

BVARs and Stochastic Volatility

Joshua C. C. Chan*

Purdue University

ORCID: 0000-0003-3632-128X

October 2023

Abstract

Bayesian vector autoregressions (BVARs) are the workhorse in macroeconomic forecasting. Research in the last decade has established the importance of allowing time-varying volatility to capture both secular and cyclical variations in macroeconomic uncertainty. This recognition, together with the growing availability of large datasets, has propelled a surge in recent research in building stochastic volatility models suitable for large BVARs. Some of these new models are also equipped with additional features that are especially desirable for large systems, such as order invariance—i.e., estimates are not dependent on how the variables are ordered in the BVAR—and robustness against COVID-19 outliers. Estimation of these large, flexible models is made possible by the recently developed equation-by-equation approach that drastically reduces the computational cost of estimating large systems. Despite these recent advances, there remains much ongoing work, such as the development of parsimonious approaches for time-varying coefficients and other types of nonlinearities in large BVARs.

Keywords: large Bayesian VAR, heavy tail, heteroskedasticity, forecasting, order invariance, COVID-19

*This review is prepared for the *Handbook of Macroeconomic Forecasting* edited by Mike Clements and Ana Galvao. I thank Todd Clark, Frank Wu, Wei Zhang and the editors Mike Clements and Ana Galvao for their helpful comments and suggestions.

1 Introduction

Bayesian vector autoregressions (BVARs) are the workhorse in empirical macroeconomics, especially for forecasting applications and structural analysis. Since the pathbreaking work by Sims (1980) and Doan, Litterman, and Sims (1984) in the early 1980s, there have been many important methodological advances in the BVAR literature. The most notable are the following two developments that motivate a large and expanding body of research in the last decade or so.

First, there is a growing interest in exploiting richer information in macroeconomic analysis, following the seminal paper by Bańbura, Giannone, and Reichlin (2010) that demonstrates the benefits of including a large number of variables in BVARs.¹ Indeed, there are many situations in which it is necessary to consider many variables simultaneously. For example, central banks and policy institutions routinely monitor and forecast dozens of key macroeconomic variables (Crump, Eusepi, Giannone, Qian, and Sbordone, 2021). In applications involving regional time-series, where a region can be a country within an economic union or a state/province within a country, it is often essential to explicitly model the interactions among the regions (Canova and Ciccarelli, 2009; Koop and Korobilis, 2016; Koop, McIntyre, and Mitchell, 2020). For nowcasting applications, it is important to incorporate the flow of data releases in real time, which necessitates a data-rich approach (Giannone, Reichlin, and Small, 2008; Bańbura, Giannone, Modugno, and Reichlin, 2013).

Many other applications also naturally call for the inclusion of a large number of time-series, such as those involving data in multiple frequencies (Schorfheide and Song, 2015; McCracken, Owyang, and Sekhposyan, 2021), disaggregated data (Giannone, Lenza, Momferatou, and Onorante, 2014; Ellahie and Ricco, 2017), firm-level data (Demirer, Diebold, Liu, and Yilmaz, 2018) and financial data (Carriero, Kapetanios, and Marcellino, 2009). Finally, the wide availability of large time-series datasets (McCracken and Ng, 2016, 2020; Baumeister, Leiva-León, and Sims, 2022; Bokun, Jackson, Kliesen, and Owyang, 2023) further propels this development.

¹The earlier paper by Leeper, Sims, and Zha (1996) develops various medium-sized structural VARs to study the effects of monetary policy, including one with 18 variables. Koop (2013) compares the forecast performance of a range of shrinkage priors for BVARs using a dataset containing 168 variables, and finds that BVARs tend to forecast better than factor models.

The second development is the increasing recognition of the need to allow for time-varying volatility in modeling most macroeconomic datasets. Influential papers by Cogley and Sargent (2005), Primiceri (2005) and Sims and Zha (2006) underscore the secular variations in volatility—e.g., many key macroeconomic variables were more volatile in the Great Inflation of the 1970s than the following Great Moderation. At the business cycle frequency, a large body of evidence has shown that macroeconomic volatility is strongly counter-cyclical (see, e.g., Bloom, 2014; Jurado, Ludvigson, and Ng, 2015). Lastly, the unexpected drastic movements in many macroeconomic variables at the onset of the COVID-19 pandemic and the subsequent heightened volatility further underline the need to allow for time-varying volatility.

Within the BVAR literature, papers such as Clark (2011), D’Agostino, Gambetti, and Giannone (2013), Clark and Ravazzolo (2015), Cross and Poon (2016) and Chan and Eisenstat (2018) highlight the empirical relevance of time-varying volatility for model-fit and forecasting in small BVARs. For large BVARs, Carriero, Clark, and Marcellino (2016) and Chan (2023a) present Bayesian model comparison results that show overwhelming data support in favor of stochastic volatility. Koop and Korobilis (2013) and Carriero, Clark, and Marcellino (2019) demonstrate that BVARs with stochastic volatility forecast better than their homoskedastic counterparts. Clark and Mertens (2023) provide a recent review on the wide range of applications of BVARs with stochastic volatility.

The convergence of these two developments drives a surge in recent research in building stochastic volatility models suitable for large BVARs and designing efficient methods for estimating these models. We first introduce in Section 2 three classes of stochastic volatility models—the common stochastic volatility, the Cholesky stochastic volatility and the factor stochastic volatility—that are especially suitable for large BVARs. We then dive into some recent research in developing BVARs with features that are particularly desirable for large systems in Section 3. These include order invariance—i.e., estimation results that are not dependent on how the variables are ordered in the BVAR—and robustness against COVID-19 outliers. Section 4 first reviews Bayesian shrinkage priors, particularly the family of Minnesota priors. It then outlines Markov chain Monte Carlo (MCMC) methods for estimating BVARs with various types of stochastic volatility, followed by a discussion on alternative Bayesian approaches and additional strategies to speed up computations. Section 5 further explores BVARs with time-varying parameters and other types of nonlinearities. Finally, Section 6 highlights a few outstanding

challenges and ongoing research directions.

2 Stochastic Volatility Models for VARs

In this section we discuss a few commonly-used stochastic volatility models in empirical studies involving macroeconomic forecasting and structural analysis. Since there is an increasing demand to incorporate a large number of time-series, we pay special attention to the trade-off between model flexibility and tractability in larger systems.

Let $\mathbf{y}_t = (y_{1,t}, \dots, y_{n,t})'$ denote the n -vector of endogenous variables for $t = 1, \dots, T$. Consider the following VAR with a generic time-varying error covariance matrix Σ_t :

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

where \mathbf{a} is an $n \times 1$ vector of intercepts, $\mathbf{A}_1, \dots, \mathbf{A}_p$ are $n \times n$ coefficient matrices and $\mathbf{0}_n$ denotes an n -vector of zeros. To model volatility clustering—e.g., the empirical observation that large changes tend to be followed by large changes, and small changes followed by small changes—it is important to specify a persistent law of motion for Σ_t . While there are many ways to construct persistent processes, the key challenge here is that for any time t , Σ_t needs to be a positive-definite covariance matrix. In what follows, we discuss various models for Σ_t that satisfy these requirements.

2.1 Common Stochastic Volatility

We first consider the common stochastic volatility model introduced in Carriero, Clark, and Marcellino (2016), which is the very first stochastic volatility model designed for large BVARs. The model is motivated by the observation that the estimated time-varying volatilities of many US macroeconomic variables have broadly similar low-frequency movements: they tend to have high volatility in the 1970s, a drastic decrease in volatility starting in the early 1980s that mark the beginning of the Great Moderation, followed by a new surge in volatility at the onset of the Great Recession of 2007-2009. In particular, Carriero, Clark, and Marcellino (2016) obtain the estimated volatilities of 14 major US macroeconomic time-series from univariate autoregressive models with stochastic volatil-

ity, and find that the first principal component explains about 66% of the variation in the individual volatility time-series.

Motivated by this common component in volatilities, Carriero, Clark, and Marcellino (2016) propose the common stochastic volatility model in which the error covariance matrix is scaled by a common, time-varying factor that may be interpreted as the overall macroeconomic volatility:

$$\Sigma_t = e^{h_t} \Sigma, \quad (2)$$

where Σ is a time-invariant covariance matrix. 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 for identification purposes, the unconditional mean of the AR(1) process is assumed to be zero.

For large systems, estimation of the generic VAR with stochastic volatility in (1) is computationally intensive because of the large number of VAR coefficients $\mathbf{A} = (\mathbf{a}_0, \mathbf{A}_1, \dots, \mathbf{A}_p)'$: it grows quadratically in the number of endogenous variables n . The key advantage of the common stochastic volatility model is that it leads to many useful analytical results that make estimation fast, with computational complexity of the order $\mathcal{O}(n^3)$ —provided that the natural conjugate prior on (\mathbf{A}, Σ) is used (see, e.g., Koop and Korobilis, 2010; Karlsson, 2013, and Chapter 2 of this volume by Hauzenberger, Huber and Koop). Consequently, estimation of BVARs with dozens endogenous variables can be done in minutes instead of hours for other multivariate stochastic volatility models.

The main drawback of the common stochastic volatility model is that it appears to be restrictive along two dimensions. First, fast estimation with computational complexity $\mathcal{O}(n^3)$ is only possible when the natural conjugate prior is used, but the natural conjugate prior can be restrictive as it rules out cross-variable shrinkage (see, e.g. Carriero, Clark, and Marcellino, 2015; Chan, 2022). Second, the common stochastic volatility model assumes in particular that all variances and covariances are scaled by a single factor. Consequently, they are always proportional to each other. Nevertheless, as discussed above, there is empirical evidence that the error variances of many US macroeconomic variables

share broadly similar low-frequency movements, and a common stochastic volatility is a parsimonious way to model this empirical feature.

The common stochastic volatility model in (2)–(3) can be extended in various ways. In particular, Chan (2020) considers a more general family of BVARs that can accommodate non-Gaussian, heteroskedastic and serially dependent errors. In addition, fast estimation of complexity $\mathcal{O}(n^3)$ can be attained for this wide family of BVARs.

Recent empirical applications using this common stochastic volatility model and its extensions include Mumtaz (2016), Mumtaz and Theodoridis (2017), Götz and Hauzenberger (2018), Poon (2018), Louzis (2019), LeSage and Hendrikz (2019), Hartwig (2021), Chan, Poon, and Zhu (2023) and Hou, Nguyen, and Zhang (2023).

2.2 Cholesky Stochastic Volatility

Instead of scaling the error covariance matrix by a common volatility factor, a more flexible approach for modeling time-varying volatilities and correlations is to incorporate multiple stochastic volatility processes. In particular, Cogley and Sargent (2005) build a multivariate stochastic volatility model based on the modified Cholesky decomposition. More specifically, consider again the VAR in (1), but now the error covariance matrix is constructed via

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

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 $h_{i,1} \sim \mathcal{N}(\mu_i, \sigma_i^2 / (1 - \phi_i^2))$. Due to its construction using the modified Cholesky decomposition, this stochastic volatility model is sometimes called the Cholesky stochastic volatility.

Comparing to the common stochastic volatility model in (2), it is clear that the Cholesky stochastic volatility is more flexible—it contains n stochastic volatility processes—and can accommodate more complex co-volatility patterns. In particular, the Cholesky stochastic

volatility can in principle recover the common stochastic volatility specification if we reparameterize the model in (4) so that \mathbf{B}_0 is the lower Cholesky factor of $\boldsymbol{\Sigma}^{-1}$ and $\mathbf{D}_t = e^{h_t} \mathbf{I}_n$, where \mathbf{I}_n is the n -dimensional identity matrix.

This modeling flexibility, however, comes at a cost of more computationally intensive posterior simulations. In particular, the conventional approach of drawing all VAR coefficients jointly, which has $\mathcal{O}(n^6)$ computational complexity, becomes excessively time-consuming when n is large. To tackle this computational problem, Carriero, Clark, and Marcellino (2019) introduce a blocking scheme that makes it possible to estimate the model equation by equation that drastically reduces the computational complexity to $\mathcal{O}(n^4)$; see Section 4 for more details. Nevertheless, the estimation of the Cholesky stochastic volatility model is still an order of magnitude slower than that of the common stochastic volatility (when the natural conjugate prior is used).

There are various useful extensions of the basic setup in (4)-(5). First, instead of the independent autoregressive processes specified in (5), \mathbf{h}_t can be modeled as a VAR(1) to allow for potential correlations among the volatility processes. In the extreme, Chan, Eisenstat, and Strachan (2020) consider a singular VAR for \mathbf{h}_t to capture the idea that only a few common shocks drive the evolution of the n volatility processes. Second, the modified Cholesky factor \mathbf{B}_0 in (4) can be made time-varying, as introduced in Primiceri (2005). In fact, the multivariate stochastic volatility model of Primiceri (2005) is perhaps the most widely used specification in empirical macroeconomics, at least for small systems.

Recent applications using large BVARs with the Cholesky stochastic volatility include Bańbura and van Vlodrop (2018), Bianchi, Guidolin, and Ravazzolo (2018), Huber and Feldkircher (2019), Cross, Hou, and Poon (2019), Koop, McIntyre, Mitchell, and Poon (2020), Tallman and Zaman (2020), Zens, Böck, and Zörner (2020), Baumeister, Korobilis, and Lee (2022), Fischer, Hauzenberger, Huber, and Pfarrhofer (2023) and Gianfreda, Ravazzolo, and Rossini (2023).

2.3 Factor Stochastic Volatility

Another approach to model the law of motion of the error covariance matrix in (1) is through the class of factor stochastic volatility models (Pitt and Shephard, 1999; Aguilar and West, 2000; Chib, Nardari, and Shephard, 2006). A conventional factor

model with homoskedastic factors and idiosyncratic errors is typically used for dimensionality reduction—to model linear relationships among a large number of variables in a very parsimonious way. For instance, an unrestricted $n \times n$ covariance matrix Σ has $n(n+1)/2$ free elements; the number of parameters thus grows quadratically in n . To reduce the number of parameters to be estimated, a factor model in essence specifies Σ using simpler matrices \mathbf{G}, \mathbf{D} and \mathbf{L} via $\Sigma = \mathbf{L}\mathbf{G}\mathbf{L}' + \mathbf{D}$, where \mathbf{G} and \mathbf{D} are $r \times r$ and $n \times n$ diagonal matrices, respectively, and \mathbf{L} is an $n \times r$ matrix, typically with $r \ll n$. Even for an unrestricted \mathbf{L} , the number of parameters in a factor model is only $nr + n + r$; it grows linearly in n . Consequently, for high-dimensional settings with a large n but a small r , a factor model can substantially reduce the number of parameters.

Factor stochastic volatility models extend the classic factor model by introducing time-varying volatility. An example is to specify the error covariance matrix Σ_t as

$$\Sigma_t = \mathbf{L}\mathbf{G}_t\mathbf{L}' + \mathbf{D}_t, \quad (6)$$

where $\mathbf{D}_t = \text{diag}(e^{h_{1,t}}, \dots, e^{h_{n,t}})$ and $\mathbf{G}_t = \text{diag}(e^{h_{n+1,t}}, \dots, e^{h_{n+r,t}})$, and $h_{i,t}, i = 1, \dots, n+r$, follows the AR(1) process in (5). To facilitate estimation, data augmentation is typically used to represent the model in terms of latent factors. More specifically, consider the same VAR in (1), but now the innovation is constructed via

$$\boldsymbol{\varepsilon}_t = \mathbf{L}\mathbf{f}_t + \mathbf{u}_t, \quad (7)$$

where $\mathbf{f}_t = (f_{1,t}, \dots, f_{r,t})'$ is an $r \times 1$ vector of latent factors and \mathbf{L} is the associated $n \times r$ factor loadings matrix. The errors \mathbf{u}_t and the latent factors \mathbf{f}_t are assumed to be independent at all leads and lags and jointly Gaussian:

$$\begin{pmatrix} \mathbf{u}_t \\ \mathbf{f}_t \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \mathbf{0} \\ \mathbf{0} \end{pmatrix}, \begin{pmatrix} \mathbf{D}_t & \mathbf{0} \\ \mathbf{0} & \mathbf{G}_t \end{pmatrix} \right), \quad (8)$$

where $\mathbf{D}_t = \text{diag}(e^{h_{1,t}}, \dots, e^{h_{n,t}})$ and $\mathbf{G}_t = \text{diag}(e^{h_{n+1,t}}, \dots, e^{h_{n+r,t}})$. Under the representation in (7)-(8), it is clear that the correlations among the elements of $\boldsymbol{\varepsilon}_t$ are induced solely by the latent factors.

For identification purposes, it is common to assume \mathbf{L} to be lower triangular with ones on the main diagonal. However, Chan, Eisenstat, and Yu (2022) show that in the presence

of stochastic volatility and $r \leq (n - 1)/2$, the factor loadings matrix \mathbf{L} is identified up to permutations and sign switches (with the restriction that $\mu_i, i = n + 1, \dots, n + r$, the unconditional means of the log-volatilities associated with the factors, are set to 0).²

In practice one also needs to determine the number of factors r . A common approach is to first fix the number of factors r in the estimation, and then select r using the marginal likelihood or some information criterion. An alternative approach is to include a large number of factors, and then use a shrinkage prior on the factor loadings so that they are increasingly shrunk to 0 as the column index increases, as considered in, e.g., Bhattacharya and Dunson (2011) and Kastner (2019). This latter approach circumvents the need to select the number of factors, but it makes the interpretation of the factors more difficult.

Compared to the Cholesky stochastic volatility, the factor stochastic volatility contains r additional volatility processes and is apparently more flexible. However, as discussed earlier, and is clear in the decomposition in (6), the latter inherits the dimensionality reduction property of the classic factor model. As such, whether it fits macroeconomic time-series better than the two other stochastic volatility models is an empirical question. In terms of estimation, the factor stochastic volatility is more computationally intensive to estimate relative to the common stochastic volatility, but it can still be fitted reasonably quickly even when n is large. In particular, given the latent factors, the VAR becomes n unrelated regressions. Consequently, the n -equation system can be estimated equation by equation.

While a variety of factor stochastic volatility models are commonly employed in financial applications, with recent examples in Jin, Maheu, and Yang (2019), Hautsch and Voigt (2019), Kastner (2019), Loaiza-Maya and Smith (2020) and McCausland, Miller, and Pelletier (2020), they are not yet widely used in the context of large BVARs. Notable exceptions are Kastner and Huber (2020) and Hauzenberger, Huber, Koop, and Mitchell (2022).

²To handle the so-called label switching problem (the latent factors are permutations invariant), one can postprocess the posterior draws to sort them into the correct categories using, e.g., the approach in Kaufmann and Schumacher (2019).

3 Order-Invariant and Outlier-Augmented Models

In this section we introduce a variety of recently proposed stochastic volatility models with features that are especially relevant for large systems and for fitting post COVID-19 data.

3.1 Order-Invariant Stochastic Volatility Models

As outlined in Section 2, the law of motion for Σ_t is typically constructed by combining multiple volatility processes using some version of matrix decomposition or factorization. This construction process often leads to order dependence—estimates might depend on the order of variables arranged in the vector \mathbf{y}_t . For instance, the Cholesky stochastic volatility, which is based on a lower triangular parameterization, is not order invariant. In particular, as explained in Carriero, Clark, and Marcellino (2019), since the model is constructed using a lower triangular impact matrix \mathbf{B}_0 and the priors are independently elicited on \mathbf{B}_0 and the volatility processes, the implied prior on Σ_t is not order invariant.

Even though this ordering issue is well-known and is explicitly discussed in Cogley and Sargent (2005) and Primiceri (2005), its empirical relevance has only been appreciated relatively recently. For example, Arias, Rubio-Ramirez, and Shin (2022) use the model of Primiceri (2005) to produce point and density forecasts of 4 macroeconomic variables. They show that while the point forecasts are essentially the same across all the permutations of the 4 variables, the density forecasts can differ substantially across different variable orderings. Since the number of permutations grows exponentially as the number of variables increases, this ordering issue is expected to be more acute in large BVARs involving dozens of dependent variables. In fact, Chan, Doucet, León-González, and Strachan (2018) find evidence that estimates of reduced-form error variances based on the Cholesky stochastic volatility can drastically change across different variable orderings in large systems.³

³The structural analysis in both Cogley and Sargent (2005) and Primiceri (2005) are based on a recursive identification scheme. As such, the order of the variables is part of the identification scheme and is therefore predetermined. However, many later studies (e.g., Benati, 2008; Baumeister and Peersman, 2013) use the BVARs in Cogley and Sargent (2005) and Primiceri (2005) as reduced-form models and implement other, non-recursive identification schemes. In those cases, the reduced-form models do not have a natural ordering. One such example is given in Bognanni (2018), who demonstrates that responses

Given the sensitivity of the estimates and forecasts on different variable orderings, it is therefore desirable to employ order-invariant stochastic volatility models in empirical applications. Which of the stochastic volatility models discussed in Section 2 are order invariant? It is clear that the common stochastic volatility is order invariant, whereas, as discussed earlier, the Cholesky stochastic volatility is not. However, the latter can be extended to be order invariant by a slight modification: avoid the use of the lower triangular parameterization of \mathbf{B}_0 in (4). Based on this observation, Chan, Koop, and Yu (2023) extend the Cholesky stochastic volatility model by allowing the impact matrix \mathbf{B}_0 to be any dense, non-degenerate matrix. They prove that the model is order invariant, and based on the results in Bertsche and Braun (2022), the impact matrix \mathbf{B}_0 is identified up to permutations and sign switches. Moreover, estimation can be done similarly as in the Cholesky stochastic volatility model, particularly that the VAR coefficients can be sampled equation by equation. Chan, Koop, and Yu (2023) also consider a version of the model with a time-varying impact matrix \mathbf{B}_0 , thus extending the model of Primiceri (2005). An alternative approach based on a different factorization of Σ_t is considered in Wu and Koop (2022). More specifically, they construct Σ_t using the eigendecomposition, where the n log-eigenvalues evolve as random walks. They develop a fast posterior sampler to estimate the model and show that it is applicable to large datasets.

Finally, whether a factor stochastic volatility model is order invariant depends on the type of restrictions imposed on the factor loadings. As noted in Kastner (2019), if one imposes the usual identification restrictions that the factor loadings matrix is triangular, then the model is not order invariant. By contrast, Chan, Eisenstat, and Yu (2022) show that, under some mild conditions discussed in the previous section, a factor stochastic volatility model is order invariant and the factor loadings are identified up to permutations and sign switches.

There are other order-invariant modeling approaches, though they are typically designed for small systems. One popular approach is to construct multivariate stochastic volatility models using Wishart or inverse-Wishart processes; examples include Philipov and Glickman (2006) and Asai and McAleer (2009). A few papers extend these multivariate stochastic volatility models to BVARs, including Chan, Doucet, León-González, and Strachan (2018), Shin and Zhong (2020) and Arias, Rubio-Ramirez, and Shin (2022).

to structural shocks identified with sign restrictions—using the reduced-form BVAR of Primiceri (2005)—can be sensitive to how the variables are ordered; see also Hartwig (2020) for another example.

More recently, Arias, Rubio-Ramirez, and Shin (2022) introduce a new order-invariant approach that builds on the new parameterization of correlation matrices proposed in Archakov and Hansen (2021). More specifically, they construct a time-varying correlation matrix based on random walk processes under the new parameterization. However, estimation of these stochastic volatility models is generally computationally intensive, as it often involves drawing from high-dimensional, non-standard distributions. Consequently, these stochastic volatility models do not scale well to large datasets.

Another order-invariant approach is based on the discounted Wishart process, which admits efficient filtering and smoothing algorithms for estimation. Stochastic volatility models constructed under this approach are considered in Uhlig (1997), West and Harrison (2006), Prado and West (2010), Bognanni (2018) and Arias, Rubio-Ramirez, and Shin (2022). However, the discounted Wishart process appears to be too tightly parameterized for typical macroeconomic data, and it tends to underperform in terms of both point and density forecasts relative to standard stochastic volatility models such as Cogley and Sargent (2005) and Primiceri (2005), as demonstrated in a forecasting exercise in Arias, Rubio-Ramirez, and Shin (2022).

3.2 Outlier-Augmented Stochastic Volatility Models

The COVID-19 pandemic has caused extreme movements in many macroeconomic and financial time-series. One consequence is that impulse response functions and forecasts from homoskedastic BVARs are heavily distorted by these extreme observations, as demonstrated in Schorfheide and Song (2021) and Lenza and Primiceri (2022) using US data and Bobeica and Hartwig (2023) using euro area data. This problem can be ameliorated by using BVARs with stochastic volatility, which downweight extreme observations. However, as pointed out by Carriero, Clark, Marcellino, and Mertens (2022), extreme observations, by definition, should reflect transitory spikes, not permanent increases, in volatility. But in a typical stochastic volatility, changes in volatility are assumed to be highly persistent. It is therefore useful to augment standard stochastic volatility models to include an explicit component to model infrequent volatility spikes.

There is a general approach to directly model infrequent volatility spikes, and it can be incorporated into the many multivariate stochastic volatility models discussed earlier.

The key idea is to replace the normal distribution in (1) with a more robust distribution that puts more mass on extreme events. A common example is a continuous heavy-tailed distribution, such as the Student- t distribution. As long as the robust distribution can be represented as a scale, finite or an infinite mixture of normals, estimation is straightforward thanks to data augmentation and the modular nature of MCMC methods.

As an example, take the Cholesky stochastic volatility model with a generalized t distribution considered in Clark and Ravazzolo (2015), Cross and Poon (2016) and Chiu, Mumtaz, and Pinter (2017), which is constructed from n univariate Student- t distributions with different degree of freedom parameters.⁴ It can be represented as a conditionally Gaussian VAR specified in (1) with Σ_t constructed as

$$\Sigma_t = \mathbf{B}_0^{-1} \mathbf{\Lambda}_t \mathbf{D}_t (\mathbf{B}_0^{-1})'. \quad (9)$$

Compared to the standard Cholesky stochastic volatility in (4), the new addition is the diagonal matrix of latent variables $\mathbf{\Lambda}_t = \text{diag}(\lambda_{1,t}, \dots, \lambda_{n,t})$, where $\lambda_{i,t} \sim \mathcal{IG}(\nu_i/2, \nu_i/2)$ and $\mathcal{IG}(a, b)$ denotes the inverse-gamma distribution with mean $b/(a - 1)$ when $a > 1$.⁵

Another example is the outlier-augmented stochastic volatility model proposed by Carriero, Clark, Marcellino, and Mertens (2022), which extends the Cholesky stochastic volatility by incorporating a discrete mixture representation first introduced in Stock and Watson (2016) for handling outliers in unobserved components models. More specifically, it has the same conditionally Gaussian VAR representation with Σ_t specified in (9), but the latent variable $\lambda_{i,t}$ is specified as $\lambda_{i,t} = \sigma_{i,t}^2$, where $\sigma_{i,t}$ follows a 2-part distribution with a point mass at 1 and a uniform distribution on the interval (2, 20). The first part of the distribution represents ‘regular’ observations with scale normalized to 1, whereas the second part captures ‘outliers’ that have 2-20 times larger standard deviations relative to regular observations.⁶

⁴A similar heavy-tailed stochastic volatility model is introduced by Cúrdia, Del Negro, and Greenwald (2014) to construct a dynamic stochastic general equilibrium model that can accommodate extreme events during the Great Recession.

⁵For the special case with $\lambda_{1,t} = \dots = \lambda_{n,t} = \lambda_t \sim \mathcal{IG}(\nu/2, \nu/2)$, the marginal distribution of the innovation $\boldsymbol{\varepsilon}_t$ (unconditional on the latent variable λ_t) is a Student- t distribution with degree of freedom parameter ν (see, e.g., Geweke, 1993; Ni and Sun, 2005).

⁶In the original setup in Stock and Watson (2016), the range of the uniform distribution is from 2 to 10. As noted in Carriero, Clark, Marcellino, and Mertens (2022), many macroeconomic variables exhibit extreme variability at the onset of the COVID-19 pandemic, and the upper range is accordingly increased to 20.

Other distributional assumptions for $\lambda_{i,t}$ are possible and they would induce different marginal distributions on the innovations. Examples include the multivariate Laplace distribution (Eltoft, Kim, and Lee, 2006), a finite mixture of scale mixtures of normals (Karlsson and Mazur, 2020) and an infinite mixture of normals (Braun, 2021). These non-Gaussian distributions can also be incorporated in other stochastic volatility models, as is done in Chan (2020) and Hartwig (2021).

4 Shrinkage Priors and Posterior Simulations

What makes BVARs Bayesian is the use of informative priors on the large number of VAR coefficients. Below we first outline a family of priors generally referred to as the Minnesota priors. We then discuss Bayesian estimation of BVARs with various types of stochastic volatility introduced earlier, particularly the posterior simulation of the log-volatilities and the VAR coefficients.

4.1 Shrinkage Priors on the VAR Coefficients

We first discuss a family of priors on the VAR coefficients that can be traced back to the work in the 1980s by Doan, Litterman, and Sims (1984) and Litterman (1986). Due to their affiliations at that time, this family of priors is commonly known as the Minnesota priors. In the original formulation for homoskedastic VARs, the error covariance matrix is fixed at some estimate, and a normal prior is elicited on the VAR coefficients $\boldsymbol{\alpha} = \text{vec}(\mathbf{A}')$. Later Kadiyala and Karlsson (1993, 1997) extend the original version to a joint prior on $\boldsymbol{\alpha}$ and the time-invariant error covariance matrix. Since then many other more flexible variants have been developed, for both homoskedastic and heteroskedastic VARs, and they can be written as a conditionally normal prior given some latent variables or hyperparameters $\boldsymbol{\xi}$:

$$(\boldsymbol{\alpha} | \boldsymbol{\xi}) \sim \mathcal{N}(\boldsymbol{\alpha}_0, \mathbf{V}), \quad (10)$$

where the prior covariance matrix \mathbf{V} may depend on $\boldsymbol{\xi}$. The prior mean $\boldsymbol{\alpha}_0$ is typically set to zero for growth rates data, such as GDP growth. This reflects the prior belief that growth rates data are generally not very persistent, and the VAR coefficients are thus shrunk to 0. For levels data such as private investment or inflation indices, $\boldsymbol{\alpha}_0$ is set to

be zero except for the coefficients associated with the first own lag, which are set to be one. This reflects the prior belief that levels data are highly persistent—in particular, it expresses the preference for a random walk specification.

In early versions of the Minnesota priors there are no latent variables and the prior hyperparameters are fixed at some subjective values. The prior covariance matrix \mathbf{V} is typically assumed to be diagonal, and the diagonal elements are elicited to reflect a range of subjective beliefs, such as cross-variable shrinkage—i.e., the idea that coefficients on other variables’ lags are on average smaller than coefficients on own lags. This is often implemented by introducing two hyperparameters, say, ξ_1 and ξ_2 , where ξ_1 controls the overall prior variance for coefficients on own lags, whereas ξ_2 for coefficients on other lags. Then, ξ_1 and ξ_2 are carefully calibrated to incorporate the belief that ξ_2 is much smaller than ξ_1 .

Naturally, one set of fixed values for ξ_1 and ξ_2 is not expected to be optimal for all datasets with widely different time-series and sample periods. Indeed, Giannone, Lenza, and Primiceri (2015) demonstrate that substantial benefits can be obtained—e.g., better out-of-sample forecast performance using real data and more accurate impulse response functions using simulated data—simply by estimating these prior hyperparameters from the data instead of fixing them at some commonly-used subjective values. In a Bayesian model comparison exercise, Chan (2023a) finds strong evidence in favor of estimating these prior hyperparameters across BVARs with a variety of stochastic volatility specifications.⁷

There is a surge of interest in the statistics literature to develop hierarchical shrinkage priors, such as the Bayesian Lasso (Park and Casella, 2008), the normal-gamma prior (Griffin and Brown, 2010), the horseshoe prior (Carvalho, Polson, and Scott, 2010) and the Dirichlet-Laplace prior (Bhattacharya, Pati, Pillai, and Dunson, 2015). These shrinkage priors are originally designed for the setting of a linear regression with a large number of arbitrary predictors. More recently, they have been introduced in BVAR settings. Applications include Huber and Feldkircher (2019), Follett and Yu (2019), Kastner and

⁷Earlier papers such as Del Negro and Schorfheide (2004) and Carriero, Kapetanios, and Marcellino (2012) have adopted an empirical Bayes approach of selecting these prior hyperparameters by maximizing the marginal likelihood of the homoskedastic BVAR, which is available analytically. For time-varying parameter models, Amir-Ahmadi, Matthes, and Wang (2020) show similar benefits in treating the prior hyperparameters as unknown parameters to be estimated.

Huber (2020), Korobilis and Pettenuzzo (2019) and Gefang, Koop, and Poon (2023). In contrast to Minnesota priors that tend to shrink all VAR coefficients, these adaptive hierarchical priors have the desirable theoretical property of only shrinking ‘small’ coefficients strongly to zero, while leaving ‘large’ coefficients mostly intact.

In practice, however, these new hierarchical priors do not seem to forecast better than some sophisticated Minnesota priors, as demonstrated in Cross, Hou, and Poon (2019). One reason for this surprising result could be because the new hierarchical priors typically treat all VAR coefficients identically, and do not take into account the many plausible prior beliefs, such as cross-variable shrinkage and that variables of higher lags are less important, incorporated into the Minnesota priors. There are a few papers that aim to adapt the new adaptive hierarchical priors so that they are more suitable for BVARs. For instance, Huber and Feldkircher (2019) develop a normal-gamma prior in which additional lag-specific shrinkage parameters are introduced with the goal of shrinking higher-order lags more strongly to zero. Korobilis and Pettenuzzo (2019) consider a class of hierarchical shrinkage priors that incorporate some features of the Minnesota priors. Chan (2021) develops a framework that explicitly nests both the Minnesota priors and the new adaptive hierarchical priors. These Minnesota-type adaptive hierarchical priors can be written as conditionally Gaussian priors with additional latent variables ξ given in (10). For example, under the Minnesota-type normal-gamma prior ξ becomes a vector of gamma random variables; for the Minnesota-type horseshoe prior, elements of ξ have a half-Cauchy distribution. Chan (2021) shows that these Minnesota-type adaptive hierarchical priors forecast better than both the conventional Minnesota priors and the standard adaptive hierarchical priors.

The literature on shrinkage priors on VAR coefficients is vast and expanding. Here we have focused on the family of Minnesota priors. More discussions on other shrinkage priors and approaches can be found in Chapter 2 of this volume by Hauzenberger, Huber and Koop and excellent reviews by Koop and Korobilis (2010) and Karlsson (2013).

4.2 Bayesian Estimation

Estimating BVARs with stochastic volatility tends to be computationally intensive, especially for large systems. But a lot of progress has been made in the last decade, and

it is now feasible to fit large BVARs with a variety of stochastic volatility specifications using MCMC methods. Below we discuss the posterior simulation of the two most computationally intensive parts: sampling the log-volatilities and the VAR coefficients.

4.2.1 Posterior Simulation of the Log-Volatilities

Conditional on the data and other model parameters, the VAR in (1) defines the observation equation for the log-volatilities. But since the log-volatilities enter the model via the time-varying covariance matrix, the state space representation is nonlinear and high-dimensional, which makes estimation very challenging in general. However, many multivariate stochastic volatility models, such as common, Cholesky and factor stochastic volatility, are constructed in such a way that they can be transformed into a linear (though non-Gaussian) state space model, for which the auxiliary mixture sampler of Kim, Shephard, and Chib (1998) can be applied to sample the log-volatilities efficiently.

More specifically, after transforming the original nonlinear Gaussian representation into a linear non-Gaussian state space model, Kim, Shephard, and Chib (1998) observe that the non-Gaussian errors in the new observation equation can be well approximated using a 7-component mixture of normals. Then, given the mixture indicators, the approximate model becomes a linear Gaussian state space model for which sampling the latent states can be done using standard algorithms, such as the Kalman-filter based smoothers of Carter and Kohn (1994) and Durbin and Koopman (2002) or the precision-based samplers of Chan and Jeliazkov (2009) and McCausland, Miller, and Pelletier (2011). Estimation details of the common, Cholesky and factor stochastic volatility models can be found in Chan (2023a). Other stochastic volatility models with additional features can also be estimated using the auxiliary mixture sampler with minor modifications; see, for example, Carriero, Clark, Marcellino, and Mertens (2022) and Chan, Koop, and Yu (2023).

One important point to emphasize is that, as pointed out in Del Negro and Primiceri (2015), when the auxiliary mixture sampler is used as part of a larger posterior simulator, the latter becomes a collapsed Gibbs sampler, where the augmented mixture indicators and all other model parameters should be sampled jointly. This joint sampling can be done by first drawing the model parameters given the log-volatilities but marginal of the mixture indicators, followed by sampling the mixture indicators given the log-volatilities

and other model parameters. In contrast to a standard Gibbs sampler, here the sampling order is important.⁸ In practice, the main implication is that the log-volatilities should be sampled immediately after the mixture indicators.⁹

For other stochastic volatility models for which the auxiliary mixture sampler cannot be applied, there are two alternatives. While they are both very general and can be applied to a wide range of stochastic volatility models, they are more computationally demanding. The first is to use sequential Monte Carlo (SMC) or particle filters. There is a large and rapidly expanding literature on sequential Monte Carlo methods; Doucet, De Freitas, and Gordon (2001) and Creal (2012) provide excellent reviews on this large family of algorithms. The most widely used member is perhaps the bootstrap particle filter, where the state equation is used to generate particle values from time t to values at time $t + 1$. While the bootstrap particle filter is easy to implement, its empirical performance is often poor. The main reason is that the generated particle values from the state equation are unlikely to be consistent with the data in typical settings. There are a number of more sophisticated filters that use information from the observation equation to construct better proposal distributions. A recent example is the tempered particle filter proposed in Herbst and Schorfheide (2019) that adaptively constructs the proposal distribution through a sequence of tempering steps.

The second general approach is to implement a Metropolis-Hastings step to sample the log-volatilities. This can be done by using either a random-walk Metropolis-Hastings algorithm as in Asai and McAleer (2009), or an independence-chain Metropolis-Hastings algorithm with a tailored proposal distribution. The first option is straightforward to implement, but is often very inefficient in terms of highly autocorrelated posterior draws, especially in high-dimensional settings. The second option requires a good proposal distribution that well approximates the conditional distribution of the log-volatilities. As such, it demands more analysis and is therefore more difficult to implement. However,

⁸It is also feasible to apply a standard Gibbs sampler by conditioning on the mixture indicators in all the Gibbs steps. This approach is typically more cumbersome as, e.g., the VAR in (1) is no longer Gaussian given the mixture indicators. See Stock and Watson (2016) for an example of this alternative approach.

⁹For stochastic volatility models with multiple time-series of log-volatilities, where each series has an associated vector of mixture indicators, it is often more convenient to sample each pair of mixture indicators and log-volatilities separately. This approach is valid if the time-series of log-volatilities are conditionally independent given the model parameters—an example is the Cholesky stochastic volatility model specified in (4)-(5). If they are not conditionally independent, then all the mixture indicators should be sampled first before the log-volatilities.

there has been a lot of progress in developing good proposal distributions. For example, McCausland (2012) introduces the HESSIAN method based on a fifth-order Taylor approximation that works remarkably well for a range of univariate stochastic volatility models. Chan and Eisenstat (2018) and Cross, Hou, Koop, and Poon (2023) instead consider quadratic (hence, Gaussian) approximations of the target distributions in the context of BVARs with stochastic volatility, which are easier to obtain and work well for their applications.¹⁰ In addition, sequential Monte Carlo methods can also be used to construct good proposal distributions; see, e.g., Andrieu, Doucet, and Holenstein (2010).

4.2.2 Posterior Simulation of the VAR Coefficients

Given the conditionally Gaussian VAR in (1) and the conditionally Gaussian prior on the VAR coefficients α in (10), the conditional posterior distribution of α is also Gaussian by standard linear regression results. Hence, in principle sampling α is straightforward. In practice, however, when the number of endogenous variables in the VAR is large, conventional ways of drawing all VAR coefficients jointly—e.g., by first obtaining the Cholesky factor of the covariance matrix or the precision matrix—become excessively computationally intensive. In particular, for an n -variable VAR with p lags, the number of VAR coefficients (ignoring intercepts) is n^2p . Since the covariance matrix of the joint distribution is of dimension $n^2p \times n^2p$, obtaining its Cholesky factor has computational complexity of the order $\mathcal{O}(n^6p^3)$.¹¹

To tackle this computational problem, Carriero, Clark, and Marcellino (2019) consider an alternative sampling scheme of drawing the VAR coefficients equation by equation that can drastically reduce the computational complexity by two orders of magnitude to $\mathcal{O}(n^4p^3)$.¹² The key reason for the speed-up is as follows. Since the number of VAR

¹⁰This approach of using Gaussian approximations is first introduced to estimate univariate stochastic volatility models in the seminal papers by Durbin and Koopman (1997) and Shephard and Pitt (1997). They formulate the approximating algorithm via a linear Gaussian state space model, which implies a Gaussian density for the states. Chan and Eisenstat (2018) and Cross, Hou, Koop, and Poon (2023) improve upon this approach by directly computing the Gaussian approximating density using Newton-Raphson method based on fast band matrix routines.

¹¹Sampling the VAR coefficients using the Cholesky factor of the precision matrix is faster, as it avoids explicitly inverting the precision matrix to obtain the covariance matrix. But since the precision matrix is dense in this case, its computational complexity is of the same order of $\mathcal{O}(n^6p^3)$.

¹²As noted in Bognanni (2022), the algorithm in Carriero, Clark, and Marcellino (2019) can only be viewed as an approximation. Carriero, Chan, Clark, and Marcellino (2022) provide an exact algorithm

coefficients in each equation is np , drawing these coefficients conditional on coefficients in other equations can be done in $\mathcal{O}(n^3p^3)$ operations. Iterating this step over n equations, the total computational complexity is therefore $\mathcal{O}(n^4p^3)$. This equation-by-equation estimation algorithm is formulated for the Cholesky stochastic volatility model, but it can be used for other types of stochastic volatility or homoskedastic VARs with minor modifications.¹³ Carriero, Chan, Clark, and Marcellino (2022) report a 10-50 times speed-up of the equation-by-equation approach for n up to 40 compared to sampling all coefficients jointly.

Another venue to improve sampling speed is to find faster ways to simulate from high-dimensional Gaussian distributions. For a linear regression with r observations and m covariates, the conventional, Cholesky-factor based sampling approach to simulate the m -vector of coefficients requires $\mathcal{O}(m^3)$ operations. Bhattacharya, Chakraborty, and Mallick (2016) propose an alternative way to sample the coefficients based on the Woodbury formula, which has computational complexity of $\mathcal{O}(r^2m)$. As such, this new algorithm is expected to be faster in cases where there are far more covariates than observations. For our BVAR setting with $r = Tn$ and $m = np$ (sampling the coefficients in one equation), the new algorithm has computational complexity of $\mathcal{O}(T^2n^3p)$ compared to the conventional sampling approach of $\mathcal{O}(n^3p^3)$. Both have the same order of complexity in n , but they have different complexity in other dimensions. For instance, when T is small and p is large, the algorithm of Bhattacharya, Chakraborty, and Mallick (2016) is expected to perform better.

4.2.3 Alternative Bayesian Approaches

So far we have focused on various MCMC methods to estimate the reduced-form VAR given in (1). There is a range of additional strategies to speed up computations so that one can fit larger BVARs with more flexible features.

One approach is to reparameterize the standard reduced-form VAR in the structural form with a recursive system (see, e.g., Chan, 2021; Gefang, Koop, and Poon, 2023). The

to draw the VAR coefficients equation by equation that has the same order of computational complexity.

¹³For some special cases, other faster algorithms are available. For example, under the natural conjugate prior, the common stochastic volatility model and its variants can be estimated using algorithms with complexity $\mathcal{O}(n^3p^3)$; see Carriero, Clark, and Marcellino (2016) and Chan (2020).

structural-form parameter estimates can then be transformed to obtain the corresponding reduced-form estimates. The main advantage of this approach is that the n -variable VAR can be directly written as a system of n unrelated univariate regressions, which can be estimated separately. This reparameterization can offer up to 10 times speed up compared to estimation of the reduced form, depending on the size of the system. Moreover, further speed-up can be achieved via parallelization. The key drawback of this approach is that to apply standard conjugacy results, one needs to use normal priors on the structural-form VAR coefficients, which often imply order-dependent priors on the reduced-form parameters. One way to ameliorate this ordering issue is to first elicit prior beliefs on the reduced-form parameters, and then work out the implied prior beliefs on the structural-form parameters; see, e.g., Chan (2022) for an example.

There are other Bayesian estimation approaches that offer substantial reduction in computation time compared to conventional MCMC methods. For example, Bognanni and Zito (2020) introduce a sequential Monte Carlo (SMC) algorithm for estimating BVARs with stochastic volatility that leverages parallelization. In addition, the algorithm is especially suitable for online learning tasks, such as out-of-sample forecasting, where new data can be incorporated to update the posterior distribution without completely reestimating the model.

While MCMC and SMC are exact methods in the sense that they can approximate the posterior distribution arbitrarily well by increasing the simulation size, there are approximate methods available in settings where neither MCMC nor SMC are computationally feasible. One prominent example is the fast growing collection of variational Bayesian methods, which are deterministic algorithms for approximating the posterior distribution using a more tractable distribution. More specifically, given a family of tractable densities, the variational Bayesian approach locates the optimal density within this family by minimizing the Kullback-Leibler divergence of the approximating density to the posterior density. Key references on variational Bayesian methods can be found in Jordan, Ghahramani, Jaakkola, and Saul (1999), Bishop (2006) and Ormerod and Wand (2010).

Hajargasht and Woźniak (2018) appear to be the first to apply the variational Bayesian approach to the estimation of a 7-variable homoskedastic VAR. Chan and Yu (2022) and Gefang, Koop, and Poon (2023) develop variational Bayesian methods for estimating large BVARs with stochastic volatility. In particular, Chan and Yu (2022) use a 96-

variable BVAR to measure global bank network connectedness, whereas Gefang, Koop, and Poon (2023) conduct an out-of-sample forecasting exercise using a large data set of 100 variables. While both Chan and Yu (2022) and Gefang, Koop, and Poon (2023) parameterize their BVARs in the structural form, Bernardi, Bianchi, and Bianco (2023) develop a variational Bayes algorithm to estimate a reduced-form BVAR with stochastic volatility. Hierarchical shrinkage priors on the reduced-form VAR coefficients seem to perform better than those elicited on the structural form, but estimation under the reduced-form parameterization is slightly more costly.

5 Time-Varying Parameter and Nonlinear VARs

Since the pioneering works of Cogley and Sargent (2001, 2005) and Primiceri (2005), BVARs with time-varying VAR coefficients have been widely used to model time-series with structural changes. In these time-varying parameter (TVP) models, it is crucial to also allow for time-varying volatility—as highlighted in Sims (2001), a failure to account for heteroskedasticity in a TVP model with constant error variances might give an appearance of time variation in the VAR coefficients. In fact, Sims and Zha (2006) compare various regime-switching BVARs and find that the best model is the one that allows for time variation in error variances only. Chan and Eisenstat (2018) report similar findings in BVARs with TVP and stochastic volatility.

Of course, TVP models are but one nonlinear approach to model the conditional mean in BVARs. For applications that study the impact of macroeconomic or financial uncertainty on key macroeconomic variables, BVARs with stochastic volatility in mean are an especially convenient framework. For example, Mumtaz and Zanetti (2013) develop a TVP-VAR with stochastic volatility where the log-volatilities and their lags enter the conditional mean as regressors. Using this new model, they then study the impact of monetary shocks volatility. An alternative approach based on a large BVAR with a multiplicative factor stochastic volatility model is proposed in Carriero, Clark, and Marcellino (2018). In their model the log-volatilities are driven by common factors representing macroeconomic and financial uncertainty. The two common factors and their lags then enter the conditional mean in the BVAR as regressors. Estimation of these stochastic volatility in mean models is typically more computationally intensive, as the auxiliary

mixture sampler does not apply. In view of this difficulty, Cross, Hou, Koop, and Poon (2023) develop an efficient precision-based sampler for posterior and predictive inference in large stochastic volatility in mean BVARs. They report substantial computational and statistical efficiency gains over a standard particle filter.

There are also a few nonparametric approaches designed for BVARs. For example, Hauzenberger, Huber, Marcellino, and Petz (2021) develop a nonparametric BVAR using a Gaussian process prior to model the relationship of each endogenous variable on its own lags and lags of all other variables. The model also features a multivariate stochastic volatility specification similar to the Cholesky stochastic volatility. An alternative nonparametric approach based on Bayesian additive regression trees is considered in Hauzenberger, Huber, Koop, and Mitchell (2022) and Clark, Huber, Koop, Marcellino, and Pfarrhofer (2022). Specifically, a TVP-VAR with stochastic volatility is first represented in the structural form, and in each equation both the time-varying conditional mean and the conditional variance are modeled using Bayesian additive regression trees. Both find forecast accuracy gains over linear BVARs. For more detailed discussions on Bayesian nonparametric methods, see Chapter 5 of this volume by Marcellino and Pfarrhofer.

6 The Road Ahead

Over the last decade substantial progress has been made in developing a wide range of stochastic volatility models and efficient estimation approaches for large datasets. There remains, however, a lot of ongoing work in building more flexible time-varying BVARs. One direction tackles the problem of modeling time variation in large systems, since standard TVP models of Cogley and Sargent (2005) and Primiceri (2005) cannot be directly applied to those settings due to computational and overparameterization concerns. In addition, allowing all VAR coefficients to be time-varying appears to be unnecessary. For example, Chan (2023b) finds evidence that in a large BVAR, the coefficients in some, but not all, equations are time-varying, and forcing all coefficients to be time-varying worsens forecast performance. Papers such as Koop and Korobilis (2013, 2018) propose fast estimation methods to approximate the posterior distributions of large TVP models. Huber, Koop, and Onorante (2019) develop an algorithm that first shrinks the time-varying

coefficients, followed by setting the small values to zero. The nonparametric approach based on Bayesian additive regression trees developed in Hauzenberger, Huber, Koop, and Mitchell (2022) can also accommodate time-varying coefficients.

Another research direction involves developing more general VAR models able to accommodate the augmentation of typical macroeconomic datasets with additional types of data. A promising recent approach is the functional VAR developed in Chang, Chen, and Schorfheide (2021) that jointly models standard macroeconomic time-series and individual-level cross-sectional data. This framework is especially suitable for studying, for example, the distributional effects of various structural macroeconomic shocks.

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.
- ANDRIEU, C., A. DOUCET, AND R. HOLENSTEIN (2010): “Particle Markov chain Monte Carlo methods,” *Journal of the Royal Statistical Society Series B*, 72(3), 269–342.
- ARCHAKOV, I., AND P. HANSEN (2021): “A new parametrization of correlation matrices,” *Econometrica*, 89(4), 1699–1715.
- ARIAS, J. E., J. F. RUBIO-RAMIREZ, AND M. SHIN (2022): “Macroeconomic forecasting and variable ordering in multivariate stochastic volatility models,” *Journal of Econometrics*, forthcoming.
- ASAI, M., AND M. MCALEER (2009): “The structure of dynamic correlations in multivariate stochastic volatility models,” *Journal of Econometrics*, 150(2), 182–192.
- BAÑBURA, 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.
- BAÑBURA, M., D. GIANNONE, AND L. REICHLIN (2010): “Large Bayesian vector autoregressions,” *Journal of Applied Econometrics*, 25(1), 71–92.
- BAÑBURA, 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 (2022): “Energy markets and global economic conditions,” *Review of Economics and Statistics*, 104(4), 828–44.
- BAUMEISTER, C., D. LEIVA-LEÓN, AND E. SIMS (2022): “Tracking weekly state-level economic conditions,” *Review of Economics and Statistics*.
- BAUMEISTER, C., AND G. PEERSMAN (2013): “Time-varying effects of oil supply shocks on the US economy,” *American Economic Journal: Macroeconomics*, 4(5), 1–28.
- BENATI, L. (2008): “The “great moderation” in the United Kingdom,” *Journal of Money, Credit and Banking*, 40(1), 121–147.

- BERNARDI, M., D. BIANCHI, AND N. BIANCO (2023): “Variational inference for large Bayesian vector autoregressions,” *Available at SSRN: <http://dx.doi.org/10.2139/ssrn.3777064>*.
- BERTSCHE, D., AND R. BRAUN (2022): “Identification of structural vector autoregressions by stochastic volatility,” *Journal of Business & Economic Statistics*, 40(1), 328–341.
- BHATTACHARYA, A., A. CHAKRABORTY, AND B. K. MALLICK (2016): “Fast sampling with Gaussian scale mixture priors in high-dimensional regression,” *Biometrika*, 103, 985–991.
- BHATTACHARYA, A., AND D. DUNSON (2011): “Sparse Bayesian infinite factor models,” *Biometrika*, pp. 291–306.
- BHATTACHARYA, A., D. PATI, N. S. PILLAI, AND D. B. DUNSON (2015): “Dirichlet–Laplace priors for optimal shrinkage,” *Journal of the American Statistical Association*, 110(512), 1479–1490.
- 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.
- BISHOP, C. M. (2006): *Pattern Recognition and Machine Learning*. springer.
- BLOOM, N. (2014): “Fluctuations in uncertainty,” *Journal of economic Perspectives*, 28(2), 153–176.
- BOBEICA, E., AND B. HARTWIG (2023): “The COVID-19 shock and challenges for inflation modelling,” *International Journal of Forecasting*, 39(1), 519–539.
- BOGNANNI, M. (2018): “A class of time-varying parameter structural VARs for inference under exact or set identification,” *FRB of Cleveland Working Paper*.
- (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.
- BOKUN, K. O., L. E. JACKSON, K. L. KLIESEN, AND M. T. OWYANG (2023): “FRED-SD: A real-time database for state-level data with forecasting applications,” *International Journal of Forecasting*, 39(1), 279–297.
- BRAUN, R. (2021): “The importance of supply and demand for oil prices: evidence from non-Gaussianity,” *Bank of England Working Paper*.

- CANOVA, F., AND M. CICCARELLI (2009): “Estimating multicountry VAR models,” *International economic review*, 50(3), 929–959.
- 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, AND M. MARCELLINO (2018): “Measuring uncertainty and its impact on the economy,” *Review of Economics and Statistics*, 100(5), 799–815.
- 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.
- CARRIERO, A., G. KAPETANIOS, AND M. MARCELLINO (2012): “Forecasting government bond yields with large Bayesian vector autoregressions,” *Journal of Banking & Finance*, 36(7), 2026–2047.
- CARTER, C. K., AND R. KOHN (1994): “On Gibbs Sampling for State Space Models,” *Biometrika*, 81, 541–553.
- CARVALHO, C. M., N. G. POLSON, AND J. G. SCOTT (2010): “The horseshoe estimator for sparse signals,” *Biometrika*, 97(2), 465–480.
- 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.
- (2022): “Asymmetric conjugate priors for large Bayesian VARs,” *Quantitative Economics*, 13(3), 1145–1169.

- (2023a): “Comparing stochastic volatility specifications for large Bayesian VARs,” *Journal of Econometrics*, 235(2), 1419–1446.
- (2023b): “Large hybrid time-varying parameter VARs,” *Journal of Business and Economic Statistics*, 41(3), 890–905.
- CHAN, J. C. C., A. DOUCET, R. LEÓN-GONZÁLEZ, AND R. W. STRACHAN (2018): “Multivariate stochastic volatility with co-heteroscedasticity,” *CAMA Working Paper*.
- CHAN, J. C. C., AND E. EISENSTAT (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, 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., E. EISENSTAT, AND X. YU (2022): “Large Bayesian VARs with factor stochastic volatility: Identification, order invariance and structural analysis,” *arXiv preprint arXiv:2207.03988*.
- 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, AND X. YU (2023): “Large order-invariant Bayesian VARs with stochastic volatility,” *Journal of Business and Economic Statistics*, forthcoming.
- CHAN, J. C. C., A. POON, AND D. ZHU (2023): “High-dimensional conditionally Gaussian state space models with missing data,” *Journal of Econometrics*, 236(1), 105468.
- CHAN, J. C. C., AND X. YU (2022): “Fast and accurate variational inference for large Bayesian VARs with stochastic volatility,” *Journal of Economic Dynamics and Control*, 143, 104505.
- CHANG, M., X. CHEN, AND F. SCHORFHEIDE (2021): “Heterogeneity and aggregate fluctuations,” Discussion paper, National Bureau of Economic Research.
- CHIB, S., F. NARDARI, AND N. SHEPHARD (2006): “Analysis of high dimensional multivariate stochastic volatility models,” *Journal of Econometrics*, 134(2), 341–371.
- CHIU, C. J., H. MUMTAZ, AND G. PINTER (2017): “Forecasting with VAR models: Fat tails and stochastic volatility,” *International Journal of Forecasting*, 33(4), 1124–1143.
- 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., F. HUBER, G. KOOP, M. MARCELLINO, AND M. PFARRHOFER (2022): “Tail forecasting with multivariate Bayesian additive regression trees,” *International Economic Review*.
- CLARK, T. E., AND E. MERTENS (2023): “Stochastic volatility in Bayesian vector autoregressions,” in *Oxford Research Encyclopedia of Economics and Finance*. Oxford University Press.
- CLARK, T. E., AND F. RAVAZZOLO (2015): “Macroeconomic forecasting performance under alternative specifications of time-varying volatility,” *Journal of Applied Econometrics*, 30(4), 551–575.
- COGLEY, T., AND T. J. SARGENT (2001): “Evolving post-world war II US inflation dynamics,” *NBER Macroeconomics Annual*, 16, 331–388.
- (2005): “Drifts and volatilities: Monetary policies and outcomes in the post WWII US,” *Review of Economic Dynamics*, 8(2), 262–302.
- CREAL, D. (2012): “A survey of sequential Monte Carlo methods for economics and finance,” *Econometric reviews*, 31(3), 245–296.
- 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.
- CROSS, J. L., C. HOU, G. KOOP, AND A. POON (2023): “Large stochastic volatility in mean VARs,” *Journal of Econometrics*, 236(1), 105469.
- CRUMP, R. K., S. EUSEPI, D. GIANNONE, E. QIAN, AND A. M. SBORDONE (2021): “A large Bayesian VAR of the United States economy,” *FRB of New York Staff Report No. 976*.
- CÚRDIA, V., M. DEL NEGRO, AND D. L. GREENWALD (2014): “Rare shocks, great recessions,” *Journal of Applied Econometrics*, 29(7), 1031–1052.
- 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.
- DEL NEGRO, M., AND F. SCHORFHEIDE (2004): “Priors from General Equilibrium Models for VARs,” *International Economic Review*, 45, 643–673.

- DEMIRER, M., F. X. DIEBOLD, L. LIU, AND K. YILMAZ (2018): “Estimating global bank network connectedness,” *Journal of Applied Econometrics*, 33(1), 1–15.
- DOAN, T., R. LITTERMAN, AND C. SIMS (1984): “Forecasting and conditional projection using realistic prior distributions,” *Econometric reviews*, 3(1), 1–100.
- DOUCET, A., N. DE FREITAS, AND N. GORDON (eds.) (2001): *Sequential Monte Carlo Methods in Practice*. Springer, New York.
- DURBIN, J., AND S. J. KOOPMAN (1997): “Monte Carlo maximum likelihood estimation for non-Gaussian state space models,” *Biometrika*, 84, 669–684.
- (2002): “A Simple and Efficient Simulation Smoother for State Space Time Series Analysis,” *Biometrika*, 89, 603–615.
- ELLAHIE, A., AND G. RICCO (2017): “Government purchases reloaded: Informational insufficiency and heterogeneity in fiscal VARs,” *Journal of Monetary Economics*, 90, 13–27.
- ELTOFT, T., T. KIM, AND T. LEE (2006): “On the multivariate Laplace distribution,” *Signal Processing Letters, IEEE*, 13(5), 300–303.
- FISCHER, M. M., N. HAUZENBERGER, F. HUBER, AND M. PFARRHOFER (2023): “General Bayesian time-varying parameter vector autoregressions for modeling government bond yields,” *Journal of Applied Econometrics*, 38(1), 69–87.
- FOLLETT, L., AND C. YU (2019): “Achieving Parsimony in Bayesian VARs with the Horseshoe Prior,” *Econometrics and Statistics*, 11, 130–144.
- GEFANG, D., G. KOOP, AND A. POON (2023): “Variational Bayesian inference in large Vector Autoregressions with hierarchical shrinkage,” *International Journal of Forecasting*, 39(1), 346–363.
- GEWEKE, J. (1993): “Bayesian Treatment of the Independent Student- t Linear Model,” *Journal of Applied Econometrics*, 8, S19–S40.
- GIANFREDA, A., F. RAVAZZOLO, AND L. ROSSINI (2023): “Large time-varying volatility models for hourly electricity prices,” *Oxford Bulletin of Economics and Statistics*, 85(3), 545–573.
- GIANNONE, D., M. LENZA, D. MOMFERATOU, AND L. ONORANTE (2014): “Short-term inflation projections: A Bayesian vector autoregressive approach,” *International journal of forecasting*, 30(3), 635–644.
- GIANNONE, D., M. LENZA, AND G. E. PRIMICERI (2015): “Prior selection for vector autoregressions,” *Review of Economics and Statistics*, 97(2), 436–451.

- GIANNONE, D., L. REICHLIN, AND D. SMALL (2008): “Nowcasting: The real-time informational content of macroeconomic data,” *Journal of monetary economics*, 55(4), 665–676.
- GÖTZ, T., AND K. HAUZENBERGER (2018): “Large mixed-frequency VARs with a parsimonious time-varying parameter structure,” *Deutsche Bundesbank Discussion Paper*.
- GRIFFIN, J., AND P. BROWN (2010): “Inference with normal-gamma prior distributions in regression problems,” *Bayesian Analysis*, 5(1), 171–188.
- HAJARGASHT, G., AND T. WOŹNIAK (2018): “Accurate computation of marginal data densities using variational Bayes,” *arXiv preprint arXiv:1805.10036*.
- HARTWIG, B. (2020): “Robust inference in time-varying structural VAR models: The DC-Cholesky multivariate stochastic volatility model,” *Deutsche Bundesbank Discussion Paper No 34/2020*.
- 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.
- HAUZENBERGER, N., F. HUBER, G. KOOP, AND J. MITCHELL (2022): “Bayesian modeling of time-varying parameters using regression trees,” *arXiv preprint arXiv:2209.11970*.
- HAUZENBERGER, N., F. HUBER, M. MARCELLINO, AND N. PETZ (2021): “Gaussian process vector autoregressions and macroeconomic uncertainty,” *arXiv preprint arXiv:2112.01995*.
- HERBST, E., AND F. SCHORFHEIDE (2019): “Tempered particle filtering,” *Journal of Econometrics*, 210(1), 26–44.
- HOU, C., B. H. NGUYEN, AND B. ZHANG (2023): “Real-time forecasting of the Australian macroeconomy using flexible Bayesian VARs,” *Journal of Forecasting*, 42(2), 418–451.
- HUBER, F., AND M. FELDKIRCHER (2019): “Adaptive shrinkage in Bayesian vector autoregressive models,” *Journal of Business and Economic Statistics*, 37(1), 27–39.
- HUBER, F., G. KOOP, AND L. ONORANTE (2019): “Inducing sparsity and shrinkage in time-varying parameter models,” *arXiv preprint arXiv:1905.10787*.
- 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.

- JORDAN, M. I., Z. GHAHRAMANI, T. S. JAAKKOLA, AND L. K. SAUL (1999): “An introduction to variational methods for graphical models,” *Machine Learning*, 37(2), 183–233.
- JURADO, K., S. C. LUDVIGSON, AND S. NG (2015): “Measuring uncertainty,” *American Economic Review*, 105(3), 1177–1216.
- 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.
- KARLSSON, S. (2013): “Forecasting with Bayesian vector autoregressions,” in *Handbook of Economic Forecasting*, ed. by G. Elliott, and A. Timmermann, vol. 2 of *Handbook of Economic Forecasting*, pp. 791–897. Elsevier.
- KARLSSON, S., AND S. MAZUR (2020): “Flexible fat-tailed vector autoregression,” Working Paper No. 5/2020, Örebro University School of Business, Örebro.
- 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. (2013): “Forecasting with medium and large Bayesian VARs,” *Journal of Applied Econometrics*, 28(2), 177–203.
- KOOP, G., AND D. KOROBILIS (2010): “Bayesian Multivariate Time Series Methods for Empirical Macroeconomics,” *Foundations and Trends in Econometrics*, 3(4), 267–358.
- (2013): “Large time-varying parameter VARs,” *Journal of Econometrics*, 177(2), 185–198.
- KOOP, G., AND D. KOROBILIS (2016): “Model uncertainty in panel vector autoregressive models,” *European Economic Review*, 81, 115–131.

- KOOP, G., AND D. KOROBILIS (2018): “Variational Bayes inference in high-dimensional time-varying parameter models,” *Available at SSRN 3246472*.
- KOOP, G., S. MCINTYRE, AND J. MITCHELL (2020): “UK regional nowcasting using a mixed frequency vector auto-regressive model with entropic tilting,” *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 183(1), 91–119.
- 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.
- KOROBILIS, D., AND D. PETTENUZZO (2019): “Adaptive hierarchical priors for high-dimensional vector autoregressions,” *Journal of Econometrics*, forthcoming.
- LEEPER, E. M., C. A. SIMS, AND T. ZHA (1996): “What does monetary policy do?,” *Brookings papers on economic activity*, 1996(2), 1–78.
- 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.
- LITTERMAN, R. (1986): “Forecasting With Bayesian Vector Autoregressions — Five Years of Experience,” *Journal of Business and Economic Statistics*, 4, 25–38.
- LOAIZA-MAYA, R., AND M. S. SMITH (2020): “Real-time macroeconomic forecasting with a heteroscedastic inversion copula,” *Journal of Business and Economic Statistics*, 38(2), 470–486.
- 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*.
- MCCAUSLAND, W. J. (2012): “The HESSIAN Method: Highly efficient simulation smoothing, in a nutshell,” *Journal of Econometrics*, 168(2), 189–206.
- MCCAUSLAND, W. J., S. MILLER, AND D. PELLETIER (2011): “Simulation smoothing for state-space models: A computational efficiency analysis,” *Computational Statistics and Data Analysis*, 55(1), 199–212.
- MCCRACKEN, M. W., AND S. NG (2016): “FRED-MD: A monthly database for macroeconomic research,” *Journal of Business and Economic Statistics*, 34(4), 574–589.

- (2020): “FRED-QD: A quarterly database for macroeconomic research,” *Federal Reserve Bank of St. Louis Working Paper*, 2020-005.
- MCCRACKEN, M. W., M. T. OWYANG, AND T. SEKHPOSYAN (2021): “Real-time forecasting and scenario analysis using a large mixed-frequency Bayesian VAR,” *International Journal of Central Banking*, 17(71), 1–41.
- 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*.
- MUMTAZ, H., AND F. ZANETTI (2013): “The Impact of the Volatility of Monetary Policy Shocks,” *Journal of Money, Credit and Banking*, 45(4), 535–558.
- NI, S., AND D. SUN (2005): “Bayesian estimates for vector autoregressive models,” *Journal of Business and Economic Statistics*, 23(1), 105–117.
- ORMEROD, J. T., AND M. P. WAND (2010): “Explaining variational approximations,” *The American Statistician*, 64(2), 140–153.
- PARK, T., AND G. CASELLA (2008): “The Bayesian Lasso,” *Journal of the American Statistical Association*, 103(482), 681–686.
- PHILIPPOV, A., AND M. GLICKMAN (2006): “Multivariate stochastic volatility via Wishart processes,” *Journal of Business and Economic Statistics*, 24, 313–328.
- PITT, M., AND N. SHEPHARD (1999): “Time varying covariances: a factor stochastic volatility approach,” *Bayesian Statistics*, 6, 547–570.
- POON, A. (2018): “Assessing the synchronicity and nature of Australian state business cycles,” *Economic Record*, 94(307), 372–390.
- PRADO, R., AND M. WEST (2010): *Time Series: Modeling, Computation, and Inference*. Chapman and Hall/CRC.
- PRIMICERI, G. E. (2005): “Time Varying Structural Vector Autoregressions and Monetary Policy,” *Review of Economic Studies*, 72(3), 821–852.
- SCHORFHEIDE, F., AND D. SONG (2015): “Real-time forecasting with a mixed-frequency VAR,” *Journal of Business and Economic Statistics*, 33(3), 366–380.
- SCHORFHEIDE, F., AND D. SONG (2021): “Real-time forecasting with a (standard) mixed-frequency VAR during a pandemic,” Discussion paper, National Bureau of Economic Research.

- SHEPHARD, N., AND M. K. PITT (1997): “Likelihood analysis of non-Gaussian measurement time series,” *Biometrika*, 84, 653–667.
- SHIN, M., AND M. ZHONG (2020): “A new approach to identifying the real effects of uncertainty shocks,” *Journal of Business and Economic Statistics*, 38(2), 367–379.
- SIMS, C. A. (1980): “Macroeconomics and reality,” *Econometrica*, 48, 1–48.
- (2001): “Comment on Sargent and Cogley’s ‘Evolving US Postwar Inflation Dynamics’,” *NBER Macroeconomics Annual*, 16, 373–379.
- SIMS, C. A., AND T. ZHA (2006): “Were There Regime Switches in U.S. Monetary Policy?,” *American Economic Review*, 96(1), 54–81.
- 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.
- UHLIG, H. (1997): “Bayesian vector autoregressions with stochastic volatility,” *Econometrica*, 65(1), 59–73.
- WEST, M., AND J. HARRISON (2006): *Bayesian Forecasting and Dynamic Models*. Springer Science & Business Media.
- WU, P., AND G. KOOP (2022): “Fast, order-invariant Bayesian inference in VARs using the eigendecomposition of the error covariance matrix,” *working paper*.
- 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.