

Huber, Florian; Koop, Gary; Onorante, Luca

Working Paper

Inducing sparsity and shrinkage in time-varying parameter models

Working Papers in Economics, No. 2019-02

Provided in Cooperation with:

Department of Social Sciences and Economics, University of Salzburg

Suggested Citation: Huber, Florian; Koop, Gary; Onorante, Luca (2019) : Inducing sparsity and shrinkage in time-varying parameter models, Working Papers in Economics, No. 2019-02, University of Salzburg, Department of Social Sciences and Economics, Salzburg

This Version is available at:

<https://hdl.handle.net/10419/201679>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Florian Huber, Gary Koop, Luca Onorante

Inducing Sparsity and Shrinkage in Time-Varying Parameter Models

No. 2019-02



WORKING PAPERS
IN ECONOMICS

Inducing Sparsity and Shrinkage in Time-Varying Parameter Models

Florian Huber*
University of Salzburg

Gary Koop†
University of Strathclyde

Luca Onorante ‡
European Central Bank

Abstract

Time-varying parameter (TVP) models have the potential to be over-parameterized, particularly when the number of variables in the model is large. Global-local priors are increasingly used to induce shrinkage in such models. But the estimates produced by these priors can still have appreciable uncertainty. Sparsification has the potential to remove this uncertainty and improve forecasts. In this paper, we develop computationally simple methods which both shrink and sparsify TVP models. In a simulated data exercise we show the benefits of our shrink-then-sparsify approach in a variety of sparse and dense TVP regressions. In a macroeconomic forecast exercise, we find our approach to substantially improve forecast performance relative to shrinkage alone.

Keywords: Sparsity, shrinkage, hierarchical priors, time varying parameter regression
JEL Codes: C11, C30, E30, D31

*Corresponding author. Email: florian.huber@sbg.ac.at. Any views expressed in this paper represent those of the authors only and not necessarily of the ECB or the Eurosystem. We are grateful to Niko Hauzenberger and Michael Pfarrhofer for helpful comments and suggestions. The first author acknowledges funding from the Austrian Science Fund (FWF): ZK 35.

†Email: gary.koop@strath.ac.uk

‡Email: Luca.Onorante@ecb.int

1 Introduction

Time-varying parameter (TVP) regressions and Vector Autoregressions (VARs) have enjoyed great popularity among econometricians in recent years as a way of modelling the parameter change that occurs in many macroeconomic and financial time series variables. These are state space models which have been found to work well in forecasting (e.g. D’Agostino et al., 2013) and been successfully used for structural economic analysis in a changing environment (e.g. Cogley and Sargent, 2005; Primiceri, 2005). They are flexible and capable of modelling almost any nonlinear relationship between explanatory and dependent variables. However, this flexibility comes at a cost: TVP models can be over-parameterized and suffer from the curse of dimensionality, particularly when the number of potential explanatory variables is large. This can lead to very good in-sample fit, but poor out-of-sample forecast performance.

There is a large and growing literature that proposes various methods for overcoming these over-parameterization concerns using Bayesian methods (see, among others, Frühwirth-Schnatter and Wagner, 2010; Belmonte et al., 2014; Kalli and Griffin, 2014; Kowal et al., 2017; Uribe and Lopes, 2017; Rockova and McAlinn, 2017; Koop and Korobilis, 2018; Bitto and Frühwirth-Schnatter, 2019; Huber et al., 2019; Eisenstat et al., 2019). These papers propose different approaches to obtain more precise inference. Much of this literature uses hierarchical global-local shrinkage priors. A key property of these priors is that they ensure shrinkage in the sense that they pull coefficients towards zero. However, they do not impose them to be exactly zero and, thus, estimation uncertainty remains. In contrast to shrinkage approaches, selection approaches seek to choose a single sparse specification. That is, they select a particular set of explanatory variables and, by doing so, impose coefficients on non-selected explanatory variables to be zero.¹

Which is better: shrinkage or sparsity? The answer to this question depends on the empirical application. In the case of constant coefficient regressions and VARs, there is debate among Bayesian econometricians as to whether models are sparse (in which case sparsification methods are appropriate) or dense (in which case shrinkage is appropriate). For instance, Giannone et al. (2017) considers a range of data sets in macroeconomics, microeconomics and finance and finds evidence mostly in favor of dense models, a finding reinforced by Cross et al. (2019). But why not do both? This is exactly what recent papers such as Hahn and Carvalho (2015) propose. That is, first shrinkage is done using a Bayesian global-local shrinkage prior and then sparsification is done on the resulting estimates. Such an approach could add the benefits of sparsity (i.e. the reduction in estimation error that is important for improving forecasts) along with the benefits of shrinkage which are so useful with dense data sets.

¹In the Bayesian literature, there are some global-local priors, such as the spike and slab prior, which do select variables, but these are less popular since Markov Chain Monte Carlo (MCMC) algorithms tend to mix poorly.

Until recently, one reason not to do both was computation. Bayesian inference with hierarchical shrinkage priors requires computationally-burdensome MCMC methods. Adding a second sparsification step can greatly increase the burden (i.e. this step often uses cross-validation methods for choosing key tuning parameters which can be computationally costly). However, a recent paper, Ray and Bhattacharya (2018), proposes a very simple algorithm, the signal adaptive variable selector (SAVS), for doing the sparsification step. This involves no tuning parameters and is computationally trivial. Ray and Bhattacharya (2018) provides a theoretical justification for SAVS and shows it to have good empirical performance in simulated and real data contexts.

The papers cited in the preceding paragraph all relate to constant coefficient regression or VAR models rather than the TVP state space models which are the focus of this paper. We develop Bayesian methods for inference and forecasting in TVP regressions and TVP-VARs which both shrink and sparsify. The shrinkage step can be done using any of the hierarchical shrinkage priors that have been used with TVP regressions. In this paper, we use the Dirichlet-Laplace prior (see Bhattacharya et al., 2015), a fully hierarchical variant of the stochastic search variable selection prior (see George and McCulloch, 1993; Ishwaran and Rao, 2005; George et al., 2008), the Horseshoe (see Carvalho et al., 2010), the Bayesian Lasso of Park and Casella (2008) and the Normal-Gamma prior of Griffin and Brown (2010). The sparsity step is done using the SAVS method of Ray and Bhattacharya (2018).

Another extension we make in this paper relative to the shrink-then-sparsify methods of Hahn and Carvalho (2015) and Ray and Bhattacharya (2018) is that we allow for uncertainty in the sparsified estimates. That is, Hahn and Carvalho (2015) and Ray and Bhattacharya (2018) take the posterior mean from the shrinkage step and use only this in the sparsification step. We sparsify every MCMC draw in the shrinkage step, thus allowing for parameter uncertainty. Our methods are illustrated with simulated and real data and we find them to improve estimation accuracy and forecast performance.

The remainder of this paper is organized as follows: Section 2 discusses various global-local shrinkage priors in the context of the regression model with constant coefficients. It describes how the sparsification strategy of Ray and Bhattacharya (2018) works in the regression model. Section 3 extends these methods to TVP regressions and TVP-VARs. Section 4 investigates the performance of our methods relative to non-sparsified alternatives using simulated data from a range of sparse and dense TVP regressions. Section 5 carries out a forecasting exercise using TVP-VARs. A comparison of forecasts which are both shrunk and sparsified to those which are only shrunk shows the benefits of doing both. Section 6 concludes the paper and a technical appendix provides further details on the specific prior setup and the posterior simulation algorithms.

2 Shrinkage and Sparsification in Regression Models

In this section we describe the shrinkage and sparsification methods for regression which we build on in this paper. In the next section, we will show how they can be adapted for dynamic regressions and multiple equation models such as VARs. Consider the regression model:

$$y_t = \boldsymbol{\beta}' \mathbf{X}_t + \varepsilon_t, \quad (1)$$

for $t = 1, \dots, T$ where y_t is a scalar dependent variable, \mathbf{X}_t is a $K \times 1$ vector of explanatory variables and $\boldsymbol{\beta}$ is a K -dimensional vector of regression coefficients. The errors are assumed to be independent and follow a zero mean Gaussian distribution with variance σ_ε^2 .

When K is large relative to T , Bayesians increasingly use hierarchical priors so as to induce shrinkage. Global-local shrinkage priors are particularly popular (see, e.g., Polson and Scott, 2010). These contain shrinkage that is both global (i.e. common to all parameters) and local (i.e. specific to each parameter). We consider priors which can be represented as scale mixtures of Gaussians. In particular, for the j^{th} regression coefficient we assume:

$$\beta_j \sim \mathcal{N}(0, \phi_j \lambda), \quad \phi_j \sim f, \quad \lambda \sim g. \quad (2)$$

Global shrinkage is controlled by λ while ϕ_j handles the shrinkage of coefficient j . f and g are mixing densities and many different choices have been proposed for them. In this paper, we consider the Horseshoe (HS) prior of Carvalho et al. (2010), the Bayesian Lasso (Lasso) of Park and Casella (2008), the Normal-Gamma (NG) prior of Griffin and Brown (2010), the Dirichlet-Laplace (DL) prior of Bhattacharya et al. (2015) and the Normal- mixture of Inverse Gamma (NMIG) prior of Ishwaran and Rao (2005), which is a variant of the stochastic search variable selection (SSVS) prior of George and McCulloch (1993, 1997). All of these are global-local shrinkage priors and differ from one another only in the choices of f and g . In addition, and unless otherwise noted, we use a weakly informative inverted Gamma prior on σ_ε^2 with hyperparameters $d_\sigma = e_\sigma = 0.01$.

Using any of these global-local shrinkage priors, MCMC methods can be used to carry out posterior inference and calculate the posterior mean, $\hat{\boldsymbol{\beta}}$. This estimate has been shrunk, but not sparsified. Even though many elements of $\hat{\boldsymbol{\beta}}$ will be near zero, they will not be precisely zero and there will be estimation uncertainty associated with them. Sparsification, as used in Hahn and Carvalho (2015) and Ray and Bhattacharya (2018), proceeds by taking $\hat{\boldsymbol{\beta}}$ and setting small elements in it to zero. We first define the SAVS estimate and then offer some explanation and

motivation for it. The SAVS estimate is

$$\bar{\gamma}_j = \text{sign}(\hat{\beta}_j) \|\mathbf{X}_j\|^{-2} \max\left(|\hat{\beta}_j| \|\mathbf{X}_j\|^2 - \kappa_j, 0\right), \quad (3)$$

with \mathbf{X}_j denoting the j^{th} column of a $T \times K$ matrix $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_T)'$, $(x)_+ = \max(x, 0)$ and $\text{sign}(x)$ returns the sign of x . Note that this is a soft-thresholding approach where all values of $\bar{\gamma}_j$ below a certain value are set to zero and that it only acts on the posterior mean.

The sparsified estimate depends on tuning parameters, κ_j , which determine the thresholds for each coefficient. Various approaches to selecting these have been proposed in the literature including computationally-intensive approaches such as cross-validation. However, a recent paper, Ray and Bhattacharya (2018), comes up with a surprisingly simple solution. This is to set:

$$\kappa_j = \frac{1}{|\hat{\beta}_j|^2}. \quad (4)$$

This choice implies a penalty for the j^{th} variable which is ranked in inverse-squared order relative to the magnitude of the j^{th} coefficient. With this choice of thresholds, the SAVS estimate is trivial to calculate.

To provide some motivation for the SAVS estimate note that (3) can be obtained by first solving an optimization problem akin to the Lasso:

$$\bar{\gamma} = \arg \min_{\gamma} \left\{ \frac{1}{2} \|\mathbf{X}\hat{\beta} - \mathbf{X}\gamma\|_2^2 + \sum_{j=1}^K \kappa_j |\gamma_j| \right\}. \quad (5)$$

Equation (5) tries to find a sparse coefficient vector γ that is close to $\hat{\beta}$ while introducing a penalty in case of non-zero elements in γ .

The typical way to solve this optimization problem is using a coordinate descent algorithm (Friedman et al., 2007). But, as shown in Ray and Bhattacharya (2018), if you initialize this algorithm at $\hat{\beta}$ and then do one iteration you get precisely the simple algorithm described in (3) and (4). It is also noted in Ray and Bhattacharya (2018) that convergence almost always occurs after one iteration and, hence, stopping after one iteration is a sensible thing to do.

One key shortcoming of computing the SAVS estimate is that uncertainty quantification about $\bar{\gamma}$ is not possible and computing non-linear functions of $\bar{\gamma}$ calls for Monte Carlo integration techniques. Ray and Bhattacharya (2018) highlight that one potential solution to this issue is to replace $\hat{\beta}$ with a draw from the full conditional posterior distribution of β . This is an insight we build upon in the context of the TVP models which are the focus of this paper.

3 Shrinkage and Sparsification in TVP Models

In this section, we develop methods for shrinkage and sparsification in state space models such as the TVP regression and the TVP-VAR. This is achieved using the non-centered parameterization of Frühwirth-Schnatter and Wagner (2010). We emphasize that the algorithms below do the sparsification at each draw from the MCMC algorithm, allowing for treatment of uncertainty in the shrinkage step. Thus, the algorithms are averaging over different sparsified estimators in a manner similar to Bayesian model averaging.

3.1 The TVP Regression Model

The TVP regression model used in this paper takes the form:

$$\begin{aligned} y_t &= \boldsymbol{\beta}'_t \mathbf{X}_t + \varepsilon_t \\ \boldsymbol{\beta}_t &= \boldsymbol{\beta}_{t-1} + \mathbf{w}_t, \end{aligned}$$

where all definitions are the same as in (1) except that $\boldsymbol{\beta}_t$ are dynamic (time-varying) regression coefficients which follow a random walk with $\mathbf{w}_t \sim \mathcal{N}(\mathbf{0}_K, \mathbf{V})$, where $\mathbf{V} = \text{diag}(v_1, \dots, v_K)$ denotes the variance-covariance matrix of the state innovations while v_j ($j = 1, \dots, K$) is a process innovation variance associated with the j^{th} coefficient.

The non-centered parameterization of this model is given by:

$$\begin{aligned} y_t &= \boldsymbol{\beta}'_0 \mathbf{X}_t + \tilde{\boldsymbol{\beta}}'_t \sqrt{\mathbf{V}} \mathbf{X}_t + \varepsilon_t, \\ \tilde{\boldsymbol{\beta}}_t &= \tilde{\boldsymbol{\beta}}_{t-1} + \boldsymbol{\eta}_t, \quad \boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0}_K, \mathbf{I}_K). \end{aligned}$$

with the j^{th} element of $\tilde{\boldsymbol{\beta}}_t$ given by $\tilde{\beta}_{jt} = \frac{\beta_{jt} - \beta_{j0}}{\sqrt{v_j}}$ and $\sqrt{\mathbf{V}} = \text{diag}(\sqrt{v_1}, \dots, \sqrt{v_K})$. This equation can be written as:

$$y_t = \boldsymbol{\alpha}' \mathbf{Z}_t + \varepsilon_t, \tag{6}$$

whereby $\boldsymbol{\alpha} = (\boldsymbol{\beta}'_0, \sqrt{v_1}, \dots, \sqrt{v_K})'$, $\mathbf{Z}_t = [\mathbf{X}'_t, (\tilde{\boldsymbol{\beta}}_t \odot \mathbf{X}_t)']'$ and \odot denotes element-wise multiplication. Conditional on knowing the full history of the states in $\tilde{\boldsymbol{\beta}}_t$, (6) resembles a standard regression model with a (partially) latent covariate vector \mathbf{Z}_t .

Well-developed MCMC methods exist to carry out Bayesian posterior and predictive inference in state space models such as the TVP regression model under various priors. In this paper, we simulate the full history of the normalized dynamic regression coefficients $\{\tilde{\boldsymbol{\beta}}_t\}_{t=1}^T$ using the forward-filtering backward-sampling algorithm

proposed in Carter and Kohn (1994) and Frühwirth-Schnatter (1994). Conditional on $\tilde{\boldsymbol{\beta}}_t$, (6) is a standard regression model, implying that we can simulate $\boldsymbol{\alpha}$ from a Gaussian full conditional posterior distribution and σ_ε^2 from an inverted Gamma distribution. The corresponding moments take standard forms and are presented in Appendix B.

We propose to do shrinkage on $\boldsymbol{\alpha}$ using the global-local mixture priors mentioned in the previous section and described in Appendix A. That is, conditional on a draw of the full history of the states, $\{\tilde{\boldsymbol{\beta}}_t\}_{t=1}^T$, we have the regression model given in (6), and shrinkage can be done exactly as described in the preceding section. For each of the global-local mixture priors we consider, MCMC methods for drawing $\boldsymbol{\alpha}$ and σ_ε^2 , conditional on draws of the states exist. For the Dirichlet-Laplace prior we follow the methods of Bhattacharya et al. (2015). For the NMIG specification, we adopt the algorithm proposed in Ishwaran and Rao (2005) while for the Horseshoe, the MCMC algorithm developed in Makalic and Schmidt (2016) is used. Since the Normal-Gamma prior nests the Bayesian Lasso, we adopt the algorithm put forth in Griffin and Brown (2010) (see Appendix A for further details).

Every draw of $\boldsymbol{\alpha}$, denoted as $\boldsymbol{\alpha}^{(n)}$, from any of the MCMC algorithms is sparsified using SAVS. Applying the SAVS estimator in (3) to each draw from the posterior of $\boldsymbol{\alpha}$ yields:

$$\bar{\gamma}_j^{(n)} = \text{sign} \left(\alpha_j^{(n)} \right) \|\mathbf{Z}_j\|^{-2} \left(|\alpha_j^{(n)}| \|\mathbf{Z}_j\|^2 - \kappa_j \right)_+, \text{ for } n = 1, \dots, N \quad (7)$$

where \mathbf{Z}_j denotes the j^{th} column of $\mathbf{Z} = (\mathbf{Z}_1, \dots, \mathbf{Z}_T)'$, $\kappa_j = |\alpha_j^{(n)}|^{-2}$ and N denotes the number of post burn-in MCMC draws. This procedure effectively allows for uncertainty quantification and the computation of potentially non-linear functions of the sparsified parameters. Thus, one can think of our proposed procedure as an approximate MCMC algorithm which draws from the sparsified conditional posterior $p(\bar{\boldsymbol{\gamma}}|\boldsymbol{\alpha}, \mathbf{Z})$.² Hence, forecasts produced will average over different sparsified models. That is, one MCMC draw will lead to one particular sparsified model which is used for forecasting, another draw may choose another sparsified model to produce forecasts. Hence, what we are proposing is similar in spirit to Bayesian model averaging. Of course, it is possible to use the SAVS algorithm directly on the posterior mean of $\boldsymbol{\alpha}$ as is done by Hahn and Carvalho (2015) and Ray and Bhattacharya (2018). This would be similar in spirit to a Bayesian model selection strategy where a single sparsified model was chosen for forecasting. But this would ignore model uncertainty and, in addition, would be problematic since \mathbf{Z}_t is partially latent.

Another point worth emphasizing about our algorithm is that it is fast. Relative to the computational time required to do MCMC, adding the SAVS step increases the computational burden by a trivial amount. For any empirical specification where MCMC is possible, our proposed algorithm is also possible. Of course, if K is too

²The algorithm is approximate since σ_ε^2 does not play a role in the SAVS algorithm. If desired, after each sparsification, one could take a draw of σ_ε^2 conditional on the sparsified estimates.

large, then MCMC methods may be computationally infeasible. In such a case, variational Bayesian methods may be a practical alternative (see Koop and Korobilis, 2018). But with variational Bayes methods, the SAVS algorithm would be applied on the approximate posterior mean and model uncertainty ignored.³

3.2 The TVP-VAR

The shrink-then-sparsify algorithm we propose for the TVP regression can be extended to handle the TVP-VAR in a straightforward fashion. The idea is to transform the TVP-VAR so that the error covariance matrix in the measurement equation is diagonal. Then the TVP regression algorithm of the preceding sub-section can be applied one equation at a time. Equation-by-equation estimation of VARs is done in several recent papers using transformations similar to the one used here (see, e.g., Carriero et al., 2016; Kastner and Huber, 2017; Koop et al., 2019) and the reader is referred to these papers for further details about the computational advantages of this approach. With macroeconomic data it is often important to add stochastic volatility, which leads us to the TVP-VAR-SV specification described in this section.

Let \mathbf{y}_t be an $M \times 1$ vector of endogenous variables for $t = 1, \dots, T$. The TVP-VAR-SV can be written as:

$$\mathbf{y}_t = (\mathbf{I}_M \otimes \mathbf{X}_t)\boldsymbol{\beta}_t + \boldsymbol{\varepsilon}_t, \quad (8)$$

where $\mathbf{X}_t = (\mathbf{y}'_{t-1}, \dots, \mathbf{y}'_{t-P}, 1)'$ contains the P lags of \mathbf{y}_t and an intercept, $\boldsymbol{\beta}_t$ is the vector $K = M(MP + 1)$ coefficients at time t which is assumed to evolve according to a multivariate random walk. The errors are independent over time with $\boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}_M, \boldsymbol{\Sigma}_t)$. $\boldsymbol{\Sigma}_t$ is the time-varying error covariance matrix with

$$\boldsymbol{\Sigma}_t = \mathbf{U}_t \mathbf{H}_t \mathbf{U}_t'.$$

Let \mathbf{U}_t denote a lower uni-triangular matrix and $\mathbf{H}_t = \text{diag}(e^{h_{1t}}, \dots, e^{h_{Mt}})$. The $M(M-1)/2$ free elements in \mathbf{U}_t follow independent random walks while the h_{jt} 's are log-volatilities that follow AR(1) processes,

$$h_{jt} = \mu_j + \rho_j(h_{jt-1} - \mu_j) + \eta_t, \quad \eta_t \sim \mathcal{N}(0, \sigma_\eta^2). \quad (9)$$

Here, we let μ_j denote the unconditional mean, ρ_j the persistence parameter and σ_η^2 the error variance of the log-volatility process. The prior specification on the parameters of the log-volatility equation closely follows Kastner and

³It would be possible to surmount this drawback of variational Bayes by first using variational Bayes to obtain an approximation to the posterior and then applying the SAVS algorithm to draws from this approximation. But this would be computationally demanding, thus undermining the main advantage of variational Bayes.

Frühwirth-Schnatter (2014). Specifically, we use a zero mean Gaussian prior with variance 10^2 on μ_j , a Beta prior on $\frac{\rho_i+1}{2} \sim \mathcal{B}(25, 5)$ and a Gamma prior on $\sigma_\eta^2 \sim \mathcal{G}(1/2, 1/2)$. The Gamma prior translates into a Gaussian prior on $\pm\sigma_\eta$ with zero mean and unit variance. In the MCMC algorithm, the full history of h_{jt} as well as the parameters of equation (9) are obtained using the algorithm proposed in Kastner and Frühwirth-Schnatter (2014). This algorithm exploits the centered and non-centered parameterization of the non-linear state space model to increase sampling efficiency and samples the full history of the log-volatilities from a $(T - 1)$ -dimensional multivariate Gaussian.

As noted in Carriero et al. (2016); Kastner and Huber (2017); Koop et al. (2019), computation is greatly simplified if the model is transformed so that the errors in different equations are independent of one another. This can be achieved by augmenting the i^{th} equation in the system with the contemporaneous values of the first $i - 1$ elements in \mathbf{y}_t . That is, if y_{it} is the i^{th} variable (for $i > 1$), we can write the TVP-VAR-SV as a set of M unrelated TVP regressions:

$$y_{it} = \mathbf{X}'_t \boldsymbol{\beta}_{it} + \sum_{j=1}^{i-1} u_{ij,t} y_{jt} + \eta_{it}, \quad \eta_{it} \sim \mathcal{N}(0, e^{h_{it}}), \quad (10)$$

where η_{it} and η_{jt} are independent for $i \neq j$, $\boldsymbol{\beta}_{it}$ denotes the elements of $\boldsymbol{\beta}_t$ in the i^{th} equation and $u_{ij,t}$ are the elements of \mathbf{U}_t^{-1} for $i = 2, \dots, M; j = 1, \dots, i - 1$.

We then write the TVP-VAR-SV using the non-centered parameterization. For equation i we obtain:

$$y_{it} = \mathbf{X}'_t \boldsymbol{\beta}_{i0} + (\sqrt{\mathbf{V}_i^\beta} \mathbf{X}_t)' \tilde{\boldsymbol{\beta}}_{it} + \sum_{j=1}^{i-1} u_{ij,0} y_{jt} + \sum_{j=1}^{i-1} \sqrt{v_{ij}^u} \tilde{u}_{ij,t} y_{jt} + \eta_{it}, \quad \eta_{it} \sim \mathcal{N}(0, e^{h_{it}}). \quad (11)$$

Here, we let $\sqrt{\mathbf{V}_i^\beta} = \text{diag} \left(\sqrt{v_{i1}^\beta}, \dots, \sqrt{v_{iK}^\beta} \right)$ and $\sqrt{v_{ij}^u}$ denotes the standard deviation of the error in the random walk state equation for the j^{th} VAR coefficient in the i^{th} equation. Similarly, $\sqrt{v_{ij}^u}$ is the standard deviation for the random walk state equation for the elements of \mathbf{U}_t . Thus, $\tilde{\boldsymbol{\beta}}_{it}$ and $\tilde{u}_{ij,t}$ are the states for equation i and the shocks in the corresponding state equations have unit standard deviation.

Since the errors in the different equations are independent of one another, estimation of one equation at a time using the algorithm of the preceding sub-section, including the SAVS step detailed in (7), can be done. Computation is also sped up since parallelization is feasible. Note also that, since the coefficients in \mathbf{U}_t^{-1} are appearing as regression coefficients in (10), these can also be shrunk and sparsified. In large TVP-VARs, where there are many such error covariance terms, this is potentially beneficial for forecasting purposes. Notice that we do not only obtain a sparse error covariance matrix but also allow for testing whether the corresponding free elements are time-varying or constant.

4 Evidence Using Artificial Data


In this section, we present evidence on the performance of the proposed methodology using artificial data generated from different TVP regression models. Across the different data generating processes (DGPs), the covariates are drawn from a Uniform distribution bounded between -1 and 1 . The β_t 's are generated using the non-centered parameterization with $\beta_0 \sim \mathcal{N}(\mathbf{0}_K, 0.1^2 \mathbf{I}_K)$ and $\pm\sqrt{v_j} \sim \mathcal{N}(0, 0.1^2), j = 1, \dots, K$, while differing percentages of the elements in α are randomly set to zero. The measurement error variance σ_ε^2 is set equal to 0.1^2 .

Before presenting results using repeated samples, the main features of sparsification are illustrated in Figure 1. The results in the three panels of the figure are based on the Horseshoe prior and use three different single artificial data sets obtained by simulating $T = 400$ observations from a large ($K = 30$) dynamic regression model. Figure 1(a) plots posterior features of β_{jt} against time for a case where it is zero (i.e. the DGP is one where j^{th} regressor is not selected) using a non-sparsified and sparsified estimator. Note that the sparsified estimator is precisely correct, it sets $\beta_{jt} = 0$ with probability one. Thus, it exactly coincides with the true value and cannot be seen in Figure 1(a). The non-sparsified posterior distribution, although the posterior mean is very close to the correct value, has a credible interval that is non-negligible. This reflects estimation uncertainty and could spill over into poor forecast performance using the non-sparsified posterior. The performance of the SAVS algorithm when β_{jt} is a non-zero constant (i.e. the DGP is one where $\beta_{jt} = \beta_{jt-1}$ for all t) is shown in Figure 1(b). In this case, the posterior distributions of the sparsified and non-sparsified models almost coincide. Notice, however, that the credible sets are constant over time for the sparsified model, indicating that the corresponding element in \sqrt{V} is set equal to zero throughout all iterations of the MCMC algorithm. In contrast, Figure 1(c) illustrates a case where β_{jt} is non-zero and time-varying. Notice that the sparsified and non-sparsified posterior distributions are close to being identical. In this case, it is not desirable to sparsify the corresponding elements in α and the SAVS algorithm is not doing so. Thus, regardless of whether a coefficient is zero, a non-zero constant or time-varying, this figure indicates that our methods estimate it well. They work better than the non-sparsified alternative in cases where there is sparsity and equally well in non-sparse cases.

Table 1 presents evidence for the importance of sparsification and shrinkage in TVP regression models using different data configurations, priors, numbers of regressors and sample sizes. The DGP described above is modified to reflect varying degrees of sparsity. These different sparsity levels are labeled sparse (with 90% zeros in α), moderate (with 70% zeros) and dense (with 30% zeros). To assess how our techniques perform across model dimensions and length of time series involved, we consider variants of each sparsity level with $K = 5, 15$ and 30 explanatory variables and $T = 250$ and 400 observations. The latter are typical values in quarterly and monthly

Figure 1: Sparsified and non-sparsified posteriors of β_{jt} for a large time-varying parameter model

(a) DGP with $\beta_{jt} = 0$ for all t (b) DGP with $\beta_{jt} = \beta_{jt-1}$ for all t (c) DGP with β_{jt} time-varying



Notes: The dark blue line denotes the true β_{jt} over time, the gray shaded area represents 5th and 95th credible sets of the posterior of β_{jt} under a HS prior. Purple lines represent 5th and 95th credible sets of the sparsified posterior.

macroeconomic data sets. For each DGP, we generate 100 artificial data sets and then run each through a sparsified and non-sparsified algorithm using each of the five global-local shrinkage priors listed in Section 2. We also include a non-informative prior (labelled Flat in the table) which does not do shrinkage. Posterior medians of β_t are produced and the absolute value of the difference between these and the true value used in the DGP is calculated. The figures in Table 1 are averages taken over three dimensions: i) the 100 simulated data sets, ii) time and iii) the K elements of β_t .

Table 1 shows the value of sparsification, particularly with sparse DGPs. With the latter, mean absolute errors (MAEs) are lower than their non-sparsified counterpart for every prior and choice for T and K . But even in moderately dense specifications, sparsification lowers MAEs in most cases. In the dense specification, sparsification does not improve upon the single best performing non-sparsified model specification. However, in that situation, accuracy differences are found to be negligible.

The benefits of shrinking and sparsifying increase with the number of explanatory variables. Note, for instance, that the unsparsified Flat prior model does not perform that poorly when $K = 3$ and 15, but displays a weak performance when $K = 30$. In fact, when $K = 3$, the Flat prior works quite well with the sparse specification, provided sparsification is done. This indicates that there are some cases where sparsification is more important than shrinkage.

Table 1: Mean absolute errors between the true realization of the β_t s and the posterior median

Sparsity level	Non-sparsified					Sparsified						
	Flat	DL	Lasso	NG	HS	NMIG	Flat	DL	Lasso	NG	HS	NMIG
Small ($K = 5$)												
$T = 250$												
Dense	7.97	7.51	8.08	7.33	7.65	7.91	7.85	7.91	8.33	7.48	7.72	8.28
Moderate	4.66	4.59	4.47	4.97	4.27	3.62	4.06	4.53	4.29	5.11	4.30	3.75
Sparse	3.63	3.05	3.46	2.73	3.48	3.33	2.84	2.96	3.04	2.58	3.15	3.17
$T = 400$												
Dense	7.57	7.14	6.89	6.40	7.65	7.03	7.45	7.65	7.15	6.42	7.91	7.46
Moderate	4.41	4.10	4.27	5.13	3.89	4.42	3.94	4.08	4.21	5.05	4.17	4.53
Sparse	2.90	2.74	3.40	2.80	2.44	3.43	2.26	2.64	2.93	2.70	2.19	3.23
Medium ($K = 15$)												
$T = 250$												
Dense	11.60	10.88	9.51	9.39	8.72	9.76	11.32	11.78	9.18	9.24	8.83	10.07
Moderate	3.46	3.98	5.22	3.27	3.64	3.03	2.29	3.48	4.84	2.88	3.39	2.90
Sparse	3.21	1.81	2.44	2.09	1.84	2.04	2.02	1.40	1.48	1.79	1.63	1.55
$T = 400$												
Dense	9.87	10.77	9.60	8.19	8.44	8.72	9.61	11.50	9.70	8.07	8.95	9.41
Moderate	2.93	3.09	3.78	3.62	3.86	3.83	1.99	2.72	3.00	3.20	3.67	3.57
Sparse	2.30	2.14	2.37	1.75	1.27	1.90	1.23	1.89	1.96	1.51	1.08	1.39
Large ($K = 30$)												
$T = 250$												
Dense	15.40	14.18	14.30	12.84	13.97	12.32	15.04	15.14	13.64	12.88	14.06	12.64
Moderate	5.24	3.48	3.84	3.26	2.44	2.38	4.27	2.83	2.74	2.64	2.04	2.11
Sparse	2.53	1.33	2.48	1.36	1.72	1.67	1.17	0.79	1.59	0.97	1.59	1.07
$T = 400$												
Dense	13.71	12.48	12.29	13.50	12.39	11.68	13.33	13.27	11.83	13.22	12.17	12.06
Moderate	4.73	3.43	3.71	4.06	2.64	2.27	3.73	2.79	2.69	3.92	2.38	1.76
Sparse	1.78	1.36	1.47	1.32	0.52	1.43	0.63	0.75	0.60	0.89	0.44	0.78

Notes: The mean is taken over time, over all parameters and over all artificial data sets. All mean absolute errors are multiplied by 100. Flat refers to a dynamic regression model with a loosely informative prior, DL to the Dirichlet-Laplace prior, Lasso to the Bayesian Lasso, NG to the Normal-Gamma prior, HS to the Horseshoe, and NMIG to the Normal-mixture of Inverse Gamma prior.

The choice of T has little impact on the results. In a regression model with constant parameters, we would expect sparsification to be less important as the sample size increases since, with longer time series, the estimation error would decrease. However, with TVP regressions, the number of parameters is also increasing with the sample size which negates this effect. Thus, even with large numbers of observations, the researcher working with TVP models can still benefit from sparsification.

With regards to the different global-local shrinkage priors, no clear pattern emerges where one performs consistently the best across different specifications. When $K = 30$ and the DGP is sparse, DL (for $T = 250$) and HS (for $T = 400$) models that are sparsified are the best performers. When $K = 30$ and the DGP is dense, the accuracy of both, the DL and the HS prior deteriorates slightly while the unsparsified NMIG model shows the best performance.

Notice that in this situation, accuracy differences across the sparsified and non-sparsified NMIG specification are, however, quite small.

From this discussion it is apparent that identifying a default prior choice is difficult. One key take away from this analysis, however, is that if the DGP is sparse, flexible shrinkage specifications such as the HS, the DL and the NMIG prior in combination with the SAVS algorithm provide accurate parameter estimates. Overall, the table tells a story of the importance of both shrinkage and sparsity, especially in large models, with the precise choice of shrinkage prior being of lesser importance.

In the next step, we assess how well the SAVS algorithm identifies true zeros in α . Table 2 shows average hit rates that measure the percentage of correctly estimated zeros using the SAVS algorithm. From this table, we observe that irrespective of the priors used, our approach works well in identifying the true level of sparsity. For sparse situations, the fraction of correctly identified zeros is often above 95% for most shrinkage priors and model sizes considered. In the case of a Flat prior, we observe values just above 90%, which is remarkable but still well below the percentages observed for the different shrinkage specifications under scrutiny. This slightly weaker performance can be traced back to the fact that without shrinkage, values in α are not pushed to zero and the corresponding penalty κ_j is too small. Consistent with the findings in Table 1, we find no discernible differences in performance across the different shrinkage priors, with all of them displaying a strong performance. In fact, in a sparse setting with $K = 30$, the SAVS algorithm identifies almost all zeros correctly, with hit rates being above 99%.

To sum up, this discussion highlights that sparsification improves estimation accuracy. These improvements tend to increase with model size and the level of sparsity of the DGP. Among the set of competing shrinkage priors, we find no single best performing specification. In terms of correctly predicting zeros in α , we found that SAVS works well across all shrinkage priors considered, often correctly identifying above 99% of the zeros. At this point, and before proceeding to the empirical application, it is worth emphasizing that our analysis only considers whether our shrink-then-sparsify approach improves accuracy of point estimates, ignoring a potential bias-variance tradeoff. One key finding is that applying SAVS never significantly decreases estimation accuracy and correctly predicts a large fraction of true zeros. In light of Figure 1, this indicates that, by zeroing out shrunk coefficients, our approach strongly reduces parameter uncertainty.

Table 2: Average hit rates of the SAVS algorithm across different prior specifications

Sparsity level	Flat	DL	Lasso	NG	HS	NMIG
Small ($K = 5$)						
$T = 250$						
Dense	86.50	73.60	81.30	79.90	78.70	79.60
Moderate	88.40	93.10	92.00	94.10	91.80	91.40
Sparse	90.40	96.10	96.40	96.70	96.40	94.00
$T = 400$						
Dense	85.60	74.80	80.20	83.10	81.50	74.90
Moderate	89.60	92.90	93.40	94.90	93.40	91.10
Sparse	91.70	97.40	96.10	97.60	97.30	95.60
Medium ($K = 15$)						
$T = 250$						
Dense	79.17	77.63	79.37	80.03	79.07	78.23
Moderate	83.93	97.97	95.70	97.87	96.50	96.47
Sparse	91.23	99.17	98.00	99.17	99.47	98.97
$T = 400$						
Dense	82.57	80.40	82.27	82.97	80.73	78.33
Moderate	87.00	98.33	96.47	97.70	96.73	96.67
Sparse	93.97	99.67	98.83	99.57	99.73	98.97
Large ($K = 30$)						
$T = 250$						
Dense	68.57	72.57	76.05	76.05	76.47	75.38
Moderate	72.40	98.83	95.42	97.75	96.93	97.13
Sparse	85.12	99.70	97.73	99.62	99.62	99.55
$T = 400$						
Dense	71.30	78.50	78.35	79.23	77.80	76.37
Moderate	80.68	99.17	96.32	98.57	97.17	97.38
Sparse	91.60	99.82	98.35	99.80	99.65	99.53

Notes: The mean is computed over all parameters and artificial data sets. Flat refers to a dynamic regression model with a loosely informative prior, DL to the Dirichlet-Laplace prior, Lasso to the Bayesian Lasso, NG to the Normal-Gamma prior, HS to the Horseshoe, and NMIG to the Normal-mixture of Inverse Gamma prior.

5 Forecasting US Macroeconomic Variables

In this section we present results from a forecasting exercise using US quarterly macroeconomic data taken from the FRED-QD database (see McCracken and Ng, 2016) that span the period from 1959Q1 to 2017Q4. We focus on forecasting GDP, inflation (based on the GDP deflator) and the Fed Funds rate (henceforth labeled focus variables). Table 4 provides details on the specific variables included alongside transformations used.

We use the TVP-VAR-SV of Sub-section 3.2 combined with the same set of global-local shrinkage priors as in the preceding section. The only specification we do not consider here is the TVP-VAR-SV with a flat prior since this model performs poorly in out-of-sample forecasting and large dimensions.

For each prior, we use non-sparsified and sparsified versions of the model in order to produce the forecasts. We forecast with small ($M = 3$), medium ($M = 8$) and large ($M = 20$) data sets and set the lag length equal to 2. Thus, the dimension of the state space in the TVP-VAR-SVs ranges from being moderate to huge. Our forecast evaluation begins in 1997Q1 and runs to the end of the sample. We use root mean squared forecast errors (RMSEs) to evaluate the quality of the point forecasts and average log predictive likelihoods (LPLs) to evaluate the quality of predictive densities. Both are benchmarked relative to a VAR-SV with DL prior, a specification that works well for US macroeconomic data (see Kastner and Huber, 2017). This is identical to the TVP-VAR-SV with DL prior except that the DL prior now applies directly to the constant VAR coefficients while $\sqrt{\mathbf{V}_i^\beta}$ and $\sqrt{v_{ij}^u}$ are set equal to zero for all i, j . The VAR is transformed to allow for equation-by-equation estimation as described in Sub-section 3.2.

Before presenting the results of our forecasting exercise, we present Figure 2 which sheds light on which variables our algorithm is choosing to predict the focus variables. This figure is produced using the large data set and the HS prior. Previously, we have discussed how doing sparsification for each MCMC draw shares similarities with Bayesian model averaging. For one draw a certain sparsified model will be used to forecast, at another draw a different model will be used. This feature allows us to calculate posterior inclusion probabilities (PIPs) for each variable. The PIP for a given coefficient is the proportion of MCMC draws for which the coefficient is not set to zero. Figure 2 is a heatmap of these PIPs at the end of the sample. Remember that, in the non-centered parameterization of the TVP-VAR-SV (see equation (11)), there are coefficients which appear on the initial states which are constant coefficients. The upper panel of the figure relates to these. The remaining coefficients determine whether there is time-variation relative to the constant coefficients. The lower half of the figure relates to these.

Figure 2 shows that our methods are inducing a high degree of sparsity in the TVP-VAR-SV in that most of the PIPs are near zero. However, a few of them are not. In terms of the VAR coefficients there is only one coefficient which is always selected (i.e. has a PIP of one). This is the first lag of the 1-year Treasury Bill rate in the equation for the Fed Funds rate. However, an appreciable number of other predictors have PIPs that are substantially above zero but much less than one. In terms of the error covariance matrix, a similar pattern emerges. There is only one error covariance term which is non-zero in every MCMC draw.⁴ This is the covariance between the errors in the equations for two different inflation measures. However, there are several other error covariances with PIPs that are substantially above zero, even if they are below one. We stress that such a finding would not be possible if we were to use the SAVS algorithm directly on the posterior mean as opposed to using it on each MCMC draw. In the former case every PIP would be either zero or one with no values in between.

⁴Note that the other green areas refer to the diagonal elements of \mathbf{U}_t .

These patterns are consistent with those found in Giannone et al. (2017) who conclude *"there seems to be a lot of uncertainty about whether certain predictors should be included in the model, which results into their selection only in a subset of the posterior draws. These findings reflect a substantial degree of collinearity among many predictors that carry similar information, hence complicating the task of structure discovery. In sum, model uncertainty is pervasive and the best prediction is obtained as a weighted average of several models."* These features seem to be exactly what our algorithm is uncovering in an automatic fashion.


Finally, it is worth noting that there is evidence of time variation in several of the coefficients and our algorithm is automatically deciding which ones to allow to be time-varying. That is most of the PIPs which are appreciably above zero in the top half of the figure are also above zero in the bottom half. This pattern indicates a non-zero coefficient zero which is time varying. But our method also allows for a coefficient to be non-zero but constant. There are some cases which provide evidence of this. For instance, in the GDP growth equation the first lag of SP500 stock returns has a PIP which is appreciably above zero in the top half of the figure, but is much closer to zero in the bottom half of the figure. This pattern indicates support for a constant coefficient on this predictor.

The evidence in Figure 2 suggests that shrinking then sparsifying is working in a sensible fashion. But the key test of our methodology is how well it forecasts. Table 3 presents the results of our forecasting exercise. A comparison of each set of sparsified forecasts to its non-sparsified counterpart shows the benefits of our shrink-then-sparsify strategy, particularly in large models. For $M = 8$ and $M = 20$, sparsification leads to substantial improvements in both RMSEs and LPLs in almost every case. These improvements are particularly noticeable for GDP forecasting for the one-step-ahead forecasts. In general, the benefits of sparsification are largest when using the DL or Lasso priors. For $M = 3$ the benefits of sparsification are less pronounced. In terms of RMSEs, there seems to be no benefits of sparsification, although it does lead to slight improvements in the density forecasts even for this already fairly parsimonious case. This smaller accuracy premium from sparsification can be traced back to the fact that, in small models, the increases in the predictive variance that arise from posterior uncertainty surrounding shrunk estimates are small relative to the variance contribution arising from the reduced-form shocks. In larger models, parameter uncertainty eventually accumulates and this seems to be detrimental for forecasting accuracy.


In relation to the benchmark VAR-SV model, it is interesting to note that it is inferior to the TVP-VAR-SV models for the small and medium data sets. Clearly, addition of time-variation in the VAR coefficients helps improve forecasts in these cases. However, in the large data set, the evidence is mixed. In this case, the RMSEs produced by the TVP-VAR-SV are substantially better than those produced by the VAR-SV. However, the density forecasts

Figure 2: Heatmaps of posterior inclusion probabilities (PIPs) for the three focus variables in $t = 2017 : Q4$

PIPs for Constant Coefficients (β_{i0} and $u_{ij,0}$ in equation 11)



PIPs for Time-Varying Coefficients (v_{ij}^β and v_{ij}^u in equation 11)



are not. This could be due to the fact that there is typically a tradeoff between model dimension and parameter change. In small models, there is often a need for a high degree of parameter change to adequately fit patterns in the data and alleviate potential omitted variable biases. But in larger models, the information provided by the additional variables can fit these patterns, leaving less of a role for parameter change. Thus, in high dimensional cases the VAR-SV might be adequate and the extra flexibility provided by a TVP-VAR-SV may not be required. Of course, if the correct specification has a zero coefficient, the non-sparsified approach would try and estimate the time-varying coefficient to be constant over time. But, as illustrated in Figure 1, estimation uncertainty (although reduced) would still exist which could potentially hurt the forecasting performance of our approach. Sparsification as done in this paper clearly helps, but in the large data set there are still some cases where the VAR-SV is superior. In such cases, a simple extension of our shrink-then-sparsify approach could help. In this paper, we have focused on sparsifying α in equation (6). But any function of the parameters of a model could be sparsified in the same manner and, in particular, sparsifying the change in the states would be possible. This would lead to the constancy of a coefficient over certain periods in time while allowing for movements in other points in time when this kind of sparsification is applied.

The results in Table 3 highlight that, when the full hold-out period is considered, sparsification improves predictive accuracy relative to a non-sparsified model specification. The magnitude of these improvements, however, depends on model size. In the next step, we ask whether accuracy differences could also be specific to certain periods in time. To this end, Figure 3(a) shows the evolution of the log predictive Bayes factor between the sparsified and non-sparsified large-scale TVP-VAR-SV with the HS prior over the hold-out period.⁵ This Bayes factor is obtained by evaluating the one-step-ahead predictive density for the three focus variables jointly after integrating out the remaining variables. To investigate whether the gains in density forecasting performance stem from capturing higher order moments in the predictive distribution or from a more precise point forecast, Figure 3(b) shows cumulative squared one-step-ahead forecast errors averaged across the focus variables over time.

Figure 3(a) indicates that accuracy premia from sparsification tend to vary significantly over the business cycle. During expansionary stages, sparsification yields modest (in the case of the medium-sized model) to sustained (in the case of the large model) improvements in density forecasting performance relative to the non-sparsified competitor. For the small-scale TVP-VAR-SV, accuracy gains are more muted during expansionary periods. During recessions, in contrast, sparse models tend to be outperformed by their non-sparsified counterparts. Our conjecture is that this

⁵Comparable figures for other shrinkage priors reveal similar patterns. Thus, for the sake of brevity, we discuss the results for the HS prior exclusively.

Table 3: Relative root mean square errors (RMSEs) to a large BVAR-SV with a DL prior: 1997Q4 to 2017Q4

		Non-sparse					Sparse				
		DL	Lasso	NG	HS	MNIG	DL	Lasso	NG	HS	MNIG
SMALL (M = 3)											
	One-step-ahead										
	GDPC1	0.516 (0.209)	0.528 (0.188)	0.526 (0.194)	0.523 (0.2)	0.517 (0.223)	0.526 (0.228)	0.529 (0.231)	0.53 (0.227)	0.533 (0.222)	0.524 (0.229)
	GDPCTPI	0.86 (0.134)	0.868 (0.125)	0.867 (0.122)	0.861 (0.125)	0.87 (0.125)	0.883 (0.114)	0.915 (0.075)	0.912 (0.076)	0.903 (0.085)	0.89 (0.105)
	FEDFUNDS	0.837 (0.495)	0.858 (0.466)	0.843 (0.509)	0.83 (0.547)	0.832 (0.554)	0.814 (0.542)	0.827 (0.564)	0.826 (0.583)	0.818 (0.604)	0.827 (0.603)
	Four-step-ahead										
	GDPC1	0.512 (0.516)	0.512 (0.501)	0.51 (0.507)	0.507 (0.514)	0.533 (0.529)	0.547 (0.542)	0.543 (0.553)	0.545 (0.545)	0.548 (0.547)	0.545 (0.537)
	GDPCTPI	0.986 (0.091)	0.979 (0.099)	0.981 (0.098)	0.981 (0.099)	0.986 (0.088)	0.988 (0.081)	0.987 (0.072)	0.986 (0.071)	0.983 (0.073)	0.987 (0.079)
	FEDFUNDS	0.745 (0.457)	0.753 (0.431)	0.745 (0.454)	0.743 (0.495)	0.758 (0.488)	0.763 (0.466)	0.734 (0.468)	0.738 (0.477)	0.742 (0.505)	0.769 (0.499)
MEDIUM (M = 8)											
	One-step-ahead										
	GDPC1	0.659 (0.056)	0.647 (0.081)	0.651 (0.091)	0.619 (0.14)	0.534 (0.199)	0.491 (0.237)	0.497 (0.273)	0.494 (0.273)	0.487 (0.277)	0.469 (0.273)
	GDPCTPI	0.874 (0.106)	0.867 (0.126)	0.858 (0.134)	0.86 (0.137)	0.859 (0.147)	0.876 (0.132)	0.912 (0.064)	0.907 (0.074)	0.908 (0.089)	0.881 (0.124)
	FEDFUNDS	0.854 (0.273)	0.872 (0.327)	0.856 (0.408)	0.824 (0.552)	0.82 (0.464)	0.78 (0.401)	0.785 (0.546)	0.79 (0.581)	0.79 (0.646)	0.785 (0.593)
	Four-step-ahead										
	GDPC1	0.616 (0.22)	0.63 (0.272)	0.62 (0.297)	0.587 (0.377)	0.535 (0.438)	0.553 (0.524)	0.549 (0.545)	0.548 (0.547)	0.552 (0.533)	0.551 (0.513)
	GDPCTPI	0.985 (-0.001)	0.98 (0.058)	0.979 (0.068)	0.98 (0.096)	0.984 (0.086)	0.995 (0.07)	0.983 (0.025)	0.987 (0.033)	0.983 (0.054)	0.985 (0.067)
	FEDFUNDS	0.831 (0.262)	0.837 (0.297)	0.814 (0.372)	0.747 (0.46)	0.742 (0.372)	0.771 (0.417)	0.738 (0.443)	0.734 (0.46)	0.735 (0.448)	0.75 (0.368)
LARGE (M = 20)											
	One-step-ahead										
	GDPC1	1.013 (-0.351)	0.851 (-0.32)	0.875 (-0.196)	0.816 (-0.002)	0.676 (0.068)	0.662 (-0.067)	0.573 (-0.004)	0.551 (0.096)	0.531 (0.207)	0.513 (0.209)
	GDPCTPI	1.123 (-0.054)	1.157 (-0.016)	1.17 (-0.015)	1.031 (0.034)	0.909 (0.064)	1.011 (0.018)	1.063 (0.011)	1.133 (-0.024)	0.941 (0.014)	0.912 (0.033)
	FEDFUNDS	1.264 (-0.55)	1.655 (-0.692)	1.442 (-0.473)	1.055 (0.229)	0.865 (0.17)	1.036 (-0.313)	1.561 (-0.556)	1.418 (-0.295)	0.891 (0.559)	0.783 (0.559)
	Four-step-ahead										
	GDPC1	0.893 (-1.185)	0.73 (-1.039)	0.627 (-0.598)	0.513 (0.191)	0.529 (0.361)	0.654 (-0.332)	0.557 (-0.167)	0.556 (0.174)	0.545 (0.57)	0.567 (0.534)
	GDPCTPI	1.063 (-0.81)	1.116 (-0.652)	1.136 (-0.331)	1.022 (0.003)	0.997 (0.055)	1.018 (-0.283)	1.034 (-0.261)	1.163 (-0.123)	0.994 (0.003)	0.987 (0.038)
	FEDFUNDS	0.937 (-1.452)	1.229 (-1.489)	1.269 (-1.037)	0.812 (0.082)	0.856 (0.089)	0.86 (-0.723)	0.844 (-0.916)	1.001 (-0.488)	0.815 (0.267)	0.835 (0.228)


Notes: Numbers in parentheses refer to the average log predictive likelihoods (LPLs) vis-à-vis the BVAR-SV with a DL prior. DL refers to a TVP-VAR-SV with a Dirichlet-Laplace prior, Lasso to the Bayesian Lasso, NG to the Normal-Gamma prior, HS to the Horseshoe, and NMIG to the Normal-mixture of Inverse Gamma prior.

stems from the fact that during turbulent times, the sparsified predictive distributions feature a smaller variance, making it harder to capture outlying observations and thus translating in lower log predictive likelihoods.


Our conjecture is confirmed when focusing on point forecasts. In terms of point predictions, we observe that forecast errors are almost identical in the period up to the global financial crisis. During the recession in 2008/2009, forecast errors increase markedly but slightly less so for the sparsified model and for the medium and large dataset. This suggests that the drop in the log predictive Bayes factor is mainly driven by higher order moments, implying that while the accuracy of the point prediction increases, adverse movements in the corresponding predictive variance offset this gain.

Figure 3: Performance differences between a sparsified and non-sparsified TVP-VAR-SV with a HS prior

(a) Evolution of log predictive Bayes factor (sparse versus non-sparse)



(b) Evolution of cumulative squared forecast errors



Notes: The log-predictive Bayes factor between the sparsified and non-sparsified model is obtained by considering the joint one-step-ahead predictive density for the three focus variables and the squared forecast errors are averages across the one-step-ahead forecast errors for the focus variables. The black line in panel (b) refers to the sparsified squared forecast error while the red line denotes the non-sparsified model. The gray shaded areas refer to NBER reference recessions in the US.

6 Conclusions

Global-local shrinkage priors have enjoyed great popularity in over-parameterized regressions and VARs involving large numbers of variables. And, increasingly, they have been used with TVP versions of these models which are potentially even more over-parameterized. Use of such priors can potentially reduce estimation error and improve

forecasts. However, estimation error is not completely eliminated and it is possible that further improvements in forecasting performance can be achieved by adding an additional sparsification step to shrunk estimates. In this paper, we have developed methods for doing so. In an artificial data exercise, we have shown that our shrink-then-sparsify approach to TVP regression leads to more accurate estimates for a variety of DGPs. Particularly large gains are found in sparse DGPs. In a macroeconomic forecasting exercise, adding sparsification to shrinkage also leads to substantial improvements in forecast performance.

References

- Belmonte M, Koop G and Korobilis D (2014) Hierarchical shrinkage in time-varying coefficient models. *Journal of Forecasting* 33(1), 80–94
- Bhattacharya A, Pati D, Pillai NS and Dunson DB (2015) Dirichlet–Laplace priors for optimal shrinkage. *Journal of the American Statistical Association* 110(512), 1479–1490
- Bitto A and Frühwirth-Schnatter S (2019) Achieving shrinkage in a time-varying parameter model framework. *Journal of Econometrics* 210(1), 75–97
- Carriero A, Clark T and Marcellino M (2016) Large Vector Autoregressions with stochastic volatility and flexible priors. *Federal Reserve Bank of Cleveland Working Paper, no. 16-17*
- Carter C and Kohn R (1994) On Gibbs sampling for state space models. *Biometrika* 81(3), 541–553
- Carvalho CM, Polson NG and Scott JG (2010) The horseshoe estimator for sparse signals. *Biometrika* 97(2), 465–480
- Cogley T and Sargent TJ (2005) Drifts and volatilities: monetary policies and outcomes in the post WWII US. *Review of Economic Dynamics* 8(2), 262 – 302
- Cross J, Hou C and Poon A (2019) Macroeconomic forecasting with large Bayesian VARs: Global-local priors and the illusion of sparsity. *Manuscript*
- D’Agostino A, Gambetti L and Giannone D (2013) Macroeconomic forecasting and structural change. *Journal of Applied Econometrics* 28(1), 82–101
- Eisenstat E, Chan J and Strachan R (2019) Reducing the state space dimension in a large TVP-VAR
- Friedman J, Hastie T, Höfling H, Tibshirani R et al. (2007) Pathwise coordinate optimization. *The Annals of Applied Statistics* 1(2), 302–332
- Frühwirth-Schnatter S (1994) Data augmentation and dynamic linear models. *Journal of Time Series Analysis* 15(2), 183–202
- Frühwirth-Schnatter S and Wagner H (2010) Stochastic model specification search for Gaussian and partial non-Gaussian state space models. *Journal of Econometrics* 154(1), 85–100

- George EI and McCulloch RE (1993) Variable selection via Gibbs sampling. *Journal of the American Statistical Association* 88(423), 881–889
- George EI and McCulloch RE (1997) Approaches for Bayesian variable selection. *Statistica sinica* , 339–373
- George EI, Sun D and Ni S (2008) Bayesian stochastic search for VAR model restrictions. *Journal of Econometrics* 142(1), 553–580
- Giannone D, Lenza M and Primiceri GE (2017) Economic predictions with big data: The illusion of sparsity
- Griffin J and Brown P (2010) Inference with normal-gamma prior distributions in regression problems. *Bayesian Analysis* 5(1), 171–188
- Hahn PR and Carvalho CM (2015) Decoupling Shrinkage and Selection in Bayesian Linear Models: A Posterior Summary Perspective. *Journal of the American Statistical Association* 110(509), 435–448
- Huber F and Feldkircher M (2019) Adaptive shrinkage in Bayesian vector autoregressive models. *Journal of Business & Economic Statistics* 37(1), 27–39
- Huber F, Kastner G and Feldkircher M (2019) Should I stay or should I go? A latent threshold approach to large-scale mixture innovation models. *Journal of Applied Econometrics* 0(0)
- Ishwaran H and Rao JS (2005) Spike and slab variable selection: frequentist and Bayesian strategies. *The Annals of Statistics* 33(2), 730–773
- Kalli M and Griffin J (2014) Time-varying sparsity in dynamic regression models. *Journal of Econometrics* 178(2), 779 – 793
- Kastner G (2016) Dealing with stochastic volatility in time series using the R package stochvol. *Journal of Statistical software* 69(5), 1–30
- Kastner G and Frühwirth-Schnatter S (2014) Ancillarity-sufficiency interweaving strategy (ASIS) for boosting MCMC estimation of stochastic volatility models. *Computational Statistics & Data Analysis* 76, 408–423
- Kastner G and Huber F (2017) Sparse Bayesian vector autoregressions in huge dimensions. *manuscript*
- Koop G and Korobilis D (2018) Variational Bayes inference in high dimensional time-varying parameter models. *manuscript*

- Koop G, Korobilis D and Pettenuzzo D (2019) Bayesian compressed Vector Autoregressions. *Journal of Econometrics* 210(1), 135–154
- Kowal D, Matteson D and Ruppert D (2017) Dynamic shrinkage processes. *arXiv:1707.00763*
- Makalic E and Schmidt DF (2016) A simple sampler for the horseshoe estimator. *IEEE Signal Processing Letters* 23(1), 179–182
- Malsiner-Walli G and Wagner H (2011) Comparing Spike and Slab Priors for Bayesian Variable Selection. *Austrian Journal of Statistics* 40(4), 241–264
- McCracken MW and Ng S (2016) FRED-MD: A monthly database for macroeconomic research. *Journal of Business & Economic Statistics* 34(4), 574–589
- Park T and Casella G (2008) The Bayesian Lasso. *Journal of the American Statistical Association* 103(482), 681–686
- Polson NG and Scott JG (2010) Shrink globally, act locally: Sparse Bayesian regularization and prediction. *Bayesian statistics* 9, 501–538
- Primiceri G (2005) Time varying structural autoregressions and monetary policy. *Oxford University Press* 72(3), 821–852
- Ray P and Bhattacharya A (2018) Signal Adaptive Variable Selector for the Horseshoe Prior. *arXiv preprint arXiv:1810.09004*
- Rockova V and McAlinn K (2017) Dynamic variable selection with spike-and-slab process priors. *Manuscript, available at <http://hedibert.org/wp-content/uploads/2018/06/uribe-lobes-Sep2017.pdf>*
- Uribe P and Lopes H (2017) Dynamic sparsity on dynamic regression models. *Manuscript, available at <http://hedibert.org/wp-content/uploads/2018/06/uribe-lobes-Sep2017.pdf>*

Appendix A Global-local Shrinkage Priors

The first four sub-sections of this appendix provide relevant details on the prior setup, briefly discussing key features of the used priors, hyperparameter choices used, and relevant information necessary to perform posterior inference.

A.1 The Dirichlet-Laplace Prior

The DL prior, originally proposed in Bhattacharya et al. (2015), assumes that each element in $\boldsymbol{\alpha}$, α_j ($j = 1, \dots, 2K$), follows a Gaussian distribution,

$$\alpha_j | \omega_j, \xi_j, \lambda \sim \mathcal{N}(0, \omega_j \xi_j^2 \zeta^2),$$

with

$$\omega_j \sim \mathcal{E}(1/2), \quad \xi_j \sim \mathcal{D}(a, \dots, a), \quad \zeta \sim \mathcal{G}(2Ka, 1/2)$$

where ω_j is a variable-specific scaling parameter that features an exponentially distributed prior, with \mathcal{E} denoting the exponential distribution, ξ_j denotes yet another local shrinkage parameter with $\boldsymbol{\xi} = (\xi_1, \dots, \xi_{2K})'$ being bounded to the $(2K - 1)$ -dimensional simplex (i.e. $\xi_j \geq 0$ and $\sum_j \xi_j = 1$). We use a Dirichlet distributed prior with intensity parameter a on ξ_j . Finally, ζ is a global shrinkage term that follows a Gamma distribution. Notice that the relationship between this prior hierarchy and the general form provided in equation (2) can be seen by defining $\phi_j = \omega_j \xi_j^2$ and $\lambda = \zeta^2$.

Bhattacharya et al. (2015) show within the stylized normal means problem that the optimal value of a is specified to be $(2K)^{-(1+\varphi)}$ with φ being a positive number close to zero. Since this hyperparameter plays a crucial role in determining the shrinkage behavior of the DL prior, we estimate it using a prior which is a uniform distribution that is bounded between $(2K)^{-1}$ and $1/2$.

Posterior simulation can be carried out using a slightly modified variant of the MCMC algorithm proposed in Bhattacharya et al. (2015). The full conditional posterior distribution of ω_j follows an inverse Gaussian distribution:

$$\omega_j | \alpha_j, \xi_j, \zeta \sim i\mathcal{G}\left(\zeta \frac{\xi_j}{|\alpha_j|}, 1\right).$$

The global shrinkage parameter ζ follows a generalized inverted Gaussian (GIG) distribution,

$$\zeta|\boldsymbol{\alpha}, \boldsymbol{\xi} \sim \mathcal{GIG}\left(2K(a-1), 1, 2\sum_{j=1}^{2K}\frac{|\alpha_j|}{\xi_j}\right).$$

Moreover, we draw the second set of local scaling parameters ξ_j by introducing auxiliary variables T_j that follow a GIG distribution:

$$T_j|a, \alpha_j \sim \mathcal{GIG}(a-1, 1, 2|\alpha_j|).$$

We then set $\xi_j = T_j / \sum_{i=1}^{2K} T_i$ to obtain a valid draw from the full conditional posterior of ξ_j .

To simulate from the conditional posterior of a , we employ a Metropolis Hastings algorithm with a Gaussian proposal distribution truncated between $(2K)^{-1}$ and $1/2$. The variance of the proposal distribution is tuned during the first 20 percent of the burn-in stage of the MCMC sampler such that the acceptance rate is between 20 and 40 percent.

A.2 The Normal-Gamma Prior and the Lasso

As compared to the DL prior, the NG prior proposed in Griffin and Brown (2010), consists of a single group of idiosyncratic scaling factors ϕ_j and a global shrinkage parameter $\lambda = 1/\tilde{\lambda}$. We assume that each α_j follows a zero mean Gaussian distribution a priori:

$$\alpha_j|\phi_j, \tilde{\lambda} \sim \mathcal{N}(0, \phi_j), \quad \phi_j|\tilde{\lambda} \sim \mathcal{G}(\vartheta, \vartheta\tilde{\lambda}/2), \quad \tilde{\lambda} \sim \mathcal{G}(d_{\tilde{\lambda}}, e_{\tilde{\lambda}}).$$

Here, we let ϑ denote a hyperparameter that controls the tail behavior of the prior, with smaller values of ϑ leading to heavier tails and increasing mass is placed on zero while larger value do the opposite. $d_{\tilde{\lambda}}$ and $e_{\tilde{\lambda}}$ are hyperparameters that control the overall degree of shrinkage, with values close to zero implying heavy shrinkage towards zero.

One key feature of the NG prior is that it nests the Bayesian Lasso of Park and Casella (2008) by setting $\vartheta = 1$. Since ϑ plays a crucial role, we follow Griffin and Brown (2010) and Huber and Feldkircher (2019) and introduce an Exponential prior on ϑ :

$$\vartheta \sim \mathcal{Exp}(\underline{\vartheta}).$$

$\underline{\vartheta}$ is set equal to 1, pushing the prior towards the Bayesian Lasso. Moreover, we set $d_{\tilde{\lambda}} = e_{\tilde{\lambda}} = 10^{-4}$, implying a disperse prior on $\tilde{\lambda}$ and thus being consistent with heavy shrinkage (by allowing large values of $\tilde{\lambda}$).

The hierarchical structure of the prior yields closed-form full conditionals for ϕ_j and $\tilde{\lambda}$. The local scaling parameters ϕ_j follow a GIG distribution:

$$\phi_j | \tilde{\lambda}, \alpha_j \sim \mathcal{G}\mathcal{F}\mathcal{G} \left(\vartheta - \frac{1}{2}, \vartheta \tilde{\lambda}, \alpha_j^2 \right).$$

For the global shrinkage parameter, we obtain a Gamma-distributed full conditional posterior distribution:

$$\tilde{\lambda} | \phi_1, \dots, \phi_{2K}, \vartheta \sim \mathcal{G} \left(d_{\tilde{\lambda}} + \vartheta 2K, e_{\tilde{\lambda}} + \frac{\vartheta}{2} \sum_{i=1}^{2K} \phi_i \right).$$

Finally, we obtain draws from the conditional posterior of ϑ by setting up a random walk MH algorithm in terms of $\log \vartheta$ (see Griffin and Brown, 2010).

A.3 The Horseshoe Prior

For the HS prior of Carvalho et al. (2010), we consider the representation based on auxiliary variables proposed in Makalic and Schmidt (2016). The corresponding prior hierarchy is given by

$$\alpha_j | \lambda, \phi_j \sim \mathcal{N}(0, \phi_j \lambda), \quad \phi_j \sim \mathcal{G}^{-1}(1/2, 1/\nu_j), \quad \lambda \sim \mathcal{G}^{-1}(1/2, 1/\varphi),$$

whereby ν_j and φ denote auxiliary variables and \mathcal{G}^{-1} denotes the inverse Gamma distribution. The auxiliary variables also follow inverse Gamma distributions,

$$\nu_1, \dots, \nu_{2K}, \varphi \sim \mathcal{G}^{-1}(1/2, 1).$$

This representation of the HS prior allows for straightforward updating of the local and global scaling parameters and involves sampling from inverted Gamma distributions exclusively. The corresponding full conditional posterior distributions are

$$\begin{aligned} \phi_j | \alpha_j, \lambda, \nu_j &\sim \mathcal{G}^{-1} \left(1, \frac{1}{\nu_j} + \frac{\alpha_j^2}{2\lambda} \right) \\ \lambda | \alpha_j, \phi_j, \varphi &\sim \mathcal{G}^{-1} \left(\frac{2K+1}{2}, \frac{1}{\varphi} + \frac{1}{2} \sum_{i=1}^{2K} \frac{\alpha_i^2}{\phi_j} \right). \end{aligned}$$

The conditional posteriors of the auxiliary variables are given by

$$\begin{aligned}\nu_j|\phi_j &\sim \mathcal{G}^{-1}\left(1, 1 + \frac{1}{\phi_j^2}\right), \\ \varphi|\lambda &\sim \mathcal{G}^{-1}\left(1, 1 + \frac{1}{\lambda}\right).\end{aligned}$$

A.4 The Normal-mixture of Inverse Gamma Prior

The NMIG prior of Ishwaran and Rao (2005) extends the original SSVS prior proposed in George and McCulloch (1993, 1997) along several dimensions. To set the stage, we use a mixture of Gaussians prior distribution on α_j :

$$\alpha_j|\delta_j, \tau_j^2 \sim \mathcal{N}(0, \tau_j^2)\delta_j + \mathcal{N}(0, c\tau_j^2)(1 - \delta_j),$$

where δ_j denotes a Bernoulli random variable with prior probability $Prob(\delta_j = 1) = \underline{p}$ while c is a constant close to zero and τ_j^2 is a coefficient-specific scaling factor. Following Ishwaran and Rao (2005), we specify an inverse Gamma prior on τ_j^2 and a Beta distributed prior on \underline{p} :

$$\begin{aligned}\tau_j^2 &\sim \mathcal{G}^{-1}(d_\tau, e_\tau), \\ \underline{p} &\sim \mathcal{B}(d_p, e_p),\end{aligned}$$

with d_τ, e_τ, d_p and e_p denoting hyperparameters. Notice that this specification implies conditional prior independence between the indicators δ_j . However, the common prior inclusion probability \underline{p} serves as a common factor, implying that marginally, the indicators are dependent.

Ishwaran and Rao (2005) notice that after integrating out τ_j^2 and \underline{p} , the two components in the prior follow t -distributions. The hyperparameter d_τ controls the degrees of freedom of the marginal prior while the variances are given by ce_τ/d_τ (for the spike component) and e_τ/d_τ (for the slab component). In the empirical applications, we set $e_p = d_p = 1$, implying a Uniform prior on \underline{p} and $c = 2.5/10^5$. Moreover, we set $d_\tau = 5$, leading to 10 degrees of freedom and $e_\tau = 4$. This is the benchmark prior setup as specified in Malsiner-Walli and Wagner (2011).

For this prior specification, all conditional posterior distributions are available in closed form. The full-conditional posterior of δ_j follows a Bernoulli distribution with posterior probability \bar{p}_j given by:

$$\bar{p}_j = \text{Prob}(\delta_j = 1 | \alpha_j, \tau_j^2, \underline{p}) = \frac{\frac{1}{\tau_j^2} \exp\left(-\frac{1}{2} \frac{\alpha_j^2}{\tau_j^2}\right)}{\frac{1}{\tau_j^2} \exp\left(-\frac{1}{2} \frac{\alpha_j^2}{\tau_j^2}\right) \underline{p} + \frac{1}{c\tau_j^2} \exp\left(-\frac{1}{2} \frac{\alpha_j^2}{c\tau_j^2}\right) (1 - \underline{p})}.$$

The scaling factors τ_j^2 follow an inverted Gamma distribution

$$\tau_j^2 | \alpha_j^2, \delta_j \sim \mathcal{G}^{-1}\left(d_\tau + \frac{1}{2}, e_\tau + \frac{\alpha_j^2}{\delta_j + (1 - \delta_j)c}\right).$$

Finally, the posterior distribution of \underline{p} is a Beta distribution:

$$\underline{p} | \delta_1, \dots, \delta_{2K} \sim \mathcal{B}\left(d_p + \sum_{j=1}^{2K} \delta_j, e_p + 2K - \sum_{j=1}^{2K} \delta_j\right).$$

Appendix B Full Conditional Posterior Simulation

For the dynamic regression models used in the main body of the text we use a relatively standard MCMC algorithm. Since we estimate the TVP-VAR-SV on an equation-by-equation basis, we describe the MCMC algorithm for the TVP regression model only. However, it is worth noting that all priors described in the previous subsection are specified to be equation-specific. This implies that instead of having a single global shrinkage parameter λ , each equation features its own global (equation-specific) shrinkage parameter. Moreover, one additional difference is that the dynamic regression model in Section 3 features homoscedastic errors. In the TVP-VAR case, we allow for stochastic volatility, implying that the MCMC algorithm differs slightly.

Our posterior simulator cycles between the following steps:

1. Simulate the full-history of $\tilde{\beta}_t$, conditional on the remaining parameters, using the forward-filtering backward-sampling algorithm proposed in Carter and Kohn (1994) and Frühwirth-Schnatter (1994) while exploiting the non-centered parameterization.
2. Sample the error variances from an inverted Gamma full conditional posterior distribution:

$$\sigma_\varepsilon^2 | \bullet \sim \mathcal{G}^{-1}\left(d_\sigma + T/2, e_\sigma + \frac{1}{2} \sum_{t=1}^T (y_t - \alpha' \mathbf{Z}_t)^2\right).$$

where the \bullet indicates conditioning on all parameters and the data.

3. Conditional on $\{\tilde{\beta}_t\}_{t=1}^T$ and σ_ε^2 , the conditional posterior of α takes a multivariate Gaussian form:

$$\alpha|\bullet \sim \mathcal{N}(\bar{\alpha}, \bar{\Omega}),$$

with

$$\begin{aligned}\bar{\Omega} &= (\sigma_\varepsilon^{-2} \mathbf{Z}' \mathbf{Z} + \underline{\Omega}^{-1})^{-1}, \\ \bar{\alpha} &= \bar{\Omega} (\sigma_\varepsilon^{-2} \mathbf{Z}' \mathbf{y})\end{aligned}$$

where \mathbf{Z} is a $T \times 2K$ matrix with the t^{th} row equal to \mathbf{Z}'_t . Likewise, $\mathbf{y} = (y_1, \dots, y_T)'$ is a T -dimensional vector. $\underline{\Omega}$ denotes a diagonal prior variance-covariance matrix with typical element depending on the specific shrinkage prior chosen.

4. Depending on the global-local shrinkage prior adopted, construct the matrix $\underline{\Omega}$ based on the conditional posterior distributions outlined in Appendix A.

In case we use a stochastic volatility specification for the error variances, we use the algorithm proposed in Kastner and Frühwirth-Schnatter (2014) and implemented in the R package `stochvol` (Kastner, 2016). For the VAR case, the main steps of this algorithm remain identical except that the different steps of the algorithm can be interpreted as being specific to a given equation of the model. In all applications, we repeat this algorithm 30,000 times and discard the first 15,000 draws as burn-in.

Appendix C Data Appendix

FRED Mnemonic	Description	Transformation Codes	SMALL	MEDIUM	LARGE
GDPC1	Real Gross Domestic Product	5	x	x	x
PCECC96	Real Personal Consumption Expenditures	5	x	x	x
FPIx	Real private fixed investment	5		x	x
GCFC1	Real Government Consumption Expenditures and Gross Investment	5			x
INDPRO	IP:Total index Industrial Production Index (Index 2012=100)	5			x
CE16OV	Civilian Employment (Thousands of Persons)	5		x	x
UNRATE	Civilian Unemployment Rate (Percent)	2			x
CES060000007	Average Weekly Hours of Production and Nonsupervisory Employees: Goods-Producing	2		x	x
HOUST	Housing Starts: Total: New Privately Owned Housing Units Started	5			x
PERMIT	New Private Housing Units Authorized by Building Permits	5			x
PCECTPI	Personal Consumption Expenditures: Chain-type Price Index	6			x
GDPCTPI	Gross Domestic Product: Chain-type Price Index	6	x	x	x
CPIAUCSL	Consumer Price Index for All Urban Consumers: All Items	6			x
CES0600000008	Average Hourly Earnings of Production and Nonsupervisory Employees:	6		x	x
FEDFUNDS	Effective Federal Funds Rate (Percent)	2	x	x	x
GS1	1-Year Treasury Constant Maturity Rate (Percent)	2			x
GS10	10-Year Treasury Constant Maturity Rate (Percent)	2			x
TOTRESNS	Total Reserves of Depository Institutions	6			x
NONBORRES	Reserves Of Depository Institutions, Nonborrowed	7			x
S.P.500	S&Ps Common Stock Price Index: Composite	5			x

Notes: Transformation codes for a series y : (1) no transformation, (2) first differences δy_t , (3) double differences $\Delta^2 y_t$ (4) logarithmic transform $\log y_t$, (5) difference of the logarithm $\Delta \log y_t$, (6) double difference of the logarithm $\Delta^2 \log y_t$, and (7) $\Delta(y_t/y_{t-1} - 1)$.

Table 4: Data description