given in (17). The conditional density $p(\phi_i | \mathbf{h}_{i,\cdot}, \mu_i, \sigma_i^2)$ is nonstandard, but a draw from it can be obtained by using an independence-chain Metropolis-Hastings step with proposal

distribution $\mathcal{N}(\widehat{\phi}_i, K_{\phi_i})1(|\phi_i| < 1)$, where

$$K_{\phi_i} = V_{\phi_i}^{-1} + \frac{1}{\sigma_i^2} \sum_{t=2}^T (h_{i,t-1} - \mu_i)^2$$

$$\widehat{\phi}_i = K_{\phi_i}^{-1} \left[V_{\phi_i}^{-1} \phi_{0,i} + \frac{1}{\sigma_i^2} \sum_{t=2}^T (h_{i,t-1} - \mu_i)(h_{i,t} - \mu_i) \right].$$

Then, given the current draw ϕ_i , a proposal ϕ_i^* is accepted with probability $\min(1, g(\phi_i^*)/g(\phi_i))$; otherwise the Markov chain stays at the current state ϕ_i .

Step 6: to sample $\sigma_1^2, \dots, \sigma_n^2$, note that each follows an inverse-gamma distribution:

$$(\sigma_i^2 \mid \mathbf{h}_{i,\cdot}, \mu_i, \phi_i) \sim \mathcal{IG} \left(\nu_i + \frac{T}{2}, \widetilde{S}_i \right),$$

where $\widetilde{S}_i = S_i + [(1 - \phi_i^2)(h_{i,1} - \mu_i)^2 + \sum_{t=2}^T (h_{i,t} - \mu_i - \phi_i(h_{i,t-1} - \mu_i))^2]/2$.

Step 7: note that κ_1, κ_2 and κ_3 only appear in their priors $\kappa_j \sim \mathcal{G}(c_{j,1}, c_{j,2}), j = 1, 2, 3$, and the prior covariance matrices of $\boldsymbol{\alpha}_i$ and $\boldsymbol{\beta}_i$ in (17). To sample κ_1, κ_2 and κ_3 , first define the index set S_{κ_1} that collects all the indexes (i, j) such that $\alpha_{i,j}$ is a coefficient associated with an own lag. That is, $S_{\kappa_1} = \{(i, j) : \alpha_{i,j} \text{ is a coefficient associated with an own lag}\}$. Similarly, define S_{κ_2} as the set that collects all the indexes (i, j) such that $\alpha_{i,j}$ is a coefficient associated with a lag of other variables. Lastly, define $S_{\kappa_3} = \{(i, j) : i = 2, \dots, n, j = 1, \dots, i - 1\}$. It is easy to check that the numbers of elements in $S_{\kappa_1}, S_{\kappa_2}$ and S_{κ_3} are, respectively, $np, (n - 1)np$ and $n(n - 1)/2$. Then, we have

$$p(\kappa_1 \mid \boldsymbol{\alpha}) \propto \prod_{(i,j) \in S_{\kappa_1}} \kappa_1^{-\frac{1}{2}} e^{-\frac{1}{2\kappa_1 C_{i,j}} (\alpha_{i,j} - \alpha_{0,i,j})^2} \times \kappa_1^{c_{1,1} - 1} e^{-\kappa_1 c_{1,2}}$$

$$= \kappa_1^{c_{1,1} - \frac{np}{2} - 1} e^{-\frac{1}{2} \left(2c_{1,2}\kappa_1 + \kappa_1^{-1} \sum_{(i,j) \in S_{\kappa_1}} \frac{(\alpha_{i,j} - \alpha_{0,i,j})^2}{C_{i,j}} \right)},$$

where $\alpha_{0,i,j}$ is the j -th element of the prior mean vector $\boldsymbol{\alpha}_{0,i}$ and $C_{i,j}$ is a constant determined by the Minnesota prior. Note that the above expression is the kernel of the

$\mathcal{GIG}\left(c_{1,1} - \frac{np}{2}, 2c_{1,2}, \sum_{(i,j) \in S_{\kappa_1}} \frac{(\alpha_{i,j} - \alpha_{0,i,j})^2}{C_{i,j}}\right)$ distribution. Similarly, we have

$$\begin{aligned} (\kappa_2 | \boldsymbol{\alpha}) &\sim \mathcal{GIG}\left(c_{2,1} - \frac{(n-1)np}{2}, 2c_{2,2}, \sum_{(i,j) \in S_{\kappa_2}} \frac{(\alpha_{i,j} - \alpha_{0,i,j})^2}{C_{i,j}}\right), \\ (\kappa_3 | \boldsymbol{\beta}) &\sim \mathcal{GIG}\left(c_{3,1} - \frac{n(n-1)}{4}, 2c_{3,2}, \sum_{(i,j) \in S_{\kappa_3}} \frac{(\beta_{i,j} - \beta_{0,i,j})^2}{\tilde{C}_{i,j}}\right), \end{aligned}$$

where $\tilde{C}_{i,j} = s_i^2/s_j^2$ and s_r^2 is the sample variance of the residuals from an AR(4) model for variable r .

Next, we provide the details of estimating the marginal likelihood of the VAR-SV model. As before, our marginal likelihood estimator of $p(\mathbf{y})$ has two parts: the conditional Monte Carlo part in which we integrate out the VAR coefficients $\boldsymbol{\alpha}$; and the adaptive importance sampling part that biases the joint distribution of $\boldsymbol{\beta}, \mathbf{h}, \boldsymbol{\mu}, \boldsymbol{\phi}, \boldsymbol{\sigma}^2$ and $\boldsymbol{\kappa}$. In what follows, we first derive an analytical expression of the conditional Monte Carlo estimator $\mathbb{E}[p(\mathbf{y} | \boldsymbol{\theta}, \mathbf{h}, \boldsymbol{\kappa}) | \boldsymbol{\beta}, \mathbf{h}, \boldsymbol{\kappa}] = p(\mathbf{y} | \boldsymbol{\beta}, \mathbf{h}, \boldsymbol{\kappa})$. To that end, rewrite (16) as

$$\mathbf{Y}\mathbf{B}'_0 = \mathbf{X}\mathbf{A}\mathbf{B}'_0 + \tilde{\boldsymbol{\varepsilon}},$$

where $\mathbf{A} = (\mathbf{a}_0, \mathbf{A}_1, \dots, \mathbf{A}_p)'$, $\tilde{\boldsymbol{\varepsilon}} = \boldsymbol{\varepsilon}\mathbf{B}'_0$ and the matrices \mathbf{Y} , \mathbf{X} and $\boldsymbol{\varepsilon}$ are, respectively, of dimensions $T \times n$, $T \times k$ and $T \times n$. Vectoring the above system, we obtain

$$\tilde{\mathbf{y}} = \tilde{\mathbf{X}}\boldsymbol{\alpha} + \text{vec}(\tilde{\boldsymbol{\varepsilon}}), \quad \text{vec}(\tilde{\boldsymbol{\varepsilon}}) \sim \mathcal{N}(\mathbf{0}, \mathbf{D}),$$

where $\tilde{\mathbf{y}} = \text{vec}(\mathbf{Y}\mathbf{B}'_0)$, $\tilde{\mathbf{X}} = \mathbf{B}_0 \otimes \mathbf{X}$ and $\boldsymbol{\alpha} = \text{vec}(\mathbf{A}) = (\boldsymbol{\alpha}'_1, \dots, \boldsymbol{\alpha}'_n)'$. Furthermore, let $\boldsymbol{\alpha}_0 = (\boldsymbol{\alpha}'_{0,1}, \dots, \boldsymbol{\alpha}'_{0,n})'$ and $\mathbf{V}_{\boldsymbol{\alpha}} = \text{diag}(\mathbf{V}_{\boldsymbol{\alpha}_1}, \dots, \mathbf{V}_{\boldsymbol{\alpha}_n})$ and $k_2 = (2\pi)^{-\frac{nT+nk}{2}} e^{-\frac{1}{2}\mathbf{1}'_{nT}\mathbf{h}} |\mathbf{V}_{\boldsymbol{\alpha}}|^{-\frac{1}{2}}$. Then, we have

$$\begin{aligned} p(\mathbf{y} | \boldsymbol{\beta}, \mathbf{h}, \boldsymbol{\kappa}) &= \int p(\mathbf{y} | \boldsymbol{\beta}, \boldsymbol{\alpha}, \mathbf{h}) p(\boldsymbol{\alpha} | \boldsymbol{\kappa}) d\boldsymbol{\alpha} \\ &= k_2 \int e^{-\frac{1}{2}(\tilde{\mathbf{y}} - \tilde{\mathbf{X}}\boldsymbol{\alpha})' \mathbf{D}^{-1} (\tilde{\mathbf{y}} - \tilde{\mathbf{X}}\boldsymbol{\alpha}) - \frac{1}{2}(\boldsymbol{\alpha} - \boldsymbol{\alpha}_0)' \mathbf{V}_{\boldsymbol{\alpha}}^{-1} (\boldsymbol{\alpha} - \boldsymbol{\alpha}_0)} d\boldsymbol{\alpha} \\ &= k_2 e^{-\frac{1}{2}(\tilde{\mathbf{y}}' \mathbf{D}^{-1} \tilde{\mathbf{y}} + \boldsymbol{\alpha}'_0 \mathbf{V}_{\boldsymbol{\alpha}}^{-1} \boldsymbol{\alpha}_0 - \tilde{\boldsymbol{\alpha}}' \mathbf{K}_{\boldsymbol{\alpha}} \tilde{\boldsymbol{\alpha}})} \int e^{-\frac{1}{2}(\boldsymbol{\alpha} - \tilde{\boldsymbol{\alpha}})' \mathbf{K}_{\boldsymbol{\alpha}} (\boldsymbol{\alpha} - \tilde{\boldsymbol{\alpha}})} d\boldsymbol{\alpha} \\ &= (2\pi)^{-\frac{nT}{2}} e^{-\frac{1}{2}\mathbf{1}'_{nT}\mathbf{h}} |\mathbf{V}_{\boldsymbol{\alpha}}|^{-\frac{1}{2}} |\mathbf{K}_{\boldsymbol{\alpha}}|^{-\frac{1}{2}} e^{-\frac{1}{2}(\tilde{\mathbf{y}}' \mathbf{D}^{-1} \tilde{\mathbf{y}} + \boldsymbol{\alpha}'_0 \mathbf{V}_{\boldsymbol{\alpha}}^{-1} \boldsymbol{\alpha}_0 - \tilde{\boldsymbol{\alpha}}' \mathbf{K}_{\boldsymbol{\alpha}} \tilde{\boldsymbol{\alpha}})}, \end{aligned}$$

where

$$\mathbf{K}_\alpha = \mathbf{V}_\alpha^{-1} + \tilde{\mathbf{X}}' \mathbf{D}^{-1} \tilde{\mathbf{X}}, \quad \hat{\alpha} = \mathbf{K}_\alpha^{-1} \left(\mathbf{V}_\alpha^{-1} \alpha_0 + \tilde{\mathbf{X}}' \mathbf{D}^{-1} \tilde{\mathbf{y}} \right).$$

The last equality holds because

$$\int e^{-\frac{1}{2}(\alpha - \hat{\alpha})' \mathbf{K}_\alpha (\alpha - \hat{\alpha})} d\alpha = (2\pi)^{\frac{nk}{2}} |\mathbf{K}_\alpha|^{-\frac{1}{2}}.$$

Note that the vector of shrinkage hyperparameters $\boldsymbol{\kappa}$ appears in the prior covariance matrix \mathbf{V}_α .¹⁴ Now, we can write the marginal likelihood as:

$$\begin{aligned} p(\mathbf{y}) &= \int p(\mathbf{y} | \boldsymbol{\beta}, \mathbf{h}, \boldsymbol{\kappa}) p(\boldsymbol{\kappa}) \prod_{i=1}^n p(\boldsymbol{\beta}_i) p(\mathbf{h}_{i,\cdot} | \mu_i, \phi_i, \sigma_i^2) p(\mu_i) p(\phi_i) p(\sigma_i^2) d(\boldsymbol{\beta}, \mathbf{h}, \boldsymbol{\mu}, \boldsymbol{\phi}, \boldsymbol{\sigma}^2, \boldsymbol{\kappa}) \\ &= \int p(\mathbf{y} | \boldsymbol{\beta}, \mathbf{h}, \boldsymbol{\kappa}) p(\boldsymbol{\kappa}) \prod_{i=1}^n p(\boldsymbol{\beta}_i) p(\mathbf{h}_{i,\cdot} | \mu_i, \phi_i) p(\mu_i) p(\phi_i) d(\boldsymbol{\beta}, \mathbf{h}, \boldsymbol{\mu}, \boldsymbol{\phi}, \boldsymbol{\kappa}), \end{aligned}$$

where $p(\mathbf{h}_{i,\cdot} | \mu_i, \phi_i, \sigma_i^2)$ is a Gaussian density implied by the state equation (5), $p(\mu_i)$, $p(\phi_i)$, $p(\sigma_i^2)$ and $p(\boldsymbol{\kappa})$ are the prior densities. Furthermore, the density $p(\mathbf{h}_{i,\cdot} | \mu_i, \phi_i)$ has the following analytical expression:

$$\begin{aligned} p(\mathbf{h}_{i,\cdot} | \mu_i, \phi_i) &= \int p(\mathbf{h}_{i,\cdot} | \mu_i, \phi_i, \sigma_i^2) p(\sigma_i^2) d\sigma_i^2 \\ &= (2\pi)^{-\frac{T}{2}} |1 - \phi_i^2|^{\frac{1}{2}} S_i^{\nu_i} \Gamma(\nu_i)^{-1} \int (\sigma_i^2)^{-(\nu_i + \frac{T}{2} + 1)} e^{-\frac{\tilde{S}_i}{\sigma_i^2}} d\sigma_i^2, \\ &= (2\pi)^{-\frac{T}{2}} |1 - \phi_i^2|^{\frac{1}{2}} \frac{\Gamma(\nu_i + \frac{T}{2}) S_i^{\nu_i}}{\Gamma(\nu_i) \tilde{S}_i^{\nu_i + \frac{T}{2}}}, \end{aligned} \tag{19}$$

where $\Gamma(\cdot)$ is the gamma function and $\tilde{S}_i = S_i + [(1 - \phi_i^2)(h_{i,1} - \mu_i)^2 + \sum_{t=2}^T (h_{i,t} - \mu_i - \phi_i(h_{i,t-1} - \mu_i))^2]/2$.

Next, we combine the conditional Monte Carlo with the importance sampling approach described in Section 3.2. Specifically, we consider the parametric family

$$\mathcal{F} = \left\{ \prod_{i=1}^n f(\boldsymbol{\beta}_i; \hat{\boldsymbol{\beta}}_i, \hat{\mathbf{K}}_{\boldsymbol{\beta}_i}^{-1}) f_{\mathcal{N}}(\mathbf{h}_{i,\cdot}; \hat{\mathbf{h}}_{i,\cdot}, \hat{\mathbf{K}}_{\mathbf{h}_{i,\cdot}}^{-1}) f_{\mathcal{N}}(\mu_i; \hat{\mu}_i, \hat{K}_{\mu_i}^{-1}) f_{\mathcal{N}}(\phi_i; \hat{\phi}_i, \hat{K}_{\phi_i}^{-1}) \prod_{j=1}^3 f_{\mathcal{G}}(\kappa_j; \hat{\nu}_{\kappa_j}, \hat{S}_{\kappa_j}) \right\}.$$

¹⁴One potential limitation of this conditional Monte Carlo method is that it requires the computation of the precision matrix \mathbf{K}_α . For very large VARs ($n \approx 100$), the computational and memory requirements for this step might become excessive. One feasible way forward is to approximate this precision matrix using some low-rank decomposition. This possibility is left to future research.

The parameters of the importance sampling densities are chosen by solving the maximization problem in (9). In particular, the T -variate Gaussian density is obtained using the method described in Section 3.4. The other importance sampling densities can be obtained similarly.

Factor Stochastic Volatility

In this section we describe the posterior sampler for estimating the VAR-FSV model:

$$\begin{aligned}\mathbf{y}_t &= \mathbf{a}_0 + \mathbf{A}_1 \mathbf{y}_{t-1} + \cdots + \mathbf{A}_p \mathbf{y}_{t-p} + \boldsymbol{\varepsilon}_t, \\ \boldsymbol{\varepsilon}_t &= \mathbf{L} \mathbf{f}_t + \mathbf{u}_t^y, \quad \mathbf{u}_t^y \sim \mathcal{N}(\mathbf{0}, \mathbf{D}_t),\end{aligned}$$

where $\mathbf{f}_t = (f_{1,t}, \dots, f_{r,t})'$ is a $r \times 1$ vector of latent factors at time t , \mathbf{L} is the associated $n \times r$ matrix of factor loadings and $\mathbf{D}_t = \text{diag}(e^{h_{1,t}}, \dots, e^{h_{n,t}})$. The factors are serially independent and distributed as $\mathbf{f}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{G}_t)$, where $\mathbf{G}_t = \text{diag}(e^{h_{n+1,t}}, \dots, e^{h_{n+r,t}})$ is diagonal. For identification purposes we assume \mathbf{L} is a lower triangular matrix with ones on the main diagonal. Finally, for $i = 1, \dots, n+r$, the log-volatility evolves as:

$$h_{i,t} = \mu_i + \phi_i(h_{i,t-1} - \mu_i) + u_{i,t}^h, \quad u_{i,t}^h \sim \mathcal{N}(0, \sigma_i^2)$$

for $t = 2, \dots, T$, and the initial conditions are assumed to follow the stationary distribution $h_{i,1} \sim \mathcal{N}(\mu_i, \sigma_i^2 / (1 - \phi_i^2))$.

Next, we specify the prior distributions on the model parameters. Let $\boldsymbol{\alpha}_i$ and \mathbf{l}_i denote the VAR coefficients and the free elements of \mathbf{L} in the i -th equation, respectively, for $i = 1, \dots, n$. We assume the following independent priors on $\boldsymbol{\alpha}_i$ and \mathbf{l}_i for $i = 1, \dots, n$:

$$\boldsymbol{\alpha}_i \sim \mathcal{N}(\boldsymbol{\alpha}_{0,i}, \mathbf{V}_{\boldsymbol{\alpha}_i}), \quad \mathbf{l}_i \sim \mathcal{N}(\mathbf{l}_{0,i}, \mathbf{V}_{\mathbf{l}_i}).$$

We set the prior mean $\mathbf{l}_{0,i}$ to be zero and the prior covariance matrix $\mathbf{V}_{\mathbf{l}_i}$ to be the identity matrix. As before, the prior mean and prior covariance matrix are elicited according to the Minnesota prior. In particular, we set $\boldsymbol{\alpha}_{0,i} = \mathbf{0}$ and $\mathbf{V}_{\boldsymbol{\alpha}_i}$ is set to be diagonal with

the k -th diagonal element $V_{\alpha_i, kk}$:

$$V_{\alpha_i, kk} = \begin{cases} \frac{\kappa_1}{l^2}, & \text{for the coefficient on the } l\text{-th lag of variable } i, \\ \frac{\kappa_2 s_i^2}{l^2 s_j^2}, & \text{for the coefficient on the } l\text{-th lag of variable } j, j \neq i, \\ 100 s_i^2, & \text{for the intercept,} \end{cases}$$

where s_r^2 denotes the sample variance of the residuals from an AR(4) model for the variable $r, r = 1, \dots, n$. Again the shrinkage hyperparameters $\boldsymbol{\kappa} = (\kappa_1, \kappa_2)'$ are treated as unknown parameters to be estimated with hierarchical gamma priors $\kappa_j \sim \mathcal{G}(c_{j,1}, c_{j,2}), j = 1, 2$. For the parameters in the stochastic volatility equations, we assume the same priors as in (17):

$$\mu_i \sim \mathcal{N}(\mu_{0,i}, V_{\mu_i}), \quad \phi_i \sim \mathcal{N}(\phi_{0,i}, V_{\phi_i})1(|\phi_i| < 1), \quad \sigma_i^2 \sim \mathcal{IG}(\nu_i, S_i), \quad i = 1, \dots, n.$$

For notational convenience, stack $\mathbf{y} = (\mathbf{y}'_1, \dots, \mathbf{y}'_T)'$, $\mathbf{f} = (\mathbf{f}'_1, \dots, \mathbf{f}'_T)'$, $\mathbf{h} = (\mathbf{h}'_1, \dots, \mathbf{h}'_{n+r})'$, $\boldsymbol{\alpha} = (\boldsymbol{\alpha}'_1, \dots, \boldsymbol{\alpha}'_n)'$ and $\mathbf{l} = (\mathbf{l}'_1, \dots, \mathbf{l}'_n)'$; similarly define $\boldsymbol{\mu}, \boldsymbol{\phi}$ and $\boldsymbol{\sigma}^2$. In addition, let $\mathbf{y}_{i,\cdot} = (y_{i,1}, \dots, y_{i,T})'$ denote the vector of observed values for the i -th variable for $i = 1, \dots, n$. We similarly define $\mathbf{h}_{i,\cdot} = (h_{i,1}, \dots, h_{i,T})', i = 1, \dots, n+r$. Then, posterior draws can be obtained by sampling sequentially from:

1. $p(\mathbf{f} | \mathbf{y}, \mathbf{l}, \boldsymbol{\alpha}, \mathbf{h}, \boldsymbol{\mu}, \boldsymbol{\phi}, \boldsymbol{\sigma}^2, \boldsymbol{\kappa}) = p(\mathbf{f} | \mathbf{y}, \mathbf{l}, \boldsymbol{\alpha}, \boldsymbol{\rho}, \mathbf{h});$
2. $p(\boldsymbol{\alpha}, \mathbf{l} | \mathbf{y}, \mathbf{f}, \mathbf{h}, \boldsymbol{\mu}, \boldsymbol{\phi}, \boldsymbol{\sigma}^2, \boldsymbol{\kappa}) = \prod_{i=1}^n p(\boldsymbol{\alpha}_i, \mathbf{l}_i | \mathbf{y}_{i,\cdot}, \mathbf{f}, \mathbf{h}_{i,\cdot}, \boldsymbol{\kappa});$
3. $p(\mathbf{h} | \mathbf{y}, \mathbf{f}, \boldsymbol{\alpha}, \mathbf{l}, \boldsymbol{\mu}, \boldsymbol{\phi}, \boldsymbol{\sigma}^2, \boldsymbol{\kappa}) = \prod_{i=1}^{n+r} p(\mathbf{h}_{i,\cdot} | \mathbf{y}, \mathbf{f}, \boldsymbol{\alpha}, \mathbf{l}, \boldsymbol{\mu}, \boldsymbol{\phi}, \boldsymbol{\sigma}^2);$
4. $p(\boldsymbol{\sigma}^2 | \mathbf{y}, \mathbf{f}, \boldsymbol{\alpha}, \mathbf{l}, \mathbf{h}, \boldsymbol{\mu}, \boldsymbol{\phi}, \boldsymbol{\kappa}) = \prod_{i=1}^{n+r} p(\sigma_i^2 | \mathbf{h}_{i,\cdot}, \mu_i, \phi_i);$
5. $p(\boldsymbol{\mu} | \mathbf{y}, \mathbf{f}, \boldsymbol{\alpha}, \mathbf{l}, \mathbf{h}, \boldsymbol{\phi}, \boldsymbol{\sigma}^2, \boldsymbol{\kappa}) = \prod_{i=1}^{n+r} p(\mu_i | \mathbf{h}_{i,\cdot}, \phi_i, \sigma_i^2);$
6. $p(\boldsymbol{\phi} | \mathbf{y}, \mathbf{f}, \boldsymbol{\alpha}, \mathbf{l}, \mathbf{h}, \boldsymbol{\mu}, \boldsymbol{\sigma}^2, \boldsymbol{\kappa}) = \prod_{i=1}^{n+r} p(\phi_i | \mathbf{h}_{i,\cdot}, \mu_i, \sigma_i^2);$
7. $p(\boldsymbol{\kappa} | \mathbf{y}, \mathbf{f}, \mathbf{l}, \boldsymbol{\alpha}, \mathbf{h}, \boldsymbol{\mu}, \boldsymbol{\phi}, \boldsymbol{\sigma}^2) = p(\boldsymbol{\kappa} | \boldsymbol{\alpha}).$

Step 1: to sample \mathbf{f} , we first stack the VAR over $t = 1, \dots, T$ and write it as:

$$\boldsymbol{\varepsilon} = (\mathbf{I}_T \otimes \mathbf{L})\mathbf{f} + \mathbf{u}^y, \quad \mathbf{u}^y \sim \mathcal{N}(\mathbf{0}, \mathbf{D}),$$

where $\boldsymbol{\varepsilon} = (\boldsymbol{\varepsilon}'_1, \dots, \boldsymbol{\varepsilon}'_T)'$ is known given $\boldsymbol{\alpha}$ and the data and $\mathbf{D} = \text{diag}(\mathbf{D}_1, \dots, \mathbf{D}_T)$. In addition, the prior on the factors can be written as $\mathbf{f} \sim \mathcal{N}(\mathbf{0}, \mathbf{G})$, where $\mathbf{G} = \text{diag}(\mathbf{G}_1, \dots, \mathbf{G}_T)$. It follows from standard linear regression results that (see, e.g., Chan, Koop, Poirier, and Tobias, 2019, chapter 12)

$$(\mathbf{f} \mid \mathbf{y}, \boldsymbol{\alpha}, \mathbf{l}, \mathbf{h}) \sim \mathcal{N}(\hat{\mathbf{f}}, \mathbf{K}_f^{-1}),$$

where

$$\mathbf{K}_f = \mathbf{G}^{-1} + (\mathbf{I}_T \otimes \mathbf{L}')\mathbf{D}^{-1}(\mathbf{I}_T \otimes \mathbf{L}), \quad \hat{\mathbf{f}} = \mathbf{K}_f^{-1}(\mathbf{I}_T \otimes \mathbf{L}')\mathbf{D}^{-1}\boldsymbol{\varepsilon}.$$

Note that the precision matrix \mathbf{K}_f is banded, and we can use the precision sampler of Chan and Jeliazkov (2009) to sample \mathbf{f} efficiently.

Step 2: to sample $\boldsymbol{\alpha}$ and \mathbf{l} jointly, first note that given the latent factors \mathbf{f} , the VAR becomes n unrelated regressions and we can sample $\boldsymbol{\alpha}$ and \mathbf{l} equation by equation. Recall that $\mathbf{y}_{i,\cdot} = (y_{i,1}, \dots, y_{i,T})'$ is defined to be the $T \times 1$ vector of observations for the i -th variable; $\boldsymbol{\alpha}_i$ and \mathbf{l}_i denote, respectively, the VAR coefficients and the free element of \mathbf{L} in the i -th equation. Note that the dimension of \mathbf{l}_i is $i - 1$ for $i \leq r$ and r for $i > r$. Then, the i -th equation of the VAR can be written as

$$\begin{aligned} \mathbf{y}_{i,\cdot} &= \mathbf{X}_i \boldsymbol{\alpha}_i + \mathbf{F}_{1:i-1} \mathbf{l}_i + \mathbf{f}_{i,\cdot} + \mathbf{u}_{i,\cdot}^y, & i \leq r, \\ \mathbf{y}_{i,\cdot} &= \mathbf{X}_i \boldsymbol{\alpha}_i + \mathbf{F}_{1:r} \mathbf{l}_i + \mathbf{u}_{i,\cdot}^y, & i > r, \end{aligned}$$

where $\mathbf{f}_{i,\cdot} = (f_{i,1}, \dots, f_{i,T})'$ is the $T \times 1$ vector of the i -th factor and $\mathbf{F}_{1:j} = (\mathbf{f}_{1,\cdot}, \dots, \mathbf{f}_{j,\cdot})$ is the $T \times j$ matrix that contains the first j factors. The vector of innovations $\mathbf{u}_{i,\cdot}^y = (u_{i,1}, \dots, u_{i,T})'$ is distributed as $\mathcal{N}(\mathbf{0}, \mathbf{D}_{\mathbf{h}_{i,\cdot}})$, where $\mathbf{D}_{\mathbf{h}_{i,\cdot}} = \text{diag}(e^{h_{i,1}}, \dots, e^{h_{i,T}})$. Letting $\boldsymbol{\theta}_i = (\boldsymbol{\alpha}'_i, \mathbf{l}'_i)'$, we can further write the VAR systems as

$$\tilde{\mathbf{y}}_{i,\cdot} = \mathbf{Z}_i \boldsymbol{\theta}_i + \mathbf{u}_{i,\cdot}^y, \tag{20}$$

where $\tilde{\mathbf{y}}_{i,\cdot} = \mathbf{y}_{i,\cdot} - \mathbf{f}_{i,\cdot}$ and $\mathbf{Z}_i = (\mathbf{X}_i, \mathbf{F}_{1:i-1})$ for $i \leq r$; $\tilde{\mathbf{y}}_i = \mathbf{y}_i$ and $\mathbf{Z}_i = (\mathbf{X}_i, \mathbf{F}_{1:r})$ for $i > r$. Again using standard linear regression results, we obtain:

$$(\boldsymbol{\theta}_i \mid \mathbf{y}_{i,\cdot}, \mathbf{f}, \mathbf{h}_{i,\cdot}, \boldsymbol{\kappa}) \sim \mathcal{N}(\hat{\boldsymbol{\theta}}_i, \mathbf{K}_{\boldsymbol{\theta}_i}^{-1}),$$

where

$$\mathbf{K}_{\theta_i} = \mathbf{V}_{\theta_i}^{-1} + \mathbf{Z}_i' \mathbf{D}_{\mathbf{h}_{i,\cdot}}^{-1} \mathbf{Z}_i, \quad \hat{\boldsymbol{\theta}}_i = \mathbf{K}_{\theta_i}^{-1} (\mathbf{V}_{\theta_i}^{-1} \boldsymbol{\theta}_{0,i} + \mathbf{Z}_i \mathbf{D}_{\mathbf{h}_{i,\cdot}}^{-1} \tilde{\mathbf{y}}_{i,\cdot})$$

with $\mathbf{V}_{\theta_i} = \text{diag}(\mathbf{V}_{\alpha_i}, \mathbf{V}_{\mathbf{1}_i})$ and $\boldsymbol{\theta}_{0,i} = (\boldsymbol{\alpha}'_{0,i}, \mathbf{1}'_{0,i})'$.

Step 3: to sample \mathbf{h} , again note that given the latent factors \mathbf{f} , the VAR becomes n unrelated regressions and we can sample each vector $\mathbf{h}_{i,\cdot} = (h_{i,1}, \dots, h_{i,T})'$ separately. More specifically, we can directly apply the auxiliary mixture sampler in Kim, Shephard, and Chib (1998) in conjunction with the precision sampler of Chan and Jeliazkov (2009) to sample from $(\mathbf{h}_{i,\cdot} | \mathbf{y}, \mathbf{f}, \boldsymbol{\alpha}, \mathbf{1}, \boldsymbol{\mu}, \boldsymbol{\phi}, \boldsymbol{\sigma}^2)$ for $i = 1, \dots, n + r$. For a textbook treatment, see, e.g., Chan, Koop, Poirier, and Tobias (2019) chapter 19.

Step 4: this step can be done easily, as the elements of $\boldsymbol{\sigma}^2$ are conditionally independent and each follows an inverse-gamma distribution:

$$(\sigma_i^2 | \mathbf{h}_{i,\cdot}, \mu_i, \phi_i) \sim \mathcal{IG}(\nu_i + T/2, \tilde{S}_i),$$

where $\tilde{S}_i = S_i + [(1 - \phi_i^2)(h_{i,1} - \mu_i)^2 + \sum_{t=2}^T (h_{i,t} - \mu_i - \phi_i(h_{i,t-1} - \mu_i))^2]/2$.

Step 5: it is also straightforward to sample $\boldsymbol{\mu}$, as the elements of $\boldsymbol{\mu}$ are conditionally independent and each follows a normal distribution:

$$(\mu_i | \mathbf{h}_{i,\cdot}, \phi_i, \sigma_i^2) \sim \mathcal{N}(\hat{\mu}_i, K_{\mu_i}^{-1}),$$

where

$$K_{\mu_i} = V_{\mu_i}^{-1} + \frac{1}{\sigma_i^2} [1 - \phi_i^2 + (T - 1)(1 - \phi_i)^2]$$

$$\hat{\mu}_i = K_{\mu_i}^{-1} \left[V_{\mu_i}^{-1} \mu_{0,i} + \frac{1}{\sigma_i^2} \left((1 - \phi_i^2) h_{i,1} + (1 - \phi_i) \sum_{t=2}^T (h_{i,t} - \phi_i h_{i,t-1}) \right) \right].$$

Step 6: to sample ϕ_i for $i = 1, \dots, n + r$, note that

$$p(\phi_i | \mathbf{h}_{i,\cdot}, \mu_i, \sigma_i^2) \propto p(\phi_i) g(\phi_i) e^{-\frac{1}{2\sigma_i^2} \sum_{t=2}^T (h_{i,t} - \mu_i - \phi_i(h_{i,t-1} - \mu_i))^2},$$

where $g(\phi_i) = (1 - \phi_i^2)^{\frac{1}{2}} e^{-\frac{1}{2\sigma_i^2} (1 - \phi_i^2)(h_{i,1} - \mu_i)^2}$ and $p(\phi_i)$ is the truncated normal prior. The conditional density $p(\phi_i | \mathbf{h}_{i,\cdot}, \mu_i, \sigma_i^2)$ is nonstandard, but a draw from it can be

obtained by using an independence-chain Metropolis-Hastings step with proposal distribution $\mathcal{N}(\widehat{\phi}_i, K_{\phi_i}^{-1})1(|\phi_i| < 1)$, where

$$K_{\phi_i} = V_{\phi_i}^{-1} + \frac{1}{\sigma_i^2} \sum_{t=2}^T (h_{i,t-1} - \mu_i)^2$$

$$\widehat{\phi}_h = K_{\phi_i}^{-1} \left[V_{\phi_i}^{-1} \phi_{0,i} + \frac{1}{\sigma_i^2} \sum_{t=2}^T (h_{i,t-1} - \mu_i)(h_{i,t} - \mu_i) \right].$$

Then, given the current draw ϕ_i , a proposal ϕ_i^* is accepted with probability $\min(1, g(\phi_i^*)/g(\phi_i))$; otherwise the Markov chain stays at the current state ϕ_i .

Step 7: lastly, sampling $\boldsymbol{\kappa} = (\kappa_1, \kappa_2)'$ can be done similarly as in other stochastic volatility models. More specifically, define the index set S_{κ_1} that collects all the indexes (i, j) such that $\alpha_{i,j}$, the j -th element of $\boldsymbol{\alpha}_i$, is a coefficient associated with an own lag and let S_{κ_2} denote the set that collects all the indexes (i, j) such that $\alpha_{i,j}$ is a coefficient associated with a lag of other variables. Then, given the priors $\kappa_j \sim \mathcal{G}(c_{j,1}, c_{j,2}), j = 1, 2$, and the prior covariance matrix of $\boldsymbol{\alpha}_i$, we have

$$(\kappa_1 | \boldsymbol{\alpha}) \sim \mathcal{GIG} \left(c_{1,1} - \frac{np}{2}, 2c_{1,2}, \sum_{(i,j) \in S_{\kappa_1}} \frac{(\alpha_{i,j} - \alpha_{0,i,j})^2}{C_{i,j}} \right)$$

$$(\kappa_2 | \boldsymbol{\alpha}) \sim \mathcal{GIG} \left(c_{2,1} - \frac{(n-1)np}{2}, 2c_{2,2}, \sum_{(i,j) \in S_{\kappa_2}} \frac{(\alpha_{i,j} - \alpha_{0,i,j})^2}{C_{i,j}} \right),$$

where $\alpha_{0,i,j}$ is the j -th element of the prior mean vector $\boldsymbol{\alpha}_{0,i}$ and $C_{i,j}$ is a constant determined by the Minnesota prior.

Next, we turn to the computation of the marginal likelihood of VAR-FSV. Here the marginal likelihood estimator has two parts: the conditional Monte Carlo where we integrate out the VAR coefficients and the latent factors; and the adaptive importance sampling that biases the joint distribution of $\mathbf{h}, \mathbf{l}, \boldsymbol{\mu}, \boldsymbol{\phi}$ and $\boldsymbol{\kappa}$. In what follows, we first derive an analytical expression of the conditional Monte Carlo estimator $\mathbb{E}[p(\mathbf{y} | \boldsymbol{\alpha}, \mathbf{l}, \mathbf{f}, \mathbf{h}, \boldsymbol{\kappa}) | \mathbf{l}, \mathbf{h}, \boldsymbol{\kappa}] = p(\mathbf{y} | \mathbf{l}, \mathbf{h}, \boldsymbol{\kappa})$. To that end, write the VAR-FSV model as

$$\mathbf{y} = \mathbf{X}\boldsymbol{\alpha} + (\mathbf{I}_T \otimes \mathbf{L})\mathbf{f}, \quad \mathbf{u}^y \sim \mathcal{N}(\mathbf{0}, \mathbf{D}),$$

where \mathbf{X} an appropriately defined covariate matrix consisting of intercepts and lagged values and $\mathbf{f} \sim \mathcal{N}(\mathbf{0}, \mathbf{G})$. Hence, the distribution of \mathbf{y} marginal of the factors \mathbf{f} is $(\mathbf{y} | \boldsymbol{\alpha}, \mathbf{l}, \mathbf{h}) \sim \mathcal{N}(\mathbf{X}\boldsymbol{\alpha}, \mathbf{S}_y)$, where $\mathbf{S}_y = \mathbf{D} + (\mathbf{I}_T \otimes \mathbf{L})\mathbf{G}(\mathbf{I}_T \otimes \mathbf{L}')$. Next, following a similar calculation as in the case for VAR-SV, one can show that

$$\begin{aligned} p(\mathbf{y} | \mathbf{l}, \mathbf{h}, \boldsymbol{\kappa}) &= \int p(\mathbf{y} | \boldsymbol{\alpha}, \mathbf{l}, \mathbf{h})p(\boldsymbol{\alpha} | \boldsymbol{\kappa})d\boldsymbol{\alpha} \\ &= (2\pi)^{-\frac{Tn}{2}}|\mathbf{S}_y|^{-\frac{1}{2}}|\mathbf{V}_\alpha|^{-\frac{1}{2}}|\mathbf{K}_\alpha|^{-\frac{1}{2}}e^{-\frac{1}{2}(\mathbf{y}'\mathbf{S}_y^{-1}\mathbf{y} + \boldsymbol{\alpha}'_0\mathbf{V}_\alpha^{-1}\boldsymbol{\alpha}_0 - \hat{\boldsymbol{\alpha}}'\mathbf{K}_\alpha\hat{\boldsymbol{\alpha}})}, \end{aligned}$$

where $\mathbf{K}_\alpha = \mathbf{V}_\alpha^{-1} + \mathbf{X}'\mathbf{S}_y^{-1}\mathbf{X}$ and $\hat{\boldsymbol{\alpha}} = \mathbf{K}_\alpha^{-1}(\mathbf{V}_\alpha^{-1}\boldsymbol{\alpha}_0 + \mathbf{X}'\mathbf{S}_y^{-1}\mathbf{y})$.

We can now write the marginal likelihood as:

$$\begin{aligned} p(\mathbf{y}) &= \int p(\mathbf{y} | \mathbf{l}, \mathbf{h}, \boldsymbol{\kappa})p(\mathbf{l})p(\boldsymbol{\mu})p(\boldsymbol{\phi})p(\boldsymbol{\sigma}^2)p(\boldsymbol{\kappa}) \prod_{j=1}^{n+r} p(\mathbf{h}_{j,\cdot} | \mu_j, \phi_j, \sigma_j^2)d(\mathbf{l}, \mathbf{h}, \boldsymbol{\mu}, \boldsymbol{\phi}, \boldsymbol{\sigma}^2, \boldsymbol{\kappa}) \\ &= \int p(\mathbf{y} | \mathbf{l}, \mathbf{h}, \boldsymbol{\kappa})p(\mathbf{l})p(\boldsymbol{\mu})p(\boldsymbol{\phi})p(\boldsymbol{\kappa}) \prod_{j=1}^{n+r} p(\mathbf{h}_{j,\cdot} | \mu_j, \phi_j)d(\mathbf{l}, \mathbf{h}, \boldsymbol{\mu}, \boldsymbol{\phi}, \boldsymbol{\kappa}), \end{aligned}$$

where $p(\mathbf{h}_{i,\cdot} | \mu_i, \phi_i) = \int p(\mathbf{h}_{i,\cdot} | \mu_i, \phi_i, \sigma_i^2)p(\sigma_i^2)d\sigma_i^2$ is available analytically as given in (19).

To combine the conditional Monte Carlo with the importance sampling approach described in Section 3.2, we consider the parametric family

$$\mathcal{F} = \left\{ f_{\mathcal{N}}(\mathbf{l}; \hat{\mathbf{l}}, \hat{\mathbf{K}}_1^{-1}) \prod_{j=1}^{n+r} f_{\mathcal{N}}(\mathbf{h}_{j,\cdot}; \hat{\mathbf{h}}_{j,\cdot}, \hat{\mathbf{K}}_{\mathbf{h}_{j,\cdot}}^{-1}) f_{\mathcal{N}}(\mu_j; \hat{\mu}_j, \hat{K}_{\mu_j}^{-1}) f_{\mathcal{N}}(\phi_j; \hat{\phi}_j, \hat{K}_{\phi_j}^{-1}) \prod_{k=1}^2 f_{\mathcal{G}}(\kappa_k; \hat{\nu}_{\kappa_k}, \hat{S}_{\kappa_k}) \right\}.$$

The parameters of the importance sampling densities are chosen by solving the maximization problem in (9). In particular, all the T -variate Gaussian densities are obtained using the procedure described in Section 3.4. Other low dimensional importance sampling densities can be obtained following Chan and Eisenstat (2015).

Appendix B: Data

The dataset is sourced from the Federal Reserve Bank of St. Louis and covers the quarters from 1959:Q1 to 2019:Q4. Table 7 lists the 30 quarterly variables and describes how they are transformed. For example, $\Delta \log$ is used to denote the first difference in the logs, i.e., $\Delta \log x = \log x_t - \log x_{t-1}$.

Table 7: Description of variables used in the empirical application.

Variable	Transformation	$n = 7$	$n = 15$	$n = 30$
Real Gross Domestic Product	400 $\Delta \log$	x	x	x
Personal Consumption Expenditures	400 $\Delta \log$		x	x
Real personal consumption expenditures: Durable goods	400 $\Delta \log$			x
Real Disposable Personal Income	400 $\Delta \log$			x
Industrial Production Index	400 $\Delta \log$	x	x	x
Industrial Production: Final Products	400 $\Delta \log$			x
All Employees: Total nonfarm	400 $\Delta \log$		x	x
All Employees: Manufacturing	400 $\Delta \log$			x
Civilian Employment	400 $\Delta \log$		x	x
Civilian Labor Force Participation Rate	no trans.			x
Civilian Unemployment Rate	no trans.	x	x	x
Nonfarm Business Section: Hours of All Persons	400 $\Delta \log$			x
Housing Starts: Total	400 $\Delta \log$		x	x
New Private Housing Units Authorized by Building Permits	400 $\Delta \log$			x
Personal Consumption Expenditures: Chain-type Price index	400 $\Delta \log$		x	x
Gross Domestic Product: Chain-type Price index	400 $\Delta \log$			x
Consumer Price Index for All Urban Consumers: All Items	400 $\Delta \log$	x	x	x
Producer Price Index for All commodities	400 $\Delta \log$			x
Real Average Hourly Earnings of Production and Nonsupervisory Employees: Manufacturing	400 $\Delta \log$	x	x	x
Nonfarm Business Section: Real Output Per Hour of All Persons	400 $\Delta \log$			x
Effective Federal Funds Rate	no trans.	x	x	x
3-Month Treasury Bill: Secondary Market Rate	no trans.			x
1-Year Treasury Constant Maturity Rate	no trans.			x
10-Year Treasury Constant Maturity Rate	no trans.	x	x	x
Moody's Seasoned Baa Corporate Bond Yield Relative to Yield on 10-Year Treasury Constant Maturity	no trans.		x	x
3-Month Commercial Paper Minus 3-Month Treasury Bill	no trans.			x
Real M1 Money Stock	400 $\Delta \log$		x	x
Real M2 Money Stock	400 $\Delta \log$			x
Total Reserves of Depository Institutions	$\Delta^2 \log$			x
S&P's Common Stock Price Index : Composite	400 $\Delta \log$		x	x

Appendix C: Additional Simulation Results

This appendix provides various additional simulation results to further assess the performance of the proposed marginal likelihood estimators.

The first set of results assesses whether one can distinguish the three stochastic volatility models—the common stochastic volatility model (VAR-CSV), the Cholesky stochastic volatility model (VAR-SV) and the factor stochastic volatility model (VAR-FSV)—in higher dimensions. More specifically, we repeat the Monte Carlo exercise in Section 4.1 of the main text, but here each VAR has $n = 20$ variables instead of $n = 10$.

We first generate 100 datasets from VAR-CSV as described in Section 4.1, and for each dataset, we compute the log marginal likelihoods of VAR-CSV and VAR-FSV relative to that of VAR-CSV. The results reported in Figure 7 show that for all datasets, the proposed marginal likelihood estimators select the correct DGP. In fact, compared to Figure 1 based on VARs with $n = 10$ variables, here the differences in the marginal likelihood values are larger. This suggests that it becomes easier to discriminate among the stochastic volatility models as the number of variables increases.

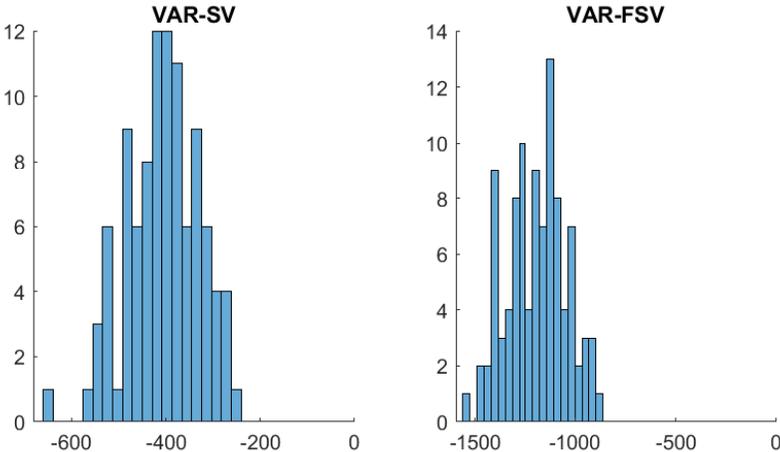


Figure 7: Histograms of log marginal likelihoods under VAR-SV (left panel) and VAR-FSV (right panel) relative to the true model (VAR-CSV) for $n = 20$. A negative value indicates that the correct model is favored.

Figures 8 and 9 present results based on datasets generated from VAR-SV and VAR-FSV, respectively. Once again, these model comparison results show that, for all datasets, the

correct model is overwhelmingly favored compared to the other two stochastic volatility specifications.

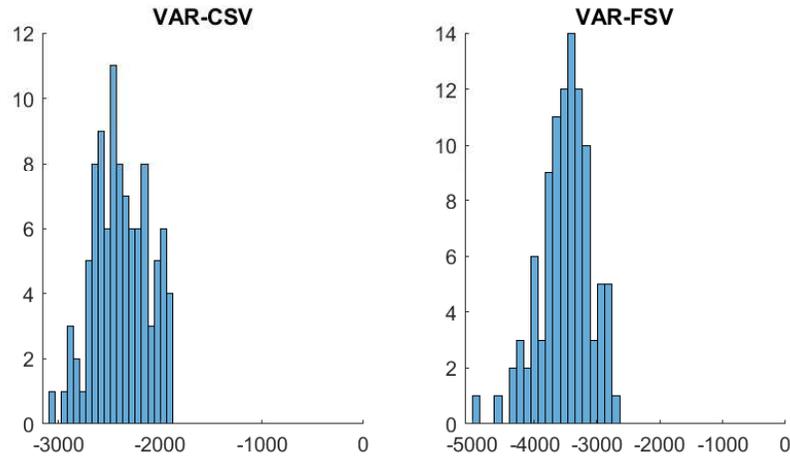


Figure 8: Histograms of log marginal likelihoods under VAR-CSV (left panel) and VAR-FSV (right panel) relative to the true model (VAR-SV) for $n = 20$. A negative value indicates that the correct model is favored.

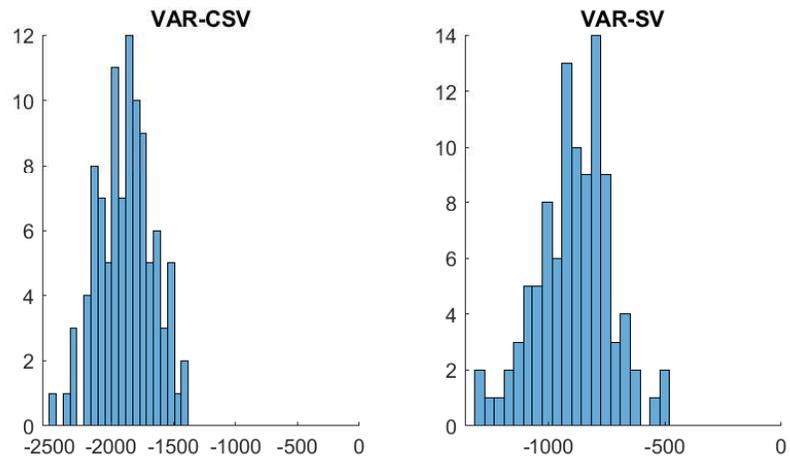


Figure 9: Histograms of log marginal likelihoods under VAR-CSV (left panel) and VAR-SV (right panel) relative to the true model (VAR-FSV) for $n = 20$. A negative value indicates that the correct model is favored.

The next set of results assesses how the proposed marginal likelihood estimators perform when the magnitude of the shocks in the VAR equation increases. More specifically, we

use the same settings as described in Section 4.1 of the main text, but here we increase the magnitude of the error covariance matrix in the VAR equation by a factor of 10. In particular, for VAR-CSV, the error covariance matrix Σ is generated from the inverse-Wishart distribution $\mathcal{IW}(n+5, 7 \times \mathbf{I}_n + 3 \times \mathbf{1}_n \mathbf{1}'_n)$; for VAR-SV and VAR-FSV, we increase the values of $\mu_i, i = 1, \dots, n$, i.e., the means of the log-volatility in the VAR equation, from -1 to 1.3 . This difference of 2.3 in natural log corresponds to approximately a 10-fold increase in volatility.

The model comparison results are reported in Figures 10-12. These results from the noisier DGPs are very similar to those reported in Section 4.1. In particular, for all datasets, the correct model is overwhelmingly favored compared to the other two stochastic volatility models. One minor difference in the two sets of results is that when the data is generated from VAR-FSV, it becomes slightly harder to distinguish between VAR-SV and VAR-FSV (reflected in the smaller differences in marginal likelihood values).

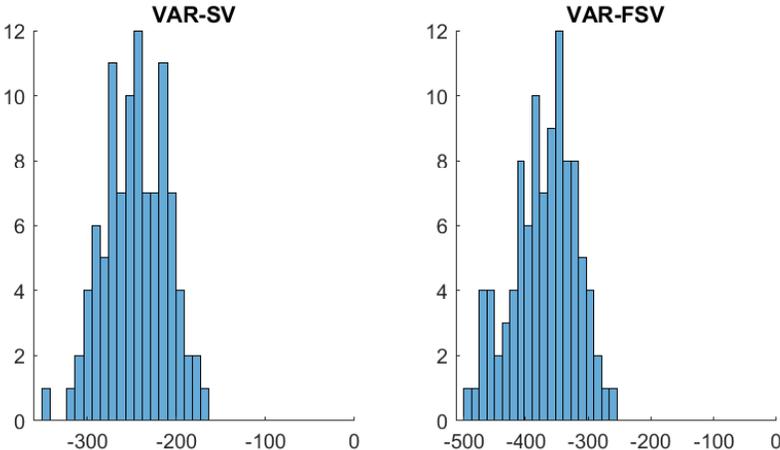


Figure 10: Histograms of log marginal likelihoods under VAR-SV (left panel) and VAR-FSV (right panel) relative to the true model (VAR-CSV) with a noisier VAR equation. A negative value indicates that the correct model is favored.

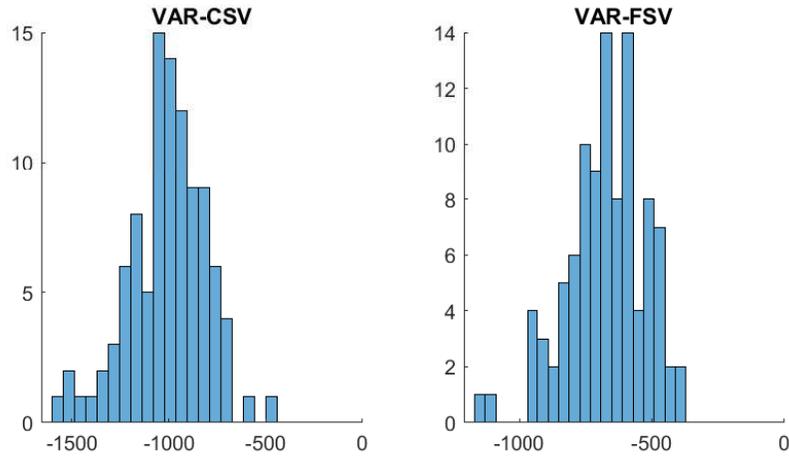


Figure 11: Histograms of log marginal likelihoods under VAR-CSV (left panel) and VAR-FSV (right panel) relative to the true model (VAR-SV) with a noisier VAR equation. A negative value indicates that the correct model is favored.

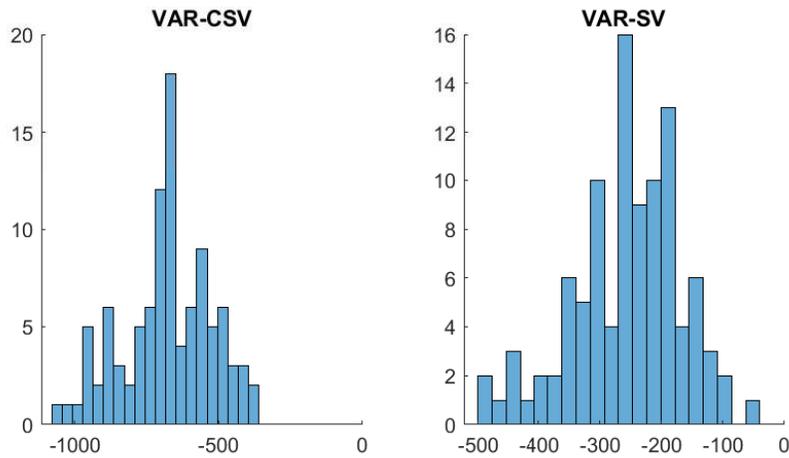


Figure 12: Histograms of log marginal likelihoods under VAR-CSV (left panel) and VAR-SV (right panel) relative to the true model (VAR-FSV) with a noisier VAR equation. A negative value indicates that the correct model is favored.

References

- AGUILAR, O., AND M. WEST (2000): “Bayesian dynamic factor models and portfolio allocation,” *Journal of Business and Economic Statistics*, 18(3), 338–357.
- AMIR-AHMADI, P., C. MATTHES, AND M.-C. WANG (2020): “Choosing prior hyperparameters: With applications to time-varying parameter models,” *Journal of Business and Economic Statistics*, 38(1), 124–136.
- ANDERSON, T. W., AND H. RUBIN (1956): “Statistical inference in factor analysis,” in *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*, vol. 5, pp. 111–150.
- BANBURA, M., D. GIANNONE, M. MODUGNO, AND L. REICHLIN (2013): “Now-casting and the real-time data flow,” in *Handbook of Economic Forecasting*, vol. 2, pp. 195–237. Elsevier.
- BANBURA, M., D. GIANNONE, AND L. REICHLIN (2010): “Large Bayesian vector autoregressions,” *Journal of Applied Econometrics*, 25(1), 71–92.
- BANBURA, M., AND A. VAN VLODROP (2018): “Forecasting with Bayesian Vector Autoregressions with Time Variation in the Mean,” *Tinbergen Institute Discussion Paper 2018-025/IV*.
- BAUMEISTER, C., D. KOROBILIS, AND T. K. LEE (2020): “Energy markets and global economic conditions,” *Review of Economics and Statistics*, forthcoming.
- BHATTACHARYA, A., AND D. DUNSON (2011): “Sparse Bayesian infinite factor models,” *Biometrika*, pp. 291–306.
- BIANCHI, D., M. GUIDOLIN, AND F. RAVAZZOLO (2018): “Dissecting the 2007–2009 Real Estate Market Bust: Systematic Pricing Correction or Just a Housing Fad?,” *Journal of Financial Econometrics*, 16(1), 34–62.
- BOGNANNI, M. (2022): “Comment on “Large Bayesian vector autoregressions with stochastic volatility and non-conjugate priors”,” *Journal of Econometrics*, 227(2), 498–505.
- BOGNANNI, M., AND J. ZITO (2020): “Sequential Bayesian inference for vector autoregressions with stochastic volatility,” *Journal of Economic Dynamics and Control*, 113, 103851.
- CARRIERO, A., J. C. C. CHAN, T. E. CLARK, AND M. G. MARCELLINO (2022): “Corrigendum to “Large Bayesian vector autoregressions with stochastic volatility and non-conjugate priors”,” *Journal of Econometrics*, 227(2), 506–512.

- CARRIERO, A., T. E. CLARK, M. MARCELLINO, AND E. MERTENS (2022): “Addressing COVID-19 outliers in BVARs with stochastic volatility,” *The Review of Economics and Statistics*, Forthcoming.
- CARRIERO, A., T. E. CLARK, AND M. G. MARCELLINO (2015): “Bayesian VARs: Specification Choices and Forecast Accuracy,” *Journal of Applied Econometrics*, 30(1), 46–73.
- (2016): “Common drifting volatility in large Bayesian VARs,” *Journal of Business and Economic Statistics*, 34(3), 375–390.
- (2019): “Large Bayesian vector autoregressions with stochastic volatility and non-conjugate priors,” *Journal of Econometrics*, 212(1), 137–154.
- CARRIERO, A., G. KAPETANIOS, AND M. MARCELLINO (2009): “Forecasting exchange rates with a large Bayesian VAR,” *International Journal of Forecasting*, 25(2), 400–417.
- CELEUX, G., F. FORBES, C. P. ROBERT, AND D. M. TITTERINGTON (2006): “Deviance Information Criteria for Missing Data Models,” *Bayesian Analysis*, 1(4), 651–674.
- CHAN, J. C. C. (2020): “Large Bayesian VARs: A Flexible Kronecker Error Covariance Structure,” *Journal of Business and Economic Statistics*, 38(1), 68–79.
- (2021): “Minnesota-type adaptive hierarchical priors for large Bayesian VARs,” *International Journal of Forecasting*, 37(3), 1212–1226.
- CHAN, J. C. C., AND E. EISENSTAT (2015): “Marginal Likelihood Estimation with the Cross-Entropy Method,” *Econometric Reviews*, 34(3), 256–285.
- (2018): “Bayesian Model Comparison for Time-Varying Parameter VARs with Stochastic Volatility,” *Journal of Applied Econometrics*, 33(4), 509–532.
- CHAN, J. C. C., E. EISENSTAT, C. HOU, AND G. KOOP (2020): “Composite likelihood methods for large Bayesian VARs with stochastic volatility,” *Journal of Applied Econometrics*, 35(6), 692–711.
- CHAN, J. C. C., E. EISENSTAT, AND R. W. STRACHAN (2020): “Reducing the State Space Dimension in a Large TVP-VAR,” *Journal of Econometrics*, 218(1), 105–118.
- CHAN, J. C. C., AND A. L. GRANT (2016): “Fast Computation of the Deviance Information Criterion for Latent Variable Models,” *Computational Statistics and Data Analysis*, 100, 847–859.
- CHAN, J. C. C., L. JACOBI, AND D. ZHU (2020): “Efficient Selection of Hyperparameters in Large Bayesian VARs Using Automatic Differentiation,” *Journal of Forecasting*, 39(6), 934–943.

- CHAN, J. C. C., AND I. JELIAZKOV (2009): “Efficient Simulation and Integrated Likelihood Estimation in State Space Models,” *International Journal of Mathematical Modelling and Numerical Optimisation*, 1(1), 101–120.
- CHAN, J. C. C., G. KOOP, D. J. POIRIER, AND J. L. TOBIAS (2019): *Bayesian Econometric Methods*. Cambridge University Press, 2 edn.
- CHIB, S. (1995): “Marginal Likelihood from the Gibbs Output,” *Journal of the American Statistical Association*, 90, 1313–1321.
- CHIB, S., AND I. JELIAZKOV (2001): “Marginal Likelihood from the Metropolis-Hastings Output,” *Journal of the American Statistical Association*, 96, 270–281.
- CHIB, S., F. NARDARI, AND N. SHEPHARD (2006): “Analysis of high dimensional multivariate stochastic volatility models,” *Journal of Econometrics*, 134(2), 341–371.
- CLARK, T. E. (2011): “Real-time density forecasts from Bayesian vector autoregressions with stochastic volatility,” *Journal of Business and Economic Statistics*, 29(3), 327–341.
- CLARK, T. E., AND F. RAVAZZOLO (2014): “Macroeconomic Forecasting Performance under alternative specifications of time-varying volatility,” *Journal of Applied Econometrics*, Forthcoming.
- COGLEY, T., AND T. J. SARGENT (2005): “Drifts and volatilities: Monetary policies and outcomes in the post WWII US,” *Review of Economic Dynamics*, 8(2), 262–302.
- CROSS, J., C. HOU, AND A. POON (2019): “Macroeconomic forecasting with large Bayesian VARs: Global-local priors and the illusion of sparsity,” *International Journal of Forecasting*, 36(3), 899–915.
- CROSS, J., AND A. POON (2016): “Forecasting structural change and fat-tailed events in Australian macroeconomic variables,” *Economic Modelling*, 58, 34–51.
- D’AGOSTINO, A., L. GAMBETTI, AND D. GIANNONE (2013): “Macroeconomic forecasting and structural change,” *Journal of Applied Econometrics*, 28, 82–101.
- DEL NEGRO, M., AND G. E. PRIMICERI (2015): “Time-varying structural vector autoregressions and monetary policy: A corrigendum,” *Review of Economic Studies*, 82(4), 1342–1345.
- DEVROYE, L. (2014): “Random variate generation for the generalized inverse Gaussian distribution,” *Statistics and Computing*, 24(2), 239–246.
- DOAN, T., R. LITTERMAN, AND C. SIMS (1984): “Forecasting and conditional projection using realistic prior distributions,” *Econometric reviews*, 3(1), 1–100.

- GELFAND, A. E., AND D. K. DEY (1994): “Bayesian Model Choice: Asymptotics and Exact Calculations,” *Journal of the Royal Statistical Society Series B*, 56(3), 501–514.
- GEWEKE, J. (1999): “Using simulation methods for Bayesian econometric models: Inference, development, and communication,” *Econometric Reviews*, 18(1), 1–73.
- GIANNONE, D., M. LENZA, AND G. E. PRIMICERI (2015): “Prior selection for vector autoregressions,” *Review of Economics and Statistics*, 97(2), 436–451.
- GÖTZ, T., AND K. HAUZENBERGER (2018): “Large mixed-frequency VARs with a parsimonious time-varying parameter structure,” *Deutsche Bundesbank Discussion Paper*.
- HARTWIG, B. (2021): “Bayesian VARs and Prior Calibration in Times of COVID-19,” *Available at SSRN 3792070*.
- HAUTSCH, N., AND S. VOIGT (2019): “Large-scale portfolio allocation under transaction costs and model uncertainty,” *Journal of Econometrics*, 212(1), 221–240.
- HUBER, F., AND M. FELDKIRCHER (2019): “Adaptive shrinkage in Bayesian vector autoregressive models,” *Journal of Business and Economic Statistics*, 37(1), 27–39.
- JIN, X., J. M. MAHEU, AND Q. YANG (2019): “Bayesian parametric and semiparametric factor models for large realized covariance matrices,” *Journal of Applied Econometrics*, 34(5), 641–660.
- JUSTINIANO, A., AND G. E. PRIMICERI (2008): “The Time-Varying Volatility of Macroeconomic Fluctuations,” *American Economic Review*, 98(3), 604–41.
- KADIYALA, R. K., AND S. KARLSSON (1993): “Forecasting with generalized Bayesian vector auto regressions,” *Journal of Forecasting*, 12(3-4), 365–378.
- (1997): “Numerical Methods for Estimation and inference in Bayesian VAR-models,” *Journal of Applied Econometrics*, 12(2), 99–132.
- KASTNER, G. (2019): “Sparse Bayesian time-varying covariance estimation in many dimensions,” *Journal of econometrics*, 210(1), 98–115.
- KASTNER, G., AND F. HUBER (2020): “Sparse Bayesian vector autoregressions in huge dimensions,” *Journal of Forecasting*, 39(7), 1142–1165.
- KAUFMANN, S., AND C. SCHUMACHER (2019): “Bayesian estimation of sparse dynamic factor models with order-independent and ex-post mode identification,” *Journal of Econometrics*, 210(1), 116–134.
- KIM, S., N. SHEPHARD, AND S. CHIB (1998): “Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models,” *Review of Economic Studies*, 65(3), 361–393.

- KOOP, G. (2003): *Bayesian Econometrics*. Wiley & Sons, New York.
- (2013): “Forecasting with medium and large Bayesian VARs,” *Journal of Applied Econometrics*, 28(2), 177–203.
- KOOP, G., AND D. KOROBILIS (2013): “Large time-varying parameter VARs,” *Journal of Econometrics*, 177(2), 185–198.
- KOOP, G., S. MCINTYRE, J. MITCHELL, AND A. POON (2020): “Regional output growth in the United Kingdom: more timely and higher frequency estimates from 1970,” *Journal of Applied Econometrics*, 35(2), 176–197.
- LENZA, M., AND G. E. PRIMICERI (2022): “How to Estimate a VAR after March 2020,” *Journal of Applied Econometrics*, Forthcoming.
- LESAGE, J. P., AND D. HENDRIKZ (2019): “Large Bayesian vector autoregressive forecasting for regions: A comparison of methods based on alternative disturbance structures,” *The Annals of Regional Science*, 62(3), 563–599.
- LI, M., AND M. SCHARTH (2020): “Leverage, Asymmetry, and Heavy Tails in the High-Dimensional Factor Stochastic Volatility Model,” *Journal of Business and Economic Statistics*, pp. 1–17.
- LI, Y., J. YU, AND T. ZENG (2020): “Deviance information criterion for latent variable models and misspecified models,” *Journal of Econometrics*, 216(2), 450–493.
- (2021): “Deviance information criterion for model selection: A theoretical justification,” .
- LITTERMAN, R. (1986): “Forecasting With Bayesian Vector Autoregressions — Five Years of Experience,” *Journal of Business and Economic Statistics*, 4, 25–38.
- LOUZIS, D. P. (2019): “Steady-state modeling and macroeconomic forecasting quality,” *Journal of Applied Econometrics*, 34(2), 285–314.
- MCCAUSLAND, W., S. MILLER, AND D. PELLETIER (2020): “Multivariate stochastic volatility using the HESSIAN method,” *Econometrics and Statistics*.
- MCCRACKEN, M. W., AND S. NG (2020): “FRED-QD: A quarterly database for macroeconomic research,” *Federal Reserve Bank of St. Louis Working Paper*, 2020-005.
- MUMTAZ, H. (2016): “The Evolving Transmission of Uncertainty Shocks in the United Kingdom,” *Econometrics*, 4(1), 16.
- MUMTAZ, H., AND K. THEODORIDIS (2017): “The Changing Transmission of Uncertainty Shocks in the U.S.,” *Journal of Business and Economic Statistics*.

- NG, S. (2021): “Modeling macroeconomic variations after COVID-19,” *National Bureau of Economic Research*.
- PITT, M., AND N. SHEPHARD (1999a): “Time varying covariances: a factor stochastic volatility approach,” *Bayesian Statistics*, 6, 547–570.
- PITT, M. K., AND N. SHEPHARD (1999b): “Filtering via simulation: Auxiliary particle filters,” *Journal of the American Statistical Association*, 94(446), 590–599.
- POON, A. (2018): “Assessing the synchronicity and nature of Australian state business cycles,” *Economic Record*, 94(307), 372–390.
- RUBINSTEIN, R. Y. (1997): “Optimization of computer simulation models with rare events,” *European Journal of Operational Research*, 99, 89–112.
- RUBINSTEIN, R. Y. (1999): “The cross-entropy method for combinatorial and continuous optimization,” *Methodology and Computing in Applied Probability*, 2, 127–190.
- RUBINSTEIN, R. Y., AND D. P. KROESE (2004): *The Cross-Entropy Method: A Unified Approach to Combinatorial Optimization Monte-Carlo Simulation, and Machine Learning*. Springer-Verlag, New York.
- SIMS, C. A. (1980): “Macroeconomics and reality,” *Econometrica*, 48, 1–48.
- SIMS, C. A., D. F. WAGGONER, AND T. ZHA (2008): “Methods for inference in large multiple-equation Markov-switching models,” *Journal of Econometrics*, 146(2), 255–274.
- SPIEGELHALTER, D. J., N. G. BEST, B. P. CARLIN, AND A. VAN DER LINDE (2002): “Bayesian Measures of Model Complexity and Fit,” *Journal of the Royal Statistical Society Series B*, 64(4), 583–639.
- STOCK, J. H., AND M. W. WATSON (2016): “Core inflation and trend inflation,” *Review of Economics and Statistics*, 98(4), 770–784.
- TALLMAN, E. W., AND S. ZAMAN (2020): “Combining survey long-run forecasts and nowcasts with BVAR forecasts using relative entropy,” *International Journal of Forecasting*, 36(2), 373–398.
- ZENS, G., M. BÖCK, AND T. O. ZÖRNER (2020): “The heterogeneous impact of monetary policy on the US labor market,” *Journal of Economic Dynamics and Control*, 119, 103989.
- ZHANG, B., AND B. H. NGUYEN (2020): “Real-time forecasting of the Australian macroeconomy using flexible Bayesian VARs,” *University of Tasmania Discussion Paper Series N 2020-12*.