

Computationally Efficient Inference in Large Bayesian Mixed Frequency VARs

Deborah Gefang, Gary Koop and Aubrey Poon

ESCoE Discussion Paper 2020-07

May 2020

Computationally Efficient Inference in Large Bayesian Mixed Frequency VARs
Deborah Gefang, Gary Koop and Aubrey Poon
ESCoE Discussion Paper No. 2020-07
May 2020

Abstract

Mixed frequency Vector Autoregressions (MF-VARs) can be used to provide timely and high frequency estimates or nowcasts of variables for which data is available at a low frequency. Bayesian methods are commonly used with MF-VARs to overcome over-parameterization concerns. But Bayesian methods typically rely on computationally demanding Markov Chain Monte Carlo (MCMC) methods. In this paper, we develop Variational Bayes (VB) methods for use with MF-VARs using Dirichlet-Laplace global-local shrinkage priors. We show that these methods are accurate and computationally much more efficient than MCMC in two empirical applications involving large MF-VARs.

Keywords: Mixed Frequency, Variational inference, Vector Autoregression, Stochastic Volatility, Hierarchical Prior, Forecasting.

JEL classification: C11, C32, C53

Deborah Gefang, University of Leicester, dg171@leicester.ac.uk, Gary Koop, University of Strathclyde and ESCoE, gary.koop@strath.ac.uk and Aubrey Poon, University of Strathclyde and ESCoE, aubrey.poon@strath.ac.uk.

Published by:
Economic Statistics Centre of Excellence
National Institute of Economic and Social Research
2 Dean Trench St
London SW1P 3HE
United Kingdom
www.escoe.ac.uk

ESCoE Discussion Papers describe research in progress by the author(s) and are published to elicit comments and to further debate. Any views expressed are solely those of the author(s) and so cannot be taken to represent those of the Economic Statistics Centre of Excellence (ESCoE), its partner institutions or the Office for National Statistics (ONS).

Computationally Efficient Inference in Large Bayesian Mixed Frequency VARs

Deborah Gefang*
University of Leicester

Gary Koop†
University of Strathclyde and Economic Statistics Centre of Excellence

Aubrey Poon ‡
University of Strathclyde and Economic Statistics Centre of Excellence

Abstract

Mixed frequency Vector Autoregressions (MF-VARs) can be used to provide timely and high frequency estimates or nowcasts of variables for which data is available at a low frequency. Bayesian methods are commonly used with MF-VARs to overcome over-parameterization concerns. But Bayesian methods typically rely on computationally demanding Markov Chain Monte Carlo (MCMC) methods. In this paper, we develop Variational Bayes (VB) methods for use with MF-VARs using Dirichlet-Laplace global-local shrinkage priors. We show that these methods are accurate and computationally much more efficient than MCMC in two empirical applications involving large MF-VARs.

Keywords: Mixed Frequency, Variational inference, Vector Autoregression, Stochastic Volatility, Hierarchical Prior, Forecasting

JEL Classifications: C11, C32, C53

*Email: dg171@leicester.ac.uk

†Email: gary.koop@strath.ac.uk

‡Email: aubrey.poon@strath.ac.uk

1 Introduction

Vector Autoregressions (VARs) have had great success for macroeconomic forecasting and structural analysis. Recently, there has been an upsurge of interest in mixed frequency VAR (MF-VAR) models which incorporate variables of different frequencies into a VAR. Important contributions include Carriero, Clark and Marcellino (2015), Eraker, Chiu, Foerster, Kim and Seoane (2015), Schorfheide and Song (2015), Ghysels (2016), McCracken, Owyang and Sekhposyan (2016), Brave, Butters and Justiniano (2018), Gotz and Hauzenberger (2018) and Koop, McIntyre, Mitchell and Poon (2020). A common case is when interest centers on a quarterly variable (e.g. GDP) and the researcher has available many monthly predictors. An MF-VAR allows for timely and frequent updating of forecasts or nowcasts of the quarterly variable. It also allows producing monthly estimates of the quarterly variable which can be useful in some structural contexts where the economic question of interest is at a high frequency (e.g. Cotter, Hallam and Yilmaz, 2017).

The availability of more and more data has also led to much recent interest in large VARs involving tens or even hundreds of variables. For the same reasons that large VARs are popular, one would expect large MF-VARs to be popular. Indeed, the mixed frequency nature of the MF-VAR greatly broadens the number of variables one could include. Large quarterly data sets could, in theory, be jointly modelled with large monthly data sets or even large daily financial data sets. But so far little has been done. Of the papers cited above, the largest MF-VAR is that of Brave, Butters and Justiniano (2018) which involves 7 quarterly and 30 monthly variables. The purpose of the present paper is to develop econometric methods for filling this gap.

Bayesian methods are usually used with MF-VARs and large VARs. Bayesian priors can be used to overcome the over-parameterization problems associated with both of these. Posterior and predictive densities are usually uncovered using MCMC methods and these can be computationally slow. It is possible that, for this reason, the MF-VARs cited above do not use large data sets. Using MCMC methods in an MF-VAR with 10 or 20 variables will be slow, but computationally feasible. Working with 50 or 100 variables may simply be too computationally costly, especially in the context of a recursive forecasting exercise which requires repeatedly using MCMC methods on an expanding or rolling window of data. Long computation times are also undesirable in policy circles where the goal is to release updated nowcasts or forecasts quickly as new high frequency information becomes available.

The goal of the present paper is to develop Variational Bayesian (VB) methods for large MF-VARs

and see how well they work in practice. VB methods provide a computationally-efficient alternative to MCMC methods. A potential drawback of VB methods is that they provide only an approximation to the posterior density. In Gefang, Koop and Poon (2019), we developed VB methods for a range of large VARs and found them to be accurate and computationally efficient (see also Hajargasht and Wozniak, 2018). In VARs with 100 variables, MCMC methods took over 100 times as much computer time as comparable VB methods for a range of priors.¹ The main contribution of the present paper is to extend the methods of our earlier work with large VARs to large MF-VARs. This extension is not trivial since MF-VARs are state space models and the parameter space thus includes large numbers of latent states. The prior we use for the VAR coefficients in the MF-VAR is the Dirichlet-Laplace (D-L) prior of Bhattacharya et al (2015) and a secondary contribution of this paper lies in developing VB methods for MF-VARs with this prior. Our derivations relating to the D-L prior can also be used with the VAR, thus extending Gefang, Koop and Poon (2019). We carry out two empirical applications and find VB methods to perform well in terms of estimation accuracy as well as computational efficiency.

2 Variational Bayesian Inference

VB methods are commonly used in many statistical fields and are increasingly being used by econometricians. Blei, Kucukelbir and McAuliffe (2017) describes the theory and practice of VB in detail. The basic idea is that a posterior, $p(\theta|y)$ using data y for parameters θ , is approximated using a simpler p.d.f. $q(\theta)$ of the form:

$$q(\theta) = \prod_{m=1}^M q_m(\theta_m), \quad (1)$$

where θ_m for $m = 1, \dots, M$ are the blocks of parameters which make up θ . The approximation which is as close as possible in a Kullback-Liebler sense, can be shown to be

$$q_m(\theta_m) = \exp[E(\log p(\theta_m|y, \theta_{-m}))], \quad (2)$$

where θ_{-m} denotes all parameters except those in θ_m and the expectation is taken over $q(\theta_{-m})$. The VB algorithm proceeds by finding the optimal arguments in the densities $q_m(\theta_m)$ in an iterative fashion. This iterative procedure is typically much faster than MCMC. There are two ways of assessing convergence. The first is to evaluate the evidence lower bound (ELBO, see the online appendix) at

¹This statement is based on 22,000 MCMC draws.

each iteration. Convergence is achieved when the change in ELBO across iterations is less than some convergence criterion. Alternatively, convergence occurs when the VB estimates of parameters or states stop changing across iterations. In our empirical work, we use the former strategy for the parameters of the model (e.g. VAR coefficients and D-L prior hyperparameters) and the latter strategy with respect to the states.

3 The Mixed Frequency VAR

Let $\mathbf{y}_t = (\mathbf{y}_t^H, \mathbf{y}_t^L)'$ be a vector of $n = n_H + n_L$ variables involving n_H high frequency variables and n_L low frequency variables. Time t subscripts measure time at the high frequency. We write the MF-VAR in a structural form which allows for equation-by-equation estimation which leads to substantial computational improvement:

$$\mathbf{A}_0 \mathbf{y}_t = \mathbf{b}_0 + \mathbf{B}_1 \mathbf{y}_{t-1} + \dots + \mathbf{B}_p \mathbf{y}_{t-p} + \epsilon_t, \epsilon_t \sim N(0, \Sigma), \quad (3)$$

for $t = 1, \dots, T$ where \mathbf{b}_0 is a $n \times 1$ vector of intercept terms, \mathbf{B}_i is the $n \times n$ matrix of lag i VAR coefficients and \mathbf{A}_0 is an $n \times n$ lower triangular matrix with ones on the diagonal.

For estimation purposes, we re-write the MF-VAR as:

$$\mathbf{y}_t = \mathbf{X}_t \beta + \mathbf{W}_t \mathbf{a} + \epsilon_t, \quad (4)$$

where $\mathbf{X}_t = \mathbf{I}_n \otimes [1, \mathbf{y}'_{t-1}, \dots, \mathbf{y}'_{t-p}]$ is an $n \times K$ matrix, $\beta = \text{vec}([\mathbf{b}_0, \mathbf{B}_1, \dots, \mathbf{B}_p]')$ is $K \times 1$ vector of coefficients, \mathbf{a} consists of the free elements of \mathbf{A}_0 stacked by rows with \mathbf{W}_t being the $n \times m$ matrix containing the appropriate contemporaneous elements of \mathbf{y}_t . Equation (4) can be written in terms of n independent equations, with the i^{th} equation being:

$$y_{i,t} = \mathbf{z}_{i,t} \theta_i + \epsilon_{i,t}, \epsilon_{i,t} \sim N(0, \sigma_i^2). \quad (5)$$

where $\mathbf{z}_{i,t}$ is a row vector with k_i elements and θ_i is a vector containing the elements of β and \mathbf{a} pertaining to the i^{th} equation.

For the mixed frequency VAR, \mathbf{y}_t^L are treated as unobserved latent variables. The relationship

between them and the observed low frequency variables is given by

$$\mathbf{y}_{i,t}^{L,O} = \mathbf{M}_{i,t}^L \Lambda_i^L \mathbf{z}_{i,t+1} \quad (6)$$

where $\mathbf{y}_{i,t}^{L,O}$ denotes the observed values of the i^{th} low frequency variable for $i = 1, \dots, n_L$.²

The form of $\mathbf{M}_{i,t}^L$ and Λ_i^L in (6) depends on the frequency mis-match and the way the variables are transformed (e.g. whether the variables are logged or log differenced). Hence, we will leave them unspecified and refer the reader to the online appendix for details. But they involve setting $\mathbf{M}_{i,t}^L = \mathbf{1}$ at times the low frequency variables are observed and setting it to an empty matrix at other times. When forecasting, it can also handle release delays by setting it to an empty matrix for periods before an observation is released.

For the high frequency variables, we have a similar specification:

$$\mathbf{y}_t^{H,O} = \mathbf{M}_t^H \Lambda^H \mathbf{y}_t. \quad (7)$$

The forms for \mathbf{M}_t^H and Λ^H are simpler since the former is equal to 1 in all periods (unless there are release delays) and the latter is defined so as to pick out the time t value of observed high frequency variable.

The mixed frequency VAR (MF-VAR) is a state space model with state equations given by (5) and measurement equations given by (6) and (7). Bayesian MCMC methods for the MF-VAR are outlined in Koop, McIntyre, Mitchell and Poon (2020).

4 Prior Shrinkage Using The Dirichlet-Laplace Prior

In previous work, Gefang, Koop and Poon (2019), we developed VB methods for large VARs using a number of popular global-local shrinkage priors such as the Least Absolute Shrinkage and Selection Operator (LASSO) and the stochastic search variable selection (SSVS) priors. However, recently Bhattacharya et al. (2015), has proposed an alternative, the D-L prior, which has a stronger theoretical justification than other priors (e.g. they show that the joint posterior distribution of the parameters concentrates at the optimal rate) and reduces the problem of prior hyperparameter choice to a single hyperparameter. This hyperparameter can either be set to recommended default values suggested in

²Depending on which transformation we use, the $\mathbf{z}_{i,t+1}$ in this equation may have to be augmented with some lags.

Bhattacharya et al. (2015) or can be treated as unknown and estimated. The D-L prior has been used in MF-VARs in Koop, McIntyre, Mitchell and Poon (2020) using MCMC methods. In this paper we develop VB methods for use with the MF-VAR with the D-L prior which are computationally practical even in very high dimensional models.

The D-L prior is given by:

$$\theta_{i,j}|\phi_{i,j}, \tau_i \sim DE(\phi_{i,j}\tau_i), \quad \phi_{i,j} \sim Dir(a, \dots, a), \quad (8)$$

where $\theta_{i,j}$ is the j^{th} coefficient in θ_i (see equation 5), DE denote the double-exponential or Laplace distribution and Dir denotes the Dirichlet distribution.

The formal properties of the D-L prior are discussed in Bhattacharya et al. (2015). Here we note informally that a DE prior distribution imposes $L1$ shrinkage as is used with the Bayesian LASSO. Relative to $L2$ shrinkage priors, this can be shown to have the desirable properties of shrinking small (unimportant) coefficients more strongly towards zero and shrinking large (important) coefficients less. But the theoretical shrinkage properties of priors based solely on $L1$ shrinkage such as the LASSO are imperfectly understood and prior hyperparameter choice can be difficult. The addition of the Dirichlet distribution for the local shrinkage prior hyperparameters surmounts these problems. The good theoretical properties of the D-L prior (e.g. as relating to posterior concentration) have been shown in Bhattacharya et al. (2015) and it involves only a single prior hyperparameter: a . In the empirical work in this paper we use a common default choice and set it to $\frac{1}{k_i+1}$. In other words, it is inversely proportional to the number of parameters in the equation being estimated. However, it is worth noting that there are other ways of selecting a or it could even be treated as an unknown parameter and estimated from the data. For instance, Zhang and Bondell (2018), reasoning that the R^2 in a regression is easier to interpret than the hyperparameter a , develop a method for choosing a which involves eliciting a prior for R^2 .

It can be shown that an equivalent way of writing D-L prior is:

$$\theta_{i,j} \sim N(0, \psi_{i,j}\phi_{i,j}^2\tau_i^2), \quad \psi_{i,j} \sim Exp(1/2) \quad (9)$$

Hence, the prior for all the coefficients in equation i is: θ_i is $N(\mathbf{0}, \mathbf{V}_i)$ where $\mathbf{V}_i = diag(\psi_{i,1}\phi_{i,1}^2\tau_i^2, \dots, \psi_{i,k_n}\phi_{i,k_n}^2\tau_i^2)$. Writing the D-L prior as a scale mixture of Normals is convenient since

MCMC methods can exploit the (conditional) conjugacy between the prior and the Normal likelihood.

Following Bhattacharya et al. (2015), we use a Gamma prior for τ_i :

$$\tau_i \sim G(k_i a, 1/2). \quad (10)$$

For σ_i^{-2} we use a standard Gamma prior

$$\sigma_i^{-2} \sim G(\underline{s}, \underline{\nu}). \quad (11)$$

5 Variational Bayes Estimation of the Mixed Frequency VAR

5.1 Notation

As discussed above, we do Bayesian estimation of the MF-VAR equation by equation. For notational simplicity, we will not explicitly note this in the following definitions and derivations. That is, all the data quantities and parameters below should have i subscripts in the following material. With this notational convention, the unknown parameters in each equation are $\theta = (\theta_1, \dots, \theta_k)'$, $\psi = (\psi_1, \dots, \psi_k)'$, $\phi = (\phi_1, \dots, \phi_k)'$, τ and σ^2 . The vector $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_T)'$ contains both observed and unobserved data values as noted above. We also use notation where $\mathbf{Z} = (\mathbf{Z}_1, \dots, \mathbf{Z}_T)'$.

Our goal is to obtain $q(\bullet)$, the VB optimal approximating for each parameter or state \bullet . We will use notation where upper bars denote the arguments of these densities. These are what are optimized across iterations in the VB algorithm. So, for instance, $\bar{\mathbf{y}}$ and $\bar{\mathbf{Z}}$ will be the dependent and explanatory variables in an equation with the VB estimates of the unobserved low frequency estimates plugged in.

We will discuss VB estimation for the parameters (conditional on the VB estimates for the states) and then estimates of the states (given VB estimates of the parameters) in the next two subsections.

5.2 VB Estimation of the Parameters

To obtain the VB approximating densities, we require the full conditional posteriors for all parameters in the model. In this sub-section, we condition on the states so the relevant dependent and explanatory variables in each equation are $\bar{\mathbf{y}}$ and $\bar{\mathbf{Z}}$. The full conditionals are easily available. For θ and σ^2 textbook sources for Bayesian results for the Normal linear regression model can be used. The others are available in Bhattacharya et al. (2015). Using these we can construct the optimal q densities and compute the

ELBO.

Remember that the VB algorithm is an iterative one and the goal of the analysis is to produce optimal choices for the quantities which have upper bars over them in the following formulae. Each of the densities below contains arguments involving such upper bar quantities (e.g. the mean and variance of $q(\theta)$ are given by $\bar{\theta}$ and $\bar{\mathbf{V}}$). But the upper bar quantities in a particular density also depend on upper bar quantities which appear in different densities (e.g. $\bar{\theta}$ depends on $\bar{\mathbf{y}}$ and $\bar{\mathbf{Z}}$). By iterating over all of the densities we can obtain the overall global optimal choices for the upper bar quantities. So, to continue the example, if we use the formula below to evaluate $\bar{\theta}$ for particular values for $\bar{\mathbf{y}}$ and $\bar{\mathbf{Z}}$ we will obtain the optimal value of $\bar{\theta}$ for those particular values, but not the global optimum. It is only by iterating between $\bar{\theta}$, $\bar{\mathbf{y}}$ and $\bar{\mathbf{Z}}$ until convergence is achieved that we will obtain the global optimum.

The VB approximating densities are given by:

5.2.1 $q(\theta)$

$$q(\theta) \sim N(\bar{\theta}, \bar{\mathbf{V}}), \quad (12)$$

where

$$\bar{\mathbf{V}} = \left(\frac{T}{\bar{s}} + \frac{\nu}{\bar{s}} \bar{\mathbf{Z}}' \bar{\mathbf{Z}} + \mathbf{V}^{-1} \right)^{-1}$$

$$\bar{\theta} = \frac{T}{\bar{s}} + \frac{\nu}{\bar{s}} \bar{\mathbf{V}} \bar{\mathbf{Z}}' \bar{\mathbf{y}}$$

with $\mathbf{V}^{-1} = \text{diag}(\overline{\psi_1^{-1} \phi_1^{-2} \tau^{-2}}, \dots, \overline{\psi_k^{-1} \phi_k^{-2} \tau^{-2}})$.

5.2.2 $q(\sigma^{-2})$

$$q(\sigma^{-2}) \sim G\left(\frac{T}{2} + \underline{\nu}, \bar{s}\right), \quad (13)$$

where

$$\bar{s} = \frac{1}{2} [\|\bar{\mathbf{y}} - \bar{\mathbf{Z}}\bar{\theta}\|^2 + \text{tr}(\bar{\mathbf{Z}}' \bar{\mathbf{Z}} \bar{\mathbf{V}})] + \underline{s}.$$

This leads to the following value being updated by the VB algorithm:

$$\overline{\sigma^{-2}} = \frac{\frac{T}{2} + \underline{\nu}}{\bar{s}}.$$

5.2.3 $q(\tau)$

$$q(\tau) \sim giG(ka - k, 1, \sum_{j=1}^k 2\overline{\phi_j^{-1}}[(\overline{\theta_j})^2 + \overline{\mathbf{V}^{jj}}]^{1/2}), \quad (14)$$

where giG denotes the generalized inverse Gaussian distribution.

Letting $\chi = \sum_{j=1}^k 2\overline{\phi_j^{-1}}[(\overline{\theta_j})^2 + \overline{\mathbf{V}^{jj}}]^{1/2}$, we obtain quantities to be iterated over in VB of

$$\overline{\tau} = \frac{\sqrt{\chi} K_{ka-k+1}(\sqrt{\chi})}{K_{ka-k}(\sqrt{\chi})}$$

and

$$\overline{\tau^2} = (\overline{\tau})^2 + \chi \left[\frac{K_{ka-k+2}(\sqrt{\chi})}{K_{ka-k}(\sqrt{\chi})} - \left(\frac{K_{ka-k+1}(\sqrt{\chi})}{K_{ka-k}(\sqrt{\chi})} \right)^2 \right]$$

where $K_*[\bullet]$ is the modified Bessel function of the second kind.

5.2.4 $q(\psi_j^{-1})$

$$q(\psi_j^{-1}) \sim iG\left(\sqrt{\frac{\overline{\phi_j^2} \overline{\tau^2}}{(\overline{\theta_j})^2 + \overline{\mathbf{V}^{jj}}}}, 1\right), \quad (15)$$

where iG denotes the inverse Gaussian distribution.

Letting $\rho = \sqrt{\frac{\overline{\phi_j^2} \overline{\tau^2}}{(\overline{\theta_j})^2 + \overline{\mathbf{V}^{jj}}}}$, we obtain the quantity to be updated by VB of $\overline{\psi_j} = 1 + 1/\rho$.

5.2.5 $q(\phi_j)$

$$q(\xi_j) \sim giG(a - 1, 1, 2\sqrt{(\overline{\theta_j})^2 + \overline{\mathbf{V}^{jj}}}) \quad (16)$$

Let $\varpi = 2\sqrt{(\overline{\theta_j})^2 + \overline{\mathbf{V}^{jj}}}$, we have $\overline{\xi_j} = \frac{\sqrt{\varpi} K_a(\sqrt{\varpi})}{K_{a-1}(\sqrt{\varpi})}$, and $var(\xi_j) = \varpi \left\{ \frac{K_{a+1}(\sqrt{\varpi})}{K_{a-1}(\sqrt{\varpi})} - \left[\frac{K_a(\sqrt{\varpi})}{K_{a-1}(\sqrt{\varpi})} \right]^2 \right\}$.

Scaling ξ_j , we have $\overline{\phi_j} = \frac{\overline{\xi_j}}{\sum_{j=1}^k \overline{\xi_j}}$ and $\overline{\phi_j^2} = (\overline{\phi_j})^2 + \frac{var(\xi_j)}{(\sum_{j=1}^k \overline{\xi_j})^2}$. Thus, the optimal q density of ϕ_j takes the following form:

$$q(\phi_j) \sim giG\left[a - 1, \sum_{j=1}^k \overline{\xi_j}, \varpi / \left(\sum_{j=1}^k \overline{\xi_j} \right)\right] \quad (17)$$

The evidence lower bound (*ELBO*) takes the following form:

$$\begin{aligned}
ELBO_n &= E\{\log p(\mathbf{y}_n, \theta_n, \sigma_n^{-2}, \phi_n, \tau_n, \psi_n)\} - E\{\log q(\theta_n, \sigma_n^{-2}, \phi_n, \tau_n, \psi_n)\} \\
&= \frac{1}{2} \log(|\bar{\mathbf{V}}_n|) - E\left[\frac{1}{2} \log(|\sigma^2 \mathbf{V}_n|)\right] - \frac{1}{2} \left[\frac{\frac{T+k_n}{2} + \underline{\nu}}{\bar{s}_n} \bar{\theta}'_n \mathbf{V}_n^{-1} \bar{\theta}_n + \text{tr}\left(\frac{\frac{T+k_n}{2} + \underline{\nu}}{\bar{s}_n} \mathbf{V}_n^{-1} \bar{\mathbf{V}}_n\right) \right] - \left(\underline{\nu} + \frac{T+k_n}{2}\right) \log(\bar{s}_n) \\
&\quad - k_n a \left[\int_0^\infty (q(\tau_n) \log \tau_n) d\tau_n \right] - 0.5 \bar{\tau}_n - \frac{1}{2} \sum_{j=1}^{k_n} \bar{\psi}_{n,j} - (a-1) \sum_{j=1}^{k_n} \left[\int_0^\infty (q(\phi_{n,j}) \log \phi_{n,j}) d\phi_{n,j} \right] \\
&\quad + \int_0^\infty q(\tau_n) \log q(\tau_n) d\tau_n + \sum_{j=1}^{k_n} \int_0^\infty q(\psi_{n,j}) \log q(\psi_{n,j}) d\psi_{n,j} + \sum_{j=1}^{k_n} \int_0^\infty q(\phi_{n,j}) \log q(\phi_{n,j}) d\phi_{n,j} + \text{Const.}
\end{aligned} \tag{18}$$

where

$$\begin{aligned}
E[\log(|\sigma^2 \mathbf{V}_n|)] &= \sum_{j=1}^{k_n} E\left[\int_0^\infty \{q(\psi_{n,j}) \log(\psi_{n,j})\} d\psi_{n,j} + 2 \int_0^\infty \{q(\phi_{n,j}) \log(\phi_{n,j})\} d\phi_{n,j} + 2k_n E\left[\int_0^\infty \{q(\tau_n) \log(\tau_n)\} d\tau_n \right] \right. \\
&\quad \left. + k_n E\left[-\psi\left(\underline{\nu} + \frac{T+k_n}{2}\right) + \log(\bar{s}_n) \right] \right],
\end{aligned} \tag{19}$$

and

$$q(\psi_{n,j}) = \left(\frac{\psi_{n,j}^{-1}}{2\pi}\right)^{1/2} \exp\left\{-\frac{(\psi_{n,j}^{-1} - \rho)^2}{2\rho^2 \psi_{n,j}^{-1}}\right\}. \tag{20}$$

5.3 VB Estimation of the States

The MF-VAR is a state space model and VB estimation of particular state space models has been done in several places. Koop and Korobilis (2018) followed Wang et al (2016). Wang et al (2016) does the general state space model with Koop and Korobilis (2018) adapting their methods for a particular TVP regression model with hierarchical shrinkage prior. However, the state space models in these papers differ from ours in that their states are the time-varying regression coefficients which follow random walk or autoregressive processes which are assumed to be uncorrelated with one another. In contrast, our states are the unobserved high frequency values for the low frequency variables which typically will be correlated with one another.

In our case, we use a shrinkage prior on the VAR coefficients, not on the states. Hence, the arguments in the VB approximating density for the states are simple. They can just be obtained using Kalman filtering methods with the VB estimates for all the remaining parameters replaced with their VB estimates as described in the preceding sub-section.

In the remainder of this sub-section, we provide details how the VB algorithm works, focussing on the US application which has a monthly/quarterly frequency mis-match. At the end of the sub-section we describe how this algorithm has to be modified for the regional UK application which involves a different frequency mis-match (i.e. quarterly/annual) and involves an additional cross-sectional restriction based on the fact that output for the UK as a whole is the sum of the regional outputs.

We write the VAR, which are the state equations in the MF-VAR, as:

$$\mathbf{s}_t = \mathbf{F}_0 + \mathbf{F}_1 \mathbf{s}_{t-1} + \epsilon_t, \epsilon_t \sim N(0, \Omega), \quad (21)$$

where $\mathbf{s}_t = (\mathbf{y}_t^H, \mathbf{y}_t^L, \dots, \mathbf{y}_{t-4}^H, \mathbf{y}_{t-4}^L)'$,

$$\mathbf{F}_0 = \begin{bmatrix} \mathbf{A}_0^{-1} \mathbf{b}_0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \quad (22)$$

$$\mathbf{F}_1 = \begin{bmatrix} \mathbf{A}_0^{-1} \mathbf{B}_1 & \mathbf{A}_0^{-1} \mathbf{B}_2 & \mathbf{A}_0^{-1} \mathbf{B}_3 & \mathbf{A}_0^{-1} \mathbf{B}_4 & \mathbf{A}_0^{-1} \mathbf{B}_5 \\ \mathbf{I}_n & 0 & 0 & 0 & 0 \\ 0 & \mathbf{I}_n & 0 & 0 & 0 \\ \vdots & 0 & \mathbf{I}_n & \vdots & \vdots \\ 0 & \dots & 0 & \mathbf{I}_n & 0 \end{bmatrix}, \quad (23)$$

$$\Omega = \text{diag}((\mathbf{A}_0' \Sigma^{-1} \mathbf{A}_0)^{-1}, 0, \dots, 0). \quad (24)$$

In months where the quarterly variables are not observed the measurement equation is:

$$\mathbf{y}_t^H = \mathbf{M}_t^H \mathbf{s}_t, \quad (25)$$

where

$$\mathbf{M}_t^H = \begin{bmatrix} \mathbf{I}_{n_H} & 0 & 0 & \dots & 0 \end{bmatrix}. \quad (26)$$

In months when both monthly and quarterly variables are observed the measurement equation is:

$$\mathbf{y}_t = \mathbf{M}\mathbf{s}_t, \quad (27)$$

where

$$\mathbf{M} = \begin{bmatrix} \mathbf{M}_t^H \\ \mathbf{M}_t^L \end{bmatrix},$$

$$\text{and } \mathbf{M}_t^L = \begin{bmatrix} \mathbf{0}_{n_H} & \frac{1}{3}\mathbf{I}_{n_L} & 0 & \frac{2}{3}\mathbf{I}_{n_L} & 0 & \mathbf{I}_{n_L} & 0 & \frac{2}{3}\mathbf{I}_{n_L} & 0 & \frac{1}{3}\mathbf{I}_{n_L} \end{bmatrix}.$$

Using (18), (22) and (24) we can then run the Kalman filter and smoother to get the interpolated monthly value for the quarterly variables $\tilde{\mathbf{y}}_t^L$. More specifically, the algorithm as it relates to the states can be summarised as below for the q -iteration:

1. The VB methods of Section 5.2 provide values of $\mathbf{F}_0^{(q)}, \mathbf{F}_1^{(q)}$.
2. Run the Kalman filter and smoother using the state space model defined in equations (18), (22) and (24). This will give us $\tilde{\mathbf{y}}_t^{L(q)}$.
3. Then we compare the previous iteration estimate of $\tilde{\mathbf{y}}_t^{L(q-1)}$ with the current iteration estimate of $\tilde{\mathbf{y}}_t^{L(q)}$, where $|\tilde{\mathbf{y}}_t^{L(q)} - \tilde{\mathbf{y}}_t^{L(q-1)}| < \varepsilon$ and ε is a very small threshold value.
4. If this criterion $|\tilde{\mathbf{y}}_t^{L(q)} - \tilde{\mathbf{y}}_t^{L(q-1)}| < \varepsilon$ is met, then we stop the algorithm. Otherwise, we repeat steps 1 to 3.

Note that for the UK application, the algorithm is exactly the same, except for the state space is now defined as:

$$\mathbf{s}_t = \mathbf{F}_0 + \mathbf{F}_1\mathbf{s}_{t-1} + \epsilon_t, \epsilon_t \sim N(0, \Omega), \quad (28)$$

where $\mathbf{s}_t = (\mathbf{y}_t^H, \mathbf{y}_t^L, \dots, \mathbf{y}_{t-7}^H, \mathbf{y}_{t-7}^L)'$,

$$\mathbf{F}_0 = \begin{bmatrix} \mathbf{A}_0^{-1}\mathbf{b}_0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \quad (29)$$

$$\mathbf{F}_1 = \begin{bmatrix} \mathbf{A}_0^{-1}\mathbf{B}_1 & \mathbf{A}_0^{-1}\mathbf{B}_2 & \mathbf{A}_0^{-1}\mathbf{B}_3 & \dots & \mathbf{A}_0^{-1}\mathbf{B}_6 & \mathbf{A}_0^{-1}\mathbf{B}_7 \\ \mathbf{I}_n & 0 & 0 & 0 & 0 & 0 \\ 0 & \mathbf{I}_n & 0 & 0 & 0 & 0 \\ \vdots & 0 & \mathbf{I}_n & \vdots & \vdots & \vdots \\ 0 & \ddots & 0 & \mathbf{I}_n & 0 & \\ 0 & 0 & 0 & 0 & \mathbf{I}_n & 0 \end{bmatrix}, \quad (30)$$

$$\Omega = \text{diag}((\mathbf{A}_0'\Sigma^{-1}\mathbf{A}_0)^{-1}, 0, \dots, 0). \quad (31)$$

In addition to (22), we also have:

$$\mathbf{y}_t = \mathbf{M}\mathbf{s}_t, \quad (32)$$

where

$$\mathbf{M} = \begin{bmatrix} \mathbf{M}_t^H \\ \mathbf{M}_t^L \end{bmatrix},$$

and

$$\mathbf{M}_t^L = \begin{bmatrix} \mathbf{0}_{n_H} & \frac{1}{4}\mathbf{I}_{n_L} & 0 & \frac{1}{2}\mathbf{I}_{n_L} & 0 & \frac{3}{4}\mathbf{I}_{n_L} & 0 & \mathbf{I}_{n_L} & 0 & \frac{3}{4}\mathbf{I}_{n_L} & 0 & \frac{1}{2}\mathbf{I}_{n_L} & 0 & \frac{1}{4}\mathbf{I}_{n_L} \end{bmatrix}, \quad (33)$$

Also, we have an additional measurement equation for the cross-sectional restriction:

$$\mathbf{y}_t^L = \frac{1}{n_L} \sum_{i=1}^{n_L} y_{i,t}^L + \eta_t, \eta_t \sim N(0, \sigma_{cs}^2), \quad (34)$$

and we set the prior for $\sigma_{cs}^2 \sim IG(1000, .001)$. The conditional posterior for σ_{cs}^2 is standard and therefore a VB approximating density for it can be easily derived.

6 Empirical Applications

In this section, we provide evidence that VB methods work well with mixed frequency VARs in two different empirical illustrations. The first is a US macroeconomic exercise involving a single low

frequency variable (quarterly GDP growth) and 50 monthly variables, thus leading to a 51 dimensional MF-VAR. The second illustration involves a different frequency mis-match (annual-quarterly) and there are many more low frequency variables than high frequency variables. In both case we produce smoothed (i.e. full sample) estimates.

6.1 Obtaining Monthly GDP Estimates for the US

Data is taken from the popular FRED data set, see McCracken and Ng (2016). Complete details of the data are provided in the appendix. The goal is to produce monthly estimates of GDP growth. Even with our very large 51 variable MF-VAR, VB methods produce reasonable estimates very quickly (in approximately ten minutes on a good personal computer), making VB methods suitable for pseudo-real time nowcasting exercises which repeatedly estimate the model on an expanding window of data. Using MCMC methods in an MF-VAR of this dimension would take days, making them unsuitable for this purpose.

To convince the reader that our VB estimates are reasonable, Figure 1 compares them to the commonly-used monthly Brave-Butters-Kelley (BBK) estimates of monthly GDP growth produced by the Chicago Fed (see <https://www.chicagofed.org/publications/bbki/index>). It can be seen that they match up very closely and the correlation between them is 0.95.

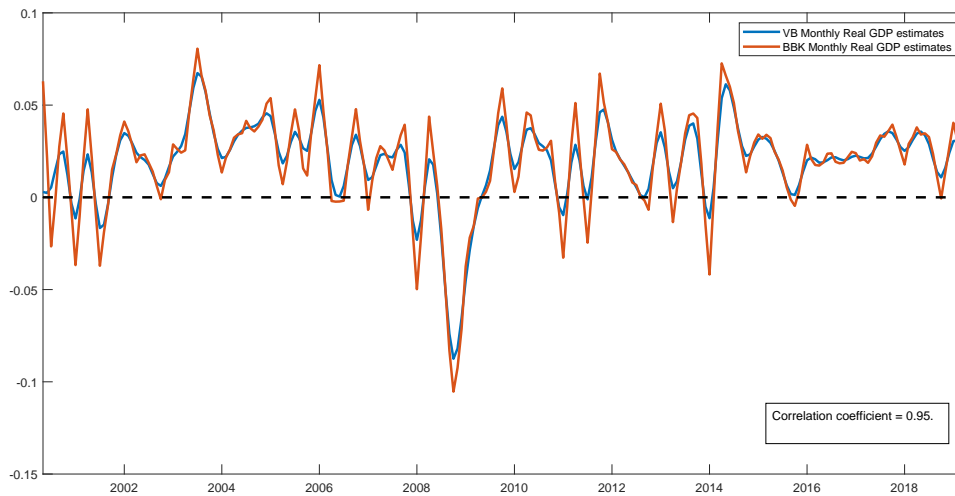


Figure 1: Monthly US GDP growth estimates: Comparison of VB vs BBK

6.2 Historical Quarterly Estimates of Regional Growth in the UK

In the UK, nominal Gross Value Added (GVA) is produced at the quarterly frequency for the UK as a whole, but is only produced annually for the 12 UK regions. Koop et al (2020) use a mixed frequency VAR with a Dirichlet-Laplace prior involving an annual-quarterly frequency mis-match to provide historical quarterly regional GVA estimates using MCMC methods. Here we repeat their analysis on the homoskedastic version of their model using VB methods. It is a 17 dimensional MF-VAR involving 12 annual regional GVA growth rates, quarterly UK GVA growth and four other quarterly UK predictors. Exact details of the data set and model are given in Koop et al (2020). Relative to the preceding sub-section which focussed on estimating a single quarterly variable, this is a more challenging empirical exercise due to the different frequency mis-match and the fact that high frequency estimates of many low frequency variables are required. But Koop et al (2020) found that, through the addition of an extra measurement equation which imposed the restriction that UK GVA is the sum of GVA for the regions, accurate estimation and good forecasting performance was achieved.

If repeat the MCMC-based empirical work of Koop et al (2020) using VB methods we obtain virtually identical results in a fraction of the time. Producing 20,000 draws from the MCMC algorithm took approximately five hours. Comparable VB estimation took 30 seconds. To illustrate the accuracy of VB methods, Figures 2 and 3 plot the historical estimates (posterior means) of quarterly GVA growth (annualized) for the regions. The high degree of similarity of MCMC and VB estimates can be seen. The correlations between VB and MCMC estimates are over 0.99 for 10 of the 12 regions. For the other two regions, the correlations are greater than 0.98.

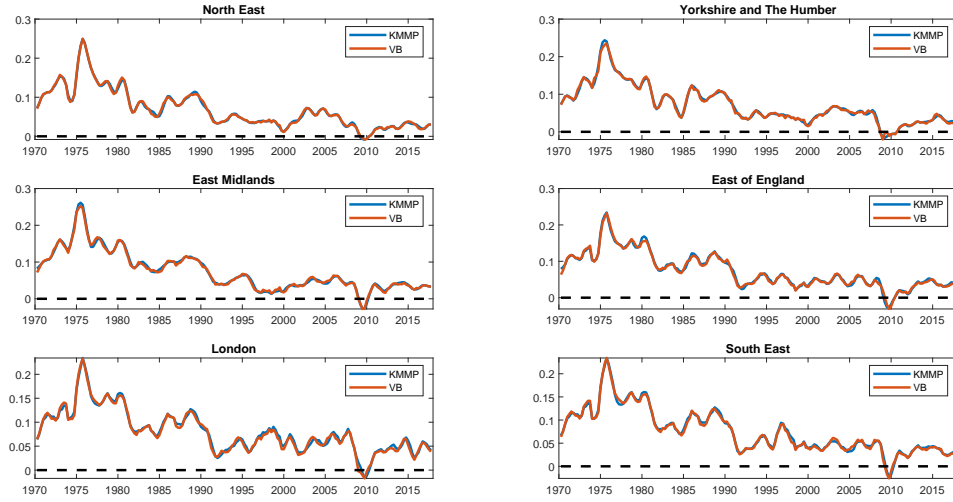


Figure 2: KMMP estimates, VB vs MCMC - Nominal GVA

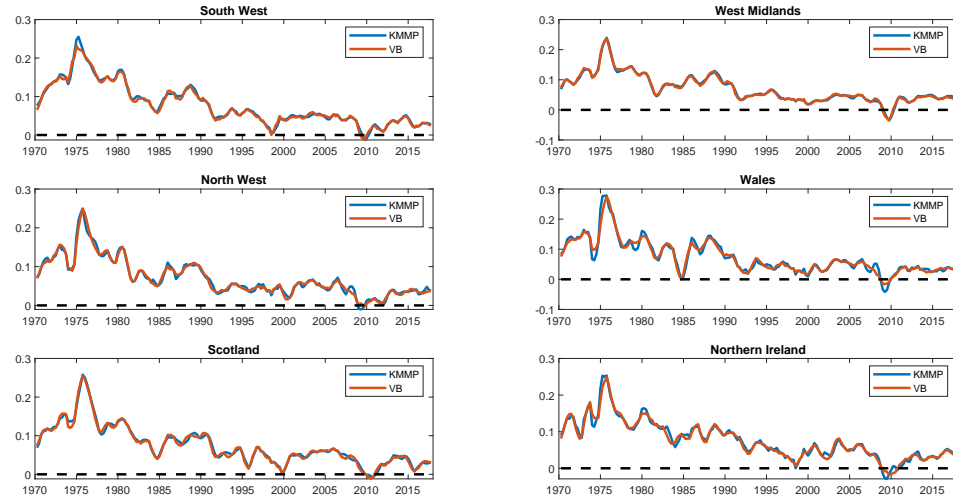


Figure 3: KMMP estimates, VB vs MCMC - Nominal GVA

7 Conclusions

This paper develops VB methods for MF-VARs. These are computationally fast and scaleable and, thus, can be used with large MF-VARs where MCMC-based methods would be impracticable. In two empirical exercises, we demonstrate the accuracy and computational efficiency of our methods.

References

- [1] Blei, D., Kucukelbir, A. and J. McAuliffe (2017). Variational inference: A review for statisticians. *Journal of the American Statistical Association* 112, 859-877.
- [2] Bhattacharya, A., Pati, D., Pillai, N. S., & Dunson, D. B. (2015). Dirichlet–Laplace priors for optimal shrinkage. *Journal of the American Statistical Association*, 110(512), 1479-1490.
- [3] Brave, S., Butters, R. and Justiniano, A. (2018). Forecasting economic activity with mixed frequency Bayesian VARs. *Federal Reserve Bank of Chicago Working Paper* 2016-05 (revised July 12, 2018).
- [4] Carriero, A., Clark, T. and Marcellino, M. (2015). Realtime nowcasting with a Bayesian mixed frequency model with stochastic volatility. *Journal of the Royal Statistical Society Series A* 178, 837-862.
- [5] Cotter, J., Hallam, M. and Yilmaz, K. (2017). Mixed frequency macro-financial spillovers. *Koc University Working Paper* 1704.
- [6] Eraker, B., Chiu, C., Foerster, A., Kim, T. and Seoane, H. (2015). Bayesian mixed frequency VAR's. *Journal of Financial Econometrics* 13, 698-721.
- [7] Gefang, D., Koop, G. and Poon, A. (2019). Variational Bayesian inference in large Vector Autoregressions with hierarchical shrinkage. *Economic Statistics Centre of Excellence discussion paper* 2019-07.
- [8] Gotz, T. and Hauzenberger, K. (2018). Large mixed-frequency VARs with a parsimonious time-varying parameter structure. *Deutsche Bundesbank Working Paper* No. 40/2018.
- [9] Ghysels, E. (2016). Macroeconomics and the reality of mixed frequency data. *Journal of Econometrics* 193, 294-314.
- [10] Hajargasht, G. and Wozniak, T. (2018). Variational Bayes inference for large vector autoregressions. Manuscript.
- [11] Koop, G. and Korobilis, D. (2018). Variational Bayes inference in high-dimensional time-varying parameter models. Manuscript available at <https://sites.google.com/site/garykoop/research>.
- [12] Koop, G., McIntyre, S., Mitchell, J. and Poon, A. (2020). Regional output growth in the United Kingdom: More timely and higher frequency estimates, 1970-2017. *Journal of Applied Econometrics*, doi.org/10.1002/jae.2748.
- [13] McCracken, M. and Ng, S. (2016). FRED-MD: A monthly database for macroeconomic research. *Journal of Business and Economic Statistics* 34, 574-589.
- [14] McCracken, M., Owyang, M. and Sekhposyan, T. (2016). Real-time forecasting with a large, mixed frequency Bayesian VAR, manuscript available at <http://www.tateviksekhposyan.org/>.
- [15] Schorfheide, F. and Song, D. (2015). Real-time forecasting with a mixed-frequency VAR. *Journal of Business and Economic Statistics* 33(3), 366-380.
- [16] Wang, H., Yu, H., Hoy, M., Dauwels, J. and Wang, H. (2016). Variational Bayesian dynamic compressive sensing, 2016 IEEE International Symposium on Information Theory.
- [17] Zhang, Y. and Bondell, H. (2018). Variable selection via penalized credible regions with DirichletLaplace global-local shrinkage priors. *Bayesian Analysis* 13, 823-844.

Data Appendix

All data for the UK regional application is exactly as described in Koop, McIntyre, Mitchell and Poon (2020). The data for the US application is taken from the popular FRED-MD depository (see <https://research.stlouisfed.org/econ/mccracken/fred-databases/>) and is summarized in the following table. All variables are transformed according to the recommendations made by FRED-MD.

Table 1: Data for US application

FRED Mnemonic	Data Frequency
PAYEMS	Monthly
CLF16OV	Monthly
CE16OV	Monthly
CLAIMS _x	Monthly
UNRATE	Monthly
AWHMAN	Monthly
DPCERA3M086SBEA	Monthly
USGOOD	Monthly
USCONS	Monthly
SRVPRD	Monthly
USTPU	Monthly
USWTRADE	Monthly
USTRADEx	Monthly
USFIRE	Monthly
USGOVT	Monthly
RPI	Monthly
W875RX1	Monthly
CMRMTSPL _x	Monthly
RETAIL _x	Monthly
INDPRO	Monthly
IPFPNSS	Monthly
IPFINAL	Monthly
IPCONGD	Monthly
IPDCONGD	Monthly
IPNCONGD	Monthly
IPBUSEQ	Monthly

Table 2: Data for US application

FRED Mnemonic	Data Frequency
IPMAT	Monthly
IPDMAT	Monthly
IPNMAT	Monthly
IPMANSICS	Monthly
IPB51222S	Monthly
IPFUELS	Monthly
CUMFNS	Monthly
AMDMNOx	Monthly
ANDENOx	Monthly
AMDMUOx	Monthly
BUSINVx	Monthly
ISRATIOx	Monthly
HOUST	Monthly
PERMIT	Monthly
UMCSENTx	Monthly
CPIAUCSL	Monthly
WPSFD49207	Monthly
PCEPI	Monthly
PPICMM	Monthly
WPSFD49502	Monthly
WPSID61	Monthly
WPSID62	Monthly
S&P 500	Monthly
S&P: industry	Monthly
GDPC1	Quarterly