CAMA

Centre for Applied Macroeconomic Analysis

Stochastic Model Specification Search for Time-Varying Parameter VARs

CAMA Working Paper 23/2014
March 2014

Eric Eisenstat

University of Bucharest and RIMIR

Joshua C.C. Chan

Research School of Economics, ANU and
Centre for Applied Macroeconomic Analysis, ANU

Rodney W. Strachan

School of Economics, The University of Queensland and
Centre for Applied Macroeconomic Analysis, ANU

Abstract

This article develops a new econometric methodology for performing stochastic model specification search (SMSS) in the vast model space of time-varying parameter VARs with stochastic volatility and correlated state transitions. This is motivated by the concern of over-fitting and the typically imprecise inference in these highly parameterized models. For each VAR coefficient, this new method automatically decides whether it is constant or time-varying. Moreover, it can be used to shrink an otherwise unrestricted time-varying parameter VAR to a stationary VAR, thus providing an easy way to (probabilistically) impose stationarity in time-varying parameter models. We demonstrate the effectiveness of the approach with a topical application, where we investigate the dynamic effects of structural shocks in government spending on U.S. taxes and GDP during a period of very low interest rates.

Keywords

Bayesian Lasso, shrinkage, fiscal policy

JEL Classification

C11, C52, E37, E47

Address for correspondence:

(E) cama.admin@anu.edu.au

[The Centre for Applied Macroeconomic Analysis](#) in the Crawford School of Public Policy has been established to build strong links between professional macroeconomists. It provides a forum for quality macroeconomic research and discussion of policy issues between academia, government and the private sector.

The Crawford School of Public Policy is the Australian National University's public policy school, serving and influencing Australia, Asia and the Pacific through advanced policy research, graduate and executive education, and policy impact.

Stochastic Model Specification Search for Time-Varying Parameter VARs

Eric Eisenstat

Faculty of Business Administration,
University of Bucharest and RIMIR

Joshua C.C. Chan

Research School of Economics and
Centre for Applied Macroeconomic Analysis
Australian National University

Rodney W. Strachan

School of Economics,
The University of Queensland
Centre for Applied Macroeconomic Analysis

February 2014

Abstract

This article develops a new econometric methodology for performing stochastic model specification search (SMSS) in the vast model space of time-varying parameter VARs with stochastic volatility and correlated state transitions. This is motivated by the concern of over-fitting and the typically imprecise inference in these highly parameterized models. For each VAR coefficient, this new method automatically decides whether it is constant or time-varying. Moreover, it can be used to shrink an otherwise unrestricted time-varying parameter VAR to a stationary VAR, thus providing an easy way to (probabilistically) impose stationarity in time-varying parameter models. We demonstrate the effectiveness of the approach with a topical application, where we investigate the dynamic effects of structural shocks in government spending on U.S. taxes and GDP during a period of very low interest rates.

Keywords: Bayesian Lasso, shrinkage, fiscal policy

JEL Classification: C11, C52, E37, E47

1 Introduction

Vector autoregressions (VARs) are widely used for modeling and forecasting in macroeconomics. In particular, VARs have been used to understand the interactions between macroeconomic variables, often through the estimation of impulse response functions that characterize the effects of a variety of structural shocks on key economic variables. In recent years there has been much interest in extending the traditional constant coefficient VARs to time-varying parameter VARs (TVP-VARs) where the VAR coefficients are allowed to gradually evolve over time (see, among many others, Cogley and Sargent, 2005; Cogley, Primiceri, and Sargent, 2010; Koop, Leon-Gonzalez, and Strachan, 2011; Koop and Korobilis, 2013). This approach is motivated by a growing body of empirical evidence that demonstrates the importance of accommodating time-varying structures for model fitting and forecasting.

However, even for a moderate size VAR, the number of parameters in the model can be quite large relative to the number of observations. This can lead to imprecise estimation of impulse response functions and poor forecast performance. This problem is exacerbated in TVP-VARs, which are much higher dimensional than constant coefficient VARs. The potential problems associated with parameter proliferation have led many researchers to use Bayesian shrinkage in VARs to improve estimates precision and forecast performance (e.g., Banbura, Giannone, and Reichlin, 2010; Carriero, Clark, and Marcellino, 2011; Koop, 2011; Korobilis, 2011). Applying shrinkage to time-varying parameter models is less straightforward and it often requires computationally demanding algorithms (Chan, Koop, Leon-Gonzalez, and Strachan, 2012; Nakajima and West, 2013) or approximate inference (Koop and Korobilis, 2012). A related issue is model specification and model selection; although empirically a TVP-VAR that allows *all* VAR coefficients to change over time typically perform better than a constant coefficient VAR, it is plausible and even likely that a TVP-VAR where only *some* coefficients are time-varying while others are time-invariant will perform better than both alternatives.¹ Reducing the number of time-varying parameters thereby achieves a degree of parsimony. However, it is unclear how one can decide a priori which coefficients are fixed and which are time-varying.

¹A referee has suggested another direction for extending such models that allows coefficients to vary some of the time but not necessarily all of the time. This approach would provide useful new information on the nature of time variation and could follow from the approach in Chan, Koop, Léon-Gonzalez and Strachan (2012). However, we do not attempt this extension in this paper.

The main goal of this paper, following the interesting work of Frühwirth-Schnatter and Wagner (2010) and Belmonte, Koop, and Korobolis (2014) (hereafter FSW and BKK, respectively), among others, is to develop a new methodology to nest both time-varying and time-invariant VARs that is applicable to high-dimensional settings. Our starting point is the stochastic model specification search (SMSS) framework introduced in FSW. Specifically, for each VAR coefficient, we introduce an indicator that chooses between a time-varying against a time-invariant parameter. This allows the model to automatically switch to a more parsimonious specification when the time-variation of the coefficient is “small”. One computational challenge in this setup is that in a typical VAR, the number of values the indicators can take can be very large — e.g., for a small VAR with 3 variables and 4 lags, the number of combinations is $2^{39} \approx 5.5 \times 10^{11}$. To circumvent this computation problem, we introduce a new hierarchical prior for the indicators under which the stochastic model specification search can be performed efficiently. In addition, this new approach also incorporates the hierarchical Lasso prior in BKK that provides additional shrinkage, while maintaining the natural and intuitive framework in FSW. The proposed approach therefore adds to the growing literature on efficient methods of ensuring parsimony in potentially over-parameterized TVP-VARs.

Another advantage of the proposed approach is that it can be used to shrink the TVP-VAR to a stationary, constant coefficient VAR. In macroeconomic applications it is often necessary to impose stationarity conditions to avoid explosive impulse-response functions or forecasts. However, imposing stationarity conditions in a TVP-VAR implies inequality constraints on the time-varying coefficients, which results in a nonlinear state space model where conventional Kalman filter-based algorithms cannot be used. As pointed out in Koop and Potter (2011), the common way to impose stationarity in such settings — estimating the unconstrained TVP-VAR using Kalman filter-based algorithms and discarding any draws that do not satisfy the stationarity conditions — may lead to invalid inference. On the other hand, the single- and multi-move samplers proposed in Koop and Potter (2011), though sampling from the correct posterior distribution, can be computationally demanding. The proposed approach therefore provides a computationally feasible alternative for imposing the stationarity conditions probabilistically. By shrinking the TVP-VAR towards a stationary VAR, the model has the features of a stationary model, while still allowing some weight on nonstationarity to capture (poorly modeled) nonlinearities or extreme events.

In the second contribution of the paper we demonstrate the overall approach with a topical application. In particular, we investigate the dynamic responses of output growth

and government revenues and expenditure to a government spending shock. We generalize earlier work by Blanchard and Perotti (2002) by allowing for general time variation in the model parameters. We study the dynamic effects of structural shocks in government spending on U.S. taxes and GDP during a period where the interest rate is close to zero. In doing so we compare the inference obtained from the SMSS and the standard TVP-VAR specifications. We find that the SMSS generally offers more accurate estimates. In particular, our results from the SMSS show clearer evidence of an implication of Ricardian equivalence that the evolution of taxes generally follows that of spending. This phenomenon is much more difficult to detect with the standard TVP-VAR alone where the impulse responses are much less precisely estimated. In contrast, the SMSS leads to a distinct, clear pattern of long run government spending and taxes changing by the same amount during the period under examination. Overall, it appears that the SMSS is able to efficiently allocate time variation among the model parameters such as to strike a balance between parsimony and flexibility: most parameters are restricted to be essentially constant, while a select few are allowed to have a large degree of time variation.

The rest of the article is organized as follows. In Section 2 we first extend the stochastic model specification search approach of FSW to a multivariate setting. We then introduce a new hierarchical prior on the indicators and highlight its advantages over competing approaches. Section 3 outlines the posterior computation. We then discuss in Section 4 two important extensions — allowing for stochastic volatility and a full error covariance matrix for state transition. The hierarchical prior we employ is particularly important in implementing the latter as neither the FSW nor the BKK specifications would suffice for this purpose. Section 5 presents the application to responses of macro variables to fiscal policy shocks and Section 6 contains some final comments.

2 SMSS and State Space Models

A popular approach for allowing for time-varying coefficients in time series models is through the state space specification. Specifically, suppose $\mathbf{y}_t$ is an $n \times 1$ vector of observations on the dependent variables, $\mathbf{X}_t$ is an $n \times m$ matrix of observations on explanatory variables and $\boldsymbol{\beta}_t$ is an $m \times 1$ vector of states. Then a generic state space

model can be written as:

$$\mathbf{y}_t = \mathbf{X}_t \boldsymbol{\beta}_t + \boldsymbol{\varepsilon}_t, \quad (1)$$

$$\boldsymbol{\beta}_t = \boldsymbol{\beta}_{t-1} + \boldsymbol{\eta}_t, \quad (2)$$

where $\boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$, $\boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Omega})$, and $\boldsymbol{\Omega}$ is typically assumed to be a diagonal matrix $\boldsymbol{\Omega} = \text{diag}(\omega_1^2, \dots, \omega_m^2)$. In Section 4 we generalize this state space model to include a full covariance matrix in the state equation and stochastic volatility in the measurement equation. The errors $\boldsymbol{\varepsilon}_t$ and $\boldsymbol{\eta}_t$ are assumed to be independent at all leads and lags. Finally, the state equation (2) is initialized with $\boldsymbol{\beta}_0 = \boldsymbol{\alpha}$, where $\boldsymbol{\alpha} \sim \mathcal{N}(\boldsymbol{\alpha}_0, \mathbf{A}_0^{-1})$. In a TVP-VAR setting, we have $\mathbf{X}_t = \mathbf{I}_n \otimes \mathbf{x}'_t$, where $\mathbf{I}_n$ is the $n \times n$ identity matrix and $\mathbf{x}_t$ is a $k \times 1$ vector of deterministic terms and lagged observations (hence $m = nk$).

This general state space framework encompasses a wide variety of commonly used time-varying parameter (TVP) regression models in macroeconomics and has become a standard framework for analyzing macroeconomic data. However, recent research has raised the concern that over-fitting might be a problem for these highly parameterized models. Moreover, these high-dimensional models typically give imprecise estimates, making any form of inference more difficult. Motivated by these concerns, researchers might wish to have a more parsimonious specification that reduces the potential problem of over-parameterization, while maintaining the flexibility of the state space framework that allows time-variation in the coefficients. For example, one might wish to have a default model with time-invariant coefficients, but where each of these coefficients can switch to being time-varying when there is strong evidence for time-variation. In this way, one can maintain a parsimonious specification that leads to more precise estimates, while minimizing the risk of model misspecification.

In what follows, we first outline two existing methods for performing a stochastic model specification search with the aim of shrinking the model towards a more parsimonious specification. We then present a new method to undertake the specification search, which has theoretical and computational advantages over existing approaches.

2.1 Existing Approaches

In an important paper Frühwirth-Schnatter and Wagner (2010) (FSW) propose the following framework for nesting both time-varying and time-invariant coefficient speci-

fications. To set the stage, consider again the generic state space model in (1)–(2), and recall that $\boldsymbol{\beta}_t = (\beta_{1,t}, \dots, \beta_{m,t})'$ is the state vector at time t , $\boldsymbol{\beta}_0 = \boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_m)'$ is the vector of initial states, and $\boldsymbol{\Omega} = \text{diag}(\omega_1^2, \dots, \omega_m^2)$ is the covariance matrix for the state transition in (2). To proceed, reparameterize $\gamma_{j,t} = (\beta_{j,t} - \alpha_j)/\omega_j$ for $j = 1, \dots, m$, so that (1)–(2) can be rewritten as

$$\mathbf{y}_t = \mathbf{X}_t \boldsymbol{\alpha} + \mathbf{X}_t \boldsymbol{\Omega}^{\frac{1}{2}} \boldsymbol{\gamma}_t + \boldsymbol{\varepsilon}_t, \quad (3)$$

$$\boldsymbol{\gamma}_t = \boldsymbol{\gamma}_{t-1} + \tilde{\boldsymbol{\eta}}_t, \quad (4)$$

where $\boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$, $\tilde{\boldsymbol{\eta}}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_m)$, and $\boldsymbol{\Omega}^{\frac{1}{2}} = \text{diag}(\omega_1, \dots, \omega_m)$. It is clear that when $\omega_1 = \dots = \omega_m = 0$, the above specification reduces to a standard time-invariant regression model. By allowing the ω_j 's to have different values, the model (3)–(4) can therefore accommodate the possibility that some coefficients are time-varying, whereas others are constant over time. More importantly, by specifying a suitable prior for the ω_j 's, this framework provides a convenient way to shrink the TVP state space model to a more parsimonious constant coefficient regression.

Specifically, FSW propose the following independent normal priors each with a point mass at 0:

$$p(\omega_j) = \pi_{0j} \mathbf{1}(\omega_j = 0) + (1 - \pi_{0j}) \phi(\omega_j; \mu_j, \tau_j^2), \quad (5)$$

where π_{0j} is the *a priori* probability that ω_j equals 0, and $\phi(\cdot; \mu, \sigma^2)$ denotes the normal density with mean μ and variance σ^2 . These priors can be equivalently specified by introducing the indicators

$$d_j = \begin{cases} 0 & \text{with probability } \pi_{0j}, \\ 1 & \text{with probability } 1 - \pi_{0j}, \end{cases}$$

and $\tilde{\omega}_j \sim \mathcal{N}(\mu_j, \tau_j^2)$. Then, $\omega_j = d_j \tilde{\omega}_j$ has the desired distribution in (5).

Each vector $\mathbf{d} = (d_1, \dots, d_m)'$ corresponds to a model in which some coefficients — specifically those with $d_j = 1$ — are time-varying while others are not. Hence, the stochastic variable selection approach of George and McCulloch (1993, 1997) can be adopted to perform the stochastic model specification search. Particularly, FSW derive a Markov Chain Monte Carlo (MCMC) algorithm for sampling the indicators simultaneously with the models parameters. However, the main drawback of this approach is that it becomes computationally infeasible when m is large. Specifically, in one step of the MCMC algorithm it is required to compute the joint probability that $d_1 = i_1, \dots, d_m = i_m$ given

the data and all the other parameters except the ω_j 's, where $(i_1, \dots, i_m) \in \{0, 1\}^m$. This step becomes computationally infeasible when m is large as there are altogether 2^m combinations.

In view of this difficulty, Belmonte, Koop, and Korobolis (2014) (BKK) specify an alternative hierarchical *Lasso* prior directly on the ω_j 's as follows:

$$\begin{aligned} (\omega_j \mid \tau_j^2) &\sim \mathcal{N}(0, \tau_j^2), \\ \tau_j^2 &\sim \mathcal{E}\left(\frac{\lambda^2}{2}\right), \end{aligned} \tag{6}$$

$$\lambda^2 \sim \mathcal{G}(\lambda_{01}, \lambda_{02}), \tag{7}$$

where $\mathcal{E}$ and $\mathcal{G}$ denote respectively the Exponential and the Gamma distributions. This Lasso prior circumvents the model proliferation problem by removing the indicators, and at the same time provides additional shrinkage towards the time-invariant model.

While the specification of BKK addresses the computation issue, this comes at the cost of losing two attractive features of the FSW specification. First, while it is true that “if ω_j is *shrunk* to 0... then we have a model with a constant parameter on predictor j ”, the probability of such shrinking ever occurring is in fact zero, i.e., $\Pr(\omega_j = 0) = 0$. Second, the probability that the indicator $d_j = 0$ has an intuitive interpretation: it is the probability that the corresponding coefficient does not change over time. As such, this is useful for prior elicitation — e.g., some coefficients might be more likely to vary over time than others apriori. On the other hand, prior elicitation in BKK's specification is more difficult. Further, given the posterior draws of the indicators in FSW's specification, it is easy to compute posterior model probabilities for model selection or other probabilities to assess various hypotheses — e.g., output persistence does not change over time — whereas these probabilities are more difficult to compute under the specification of BKK.

2.2 A Tobit Prior

Both the FSW and BKK prior specifications have attractive features: the former has a natural, intuitive interpretation that is useful for prior elicitation and computing a variety of interesting posterior quantities, whereas the latter provides additional shrinkage and leads to a computationally feasible MCMC sampler even when the dimension of the states is large. We therefore aim to have a setup that has the advantages of both

specifications, while circumventing their drawbacks. More specifically, we propose the following censored or Tobit prior on the ω_j 's. For $j = 1, \dots, m$, introduce the latent variable

$$\omega_j^* \sim \mathcal{N}(\mu_j, \tau_j^2),$$

and set

$$\omega_j = \begin{cases} 0 & \text{if } \omega_j^* \leq 0 \\ \omega_j^* & \text{if } \omega_j^* > 0. \end{cases}$$

It is easy to see that the marginal density of ω_j unconditional of ω_j^* is given by

$$p(\omega_j | \tau_j^2) = \Phi\left(-\frac{\mu_j}{\tau_j}\right) \mathbf{1}(\omega_j = 0) + \phi(\omega_j; \mu_j, \tau_j^2) \mathbf{1}(\omega_j > 0), \quad (8)$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution. Additionally, we assume τ_j^2 has the prior as in (6)-(7) to incorporate the Lasso structure.

This proposed prior specification has several appealing features. First, it incorporates the essential elements of both the priors in FSW and BKK. To see the connection to FSW, let $\pi_{0j} = \Phi\left(-\frac{\mu_j}{\tau_j}\right)$ and rewrite the proposed prior in (8) as

$$p(\omega_j | \tau_j^2) = \pi_{0j} \mathbf{1}(\omega_j = 0) + (1 - \pi_{0j}) \phi_{(0, \infty)}(\omega_j; \mu_j, \tau_j^2),$$

where $\phi_{\mathbb{X}}(\cdot; \mu, \sigma^2)$ denotes the density of the truncated $\mathcal{N}(\mu, \sigma^2)$ distribution with support in $\mathbb{X}$. Therefore, it allows for both a non-trivial probability of a certain parameter to be time-invariant and simultaneously the *hierarchical shrinking* that BKK so strongly argue for in forecasting applications. Moreover, the fact that ω_j is automatically restricted to be positive under the Tobit prior—in contrast to both FSW and BKK—turns out to play an important role in extending this approach to state space settings with a full covariance matrix in the state equation, as discussed in further detail in Section 4.

Second, the proposed specification imposes a sensible relationship between the probability of time-invariance π_{0j} and the distribution of ω_j in the time-varying case. Specifically, the farther the truncated mean (normalized by standard deviation) is away from zero, the lower the probability that ω_j is zero, and vice-versa. This means that for values of ω_j that are close to zero, the prior elicits a strong preference for a time-invariant specification, while at the same time, penalizing situations where a high degree of time variation (i.e., a large ω_j) would be associated with a high probability of time invariance. In doing so, it imposes an important and desirable element of shrinkage on the model

space. On the other hand, setting $\mu_j = 0$ implies that π_{0j} is constant and equal to 0.5. Consequently, tuning the hyper-parameters $\lambda_{01}, \lambda_{02}$ such that τ_j^2 is “large” with probability ≈ 1 leads to the FSW specification where d_j is zero with probability 0.5 and $\tilde{\omega}_j$ follows a half-normal, weakly-informative distribution.

Lastly, this prior leads to a straightforward Gibbs sampler. Importantly, in contrast to the MCMC sampler in FSW that becomes computationally intensive in high-dimensional state space models, this Gibbs sampler is fast and is applicable to such settings.

3 Priors and Posterior Computation

To complete the model specification, we assume the following standard independent priors for $\boldsymbol{\alpha}$ and $\boldsymbol{\Sigma}$:

$$\boldsymbol{\alpha} \sim \mathcal{N}(\boldsymbol{\alpha}_0, \boldsymbol{A}_0^{-1}), \quad \boldsymbol{\Sigma} \sim \mathcal{IW}(\nu_0, \boldsymbol{\Sigma}_0),$$

where $\mathcal{IW}$ denotes the inverse Wishart distribution. In what follows, we outline the posterior computation using MCMC, and we give the details in Appendix B. First, define the $m \times 1$ vectors $\boldsymbol{\omega}^* = (\omega_1^*, \dots, \omega_m^*)'$, $\boldsymbol{\omega} = (\omega_1, \dots, \omega_m)'$ and $\boldsymbol{\tau} = (\tau_1, \dots, \tau_m)'$, and stack

$$\boldsymbol{y} = \begin{pmatrix} \boldsymbol{y}_1 \\ \vdots \\ \boldsymbol{y}_T \end{pmatrix}, \quad \boldsymbol{X} = \begin{pmatrix} \boldsymbol{X}_1 \\ \vdots \\ \boldsymbol{X}_T \end{pmatrix}, \quad \boldsymbol{\gamma} = \begin{pmatrix} \gamma_1 \\ \vdots \\ \gamma_T \end{pmatrix}, \quad \boldsymbol{\varepsilon} = \begin{pmatrix} \varepsilon_1 \\ \vdots \\ \varepsilon_T \end{pmatrix}.$$

Then, posterior draws are obtained by sequentially sampling from:

1. $p(\boldsymbol{\alpha} \mid \boldsymbol{y}, \boldsymbol{\gamma}, \boldsymbol{\omega}^*, \boldsymbol{\Sigma}, \boldsymbol{\tau}, \lambda);$
2. $p(\boldsymbol{\gamma} \mid \boldsymbol{y}, \boldsymbol{\alpha}, \boldsymbol{\omega}^*, \boldsymbol{\Sigma}, \boldsymbol{\tau}, \lambda);$
3. $p(\boldsymbol{\Sigma} \mid \boldsymbol{y}, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\omega}^*, \boldsymbol{\tau}, \lambda);$
4. $p(\boldsymbol{\omega}^* \mid \boldsymbol{y}, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\Sigma}, \boldsymbol{\tau}, \lambda);$
5. $p(\boldsymbol{\tau} \mid \boldsymbol{y}, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\omega}^*, \boldsymbol{\Sigma}, \lambda);$
6. $p(\lambda \mid \boldsymbol{y}, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\omega}^*, \boldsymbol{\Sigma}, \boldsymbol{\tau}).$

We note that each ω_j is completely determined by ω_j^* . In fact, under the Tobit prior discussed in Section 2.2, we have $\omega_j = 0$ if $\omega_j^* \leq 0$ and $\omega_j = \omega_j^*$ otherwise. Now, given $\boldsymbol{\omega}$, the model in (3)–(4) is a conventional TVP-VAR. Hence, Steps 1-3 are standard, and we leave the details to Appendix B. Here we discuss the implementations of Steps 4-6.

One feasible approach to sample $\boldsymbol{\omega}^*$ from its full conditional distribution is to simulate each element ω_j^* at a time. To that end, first recall that $\mathbf{X}_t = \mathbf{I}_n \otimes \mathbf{x}'_t$, and $\mathbf{x}_t = (x_{1,t}, \dots, x_{k,t})'$ is a $k \times 1$ vector of deterministic terms and lagged observations. Now, defining $\mathbf{G}_t = \mathbf{X}_t \text{diag}(\boldsymbol{\gamma}_t)$, let

$$\mathbf{G} = \begin{pmatrix} \mathbf{G}_1 \\ \vdots \\ \mathbf{G}_T \end{pmatrix}.$$

Stack (3) over t and rewrite the measurement equation as

$$\mathbf{y} = \mathbf{X}\boldsymbol{\alpha} + \mathbf{G}\boldsymbol{\omega} + \boldsymbol{\varepsilon}. \quad (9)$$

Further, let $\mathbf{g}_j$ denote the j -th column of $\mathbf{G}$, and accordingly, let $\mathbf{G}_{\setminus j}$ be the $Tn \times (m-1)$ matrix obtained by deleting the j -th column in $\mathbf{G}$. Similarly, let $\boldsymbol{\omega}_{\setminus j}$ represent $\boldsymbol{\omega}$ with the j -th row removed. Then,

$$\mathbf{v}_j \equiv \mathbf{y} - \mathbf{X}\boldsymbol{\alpha} - \mathbf{G}_{\setminus j}\boldsymbol{\omega}_{\setminus j} = \mathbf{g}_j\omega_j + \boldsymbol{\varepsilon}.$$

Compute further the following posterior quantities:

$$\begin{aligned} \hat{\tau}_j^2 &= (\tau_j^{-2} + \mathbf{g}'_j(\mathbf{I}_T \otimes \boldsymbol{\Sigma}^{-1})\mathbf{g}_j)^{-1} \\ \hat{\mu}_j &= \hat{\tau}_j^2 (\mu_j/\tau_j^2 + \mathbf{g}'_j(\mathbf{I}_T \otimes \boldsymbol{\Sigma}^{-1})\mathbf{v}_j) \\ \hat{\psi}_j &= \frac{\Phi(\hat{\mu}_j/\hat{\tau}_j)}{\Phi(-\mu_j/\tau_j)} \frac{\hat{\tau}_j}{\tau_j} \exp \left\{ \frac{1}{2} \left(\frac{\hat{\mu}_j^2}{\hat{\tau}_j^2} - \frac{\mu_j^2}{\tau_j^2} \right) \right\} \\ \hat{\pi}_j &= (1 + \hat{\psi}_j)^{-1}. \end{aligned}$$

Then, the conditional density of ω_j^* is the following 2-component mixture of truncated normals

$$p(\omega_j^* | \mathbf{y}, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\omega}_{\setminus j}^*, \boldsymbol{\Sigma}, \boldsymbol{\tau}, \lambda) = \hat{\pi}_j \phi_{(-\infty, 0)}(\omega_j^* | \mu_j, \tau_j^2) + (1 - \hat{\pi}_j) \phi_{(0, \infty)}(\omega_j^* | \hat{\mu}_j, \hat{\tau}_j^2).$$

The derivations can be found in Appendix B. A draw from the above mixture distribution can be obtained as follows. First, get a Bernoulli draw Z with $\mathbb{P}(Z = 0) = \hat{\pi}_j$. If $Z = 0$, sample ω_j^* from the truncated $\mathcal{N}(\omega_j^*; \mu_j, \tau_j^2)$ distribution with support $(-\infty, 0)$; if $Z = 1$, sample ω_j^* instead from the truncated $\mathcal{N}(\omega_j^*; \hat{\mu}_j, \hat{\tau}_j^2)$ distribution with support $(0, \infty)$.

Finally, τ_j^2 and λ conditional on ω_j^* may be sampled exactly as discussed in BKK:

$$\begin{aligned} (\tau_j^{-2} | \lambda, \omega_j^*) &\sim \mathcal{IG} \left(\sqrt{\frac{\lambda^2}{(\omega_j^* - \mu_j)^2}}, \lambda^2 \right), \\ (\lambda^2 | \boldsymbol{\tau}) &\sim \mathcal{G} \left(\lambda_{01} + m, \lambda_{02} + \frac{1}{2} \sum_{j=1}^m \tau_j^2 \right), \end{aligned}$$

where $\mathcal{IG}$ denotes the inverse Gaussian distribution.

4 Extensions

In this section we discuss two important extensions of the basic framework in (1)–(2). Firstly, we allow the error covariance matrix in the state equation to be a full matrix as opposed to a diagonal matrix. This is a main departure from both FSW and BKK, who only consider the case of a diagonal error covariance matrix. Secondly, we generalize the homoskedastic state space model to include stochastic volatility in the measurement equation (1).

Now, consider the following TVP-SVAR:

$$\mathbf{B}_{0,t} \mathbf{y}_t = \mathbf{X}_t^* \boldsymbol{\beta}_t^* + \boldsymbol{\varepsilon}_t, \quad (10)$$

$$\boldsymbol{\beta}_t^* = \boldsymbol{\beta}_{t-1}^* + \boldsymbol{\eta}_t^*, \quad (11)$$

where $\boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_t)$ and $\boldsymbol{\eta}_t^* \sim \mathcal{N}(\mathbf{0}, \tilde{\boldsymbol{\Omega}})$ are independent of each other for all leads and lags and $\mathbf{B}_{0,t}$ is a lower unitriangular matrix—i.e., a lower triangular matrix with unit diagonal elements.² Since the model (10)–(11) is in the structural form, we may assume that $\boldsymbol{\Sigma}_t$ is diagonal with diagonal elements $\sigma_{i,t}^2 = \exp(h_{i,t})$, where defining $\mathbf{h}_t =$

²The advantage of the structural specification in the context of the Tobit prior is that it allows us to work with a single block of coefficients, which greatly simplifies computation. In the course of our empirical application, we compared estimates obtained under the benchmark TVP-SVAR (i.e. without time invariance selection) to those of a reduced form TVP-VAR as in Primiceri (2005). We find the two specifications to give very similar results.

$(h_{1,t}, \dots, h_{n,t})'$ the log-volatilities follow a random walk:

$$\mathbf{h}_t = \mathbf{h}_{t-1} + \boldsymbol{\eta}_t^h, \quad \boldsymbol{\eta}_t^h \sim \mathcal{N}(\mathbf{0}, \mathbf{R}). \quad (12)$$

We assume the random walk is initialized with $\mathbf{h}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{V}_0^h)$ and the prior on the transition covariance is $\mathbf{R} \sim \mathcal{IW}(\nu_0, \mathbf{R}_0)$.

Since $\mathbf{B}_{0,t}$ is lower unitriangular, (10)–(11) can be rewritten as (1)–(2) by appropriately rearranging $\mathbf{X}_t^*$ to include contemporaneous values $\mathbf{y}_t$, as well as $\boldsymbol{\beta}_t^*$ to include the free elements in $\mathbf{B}_{0,t}$. The error covariance matrix $\tilde{\boldsymbol{\Omega}}$ is now a full matrix. We decompose it as follows:

$$\tilde{\boldsymbol{\Omega}} = \boldsymbol{\Omega}^{\frac{1}{2}} \boldsymbol{\Phi} \boldsymbol{\Phi}' \boldsymbol{\Omega}^{\frac{1}{2}},$$

where $\boldsymbol{\Omega}^{\frac{1}{2}} = \text{diag}(\omega_1, \dots, \omega_m)$ is the same as defined in Section 2.1 and $\boldsymbol{\Phi}$ is lower unitriangular.

This type of covariance decomposition was used by Chen and Dunson (2003) and Kinney and Dunson (2007) for a similar purpose in a random-effects setting, and constitutes an extension of the Cholesky decomposition. Note, however, that in order for the decomposition to be unique, sign restrictions must be imposed on $\omega_1, \dots, \omega_m$ (see, for example, Seber, 1977). In their approach, Chen and Dunson (2003) and Kinney and Dunson (2007) ensure this by specifying truncated (i.e. above zero) normal priors on ω_j . However, they do not consider hierarchical shrinkage. In our setting, building a hierarchy upon the truncated normal distribution would lead to computational difficulties since the hyper-parameters would enter the truncated normal integrand, thereby eliminating conjugacy.³ Fortunately, as pointed out in Section 2, the Tobit prior automatically restricts $\omega_j \geq 0$ while still allowing for a straightforward implementation of hierarchical shrinkage through the hyper-parameters of ω_j^* .

Applying the above decomposition, (3)–(4) becomes

$$\mathbf{y}_t = \mathbf{X}_t \boldsymbol{\alpha} + \mathbf{X}_t \boldsymbol{\Omega}^{\frac{1}{2}} \boldsymbol{\Phi} \boldsymbol{\gamma}_t + \boldsymbol{\varepsilon}_t, \quad (13)$$

$$\boldsymbol{\gamma}_t = \boldsymbol{\gamma}_{t-1} + \tilde{\boldsymbol{\eta}}_t, \quad (14)$$

where $\boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_t)$ and $\tilde{\boldsymbol{\eta}}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_m)$ are independent of each other for all leads and lags. The estimation algorithm described in Section 3 can be easily modified to fit the more general model in (13)–(14). For example, to draw from the full conditional

³One alternative in this respect may be to consider the non-linear, non-Gaussian state space algorithm of Chan and Strachan (2012).

distribution of ω , we only need to redefine $\mathbf{G}_t = \mathbf{X}_t \text{diag}(\Phi \gamma_t)$ and proceed as before. Similarly, to sample γ from its full conditional distribution, we can simply redefine $\mathbf{W}_t = \mathbf{X}_t \Omega^{\frac{1}{2}} \Phi$ and everything else follows identically.

To sample the free elements in Φ from its full conditional distribution, first define the matrix

$$\mathbf{F}_t = \begin{pmatrix} \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} \\ \gamma'_{1,t} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \gamma'_{[1,\dots,2],t} & & \vdots \\ \vdots & & \ddots & \mathbf{0} \\ \mathbf{0} & \cdots & \mathbf{0} & \gamma'_{[1,\dots,m-1],t} \end{pmatrix},$$

where $\gamma_{[1,\dots,j],t}$ denotes the row vector $(\gamma_{1,t}, \dots, \gamma_{j,t})$. Let

$$\phi = (\phi_{2,1}, \phi_{3,1}, \phi_{3,2}, \dots, \phi_{m,1}, \dots, \phi_{m,m-1})',$$

i.e., the vector obtained by squeezing the zeros and unit diagonal elements out of $\text{vec}(\Phi')$. Then, rewrite (13) as

$$\mathbf{y}_t = \mathbf{X}_t \alpha + \mathbf{X}_t \Omega^{\frac{1}{2}} \gamma_t + \mathbf{X}_t \Omega^{\frac{1}{2}} \mathbf{F}_t \phi + \epsilon_t. \quad (15)$$

Since (15) is linear in ϕ , if we assume a normal prior for ϕ , sampling ϕ can be done using standard linear regression results (see, e.g., Kroese and Chan, 2014, pp. 237-240). However, it is important to observe that ϕ will typically be very high-dimensional. For example, with $n = 3, p = 3$, ϕ contains 741 elements. Therefore, some form of shrinkage or a tight prior specification will typically be necessary to estimate (13)-(14). In our application of Section 5, and consistent with the main theme of this paper, we specify a Lasso prior on the individual elements of ϕ . Finally, sampling the log-volatilities can also be done using standard algorithms. We follow Primeri (2005) in extending the algorithm of Kim, Shephard, and Chib (1998) to a multivariate setting.

5 Application

In this section we investigate the effects of a fiscal shock on output, taxes and government spending. We estimate the impulse responses of these variables to a shock to government spending. Perotti (2005) and Galí, Vallés and López-Salido (2007) point out that models based upon different theories — neoclassical, Keynesian or neo-Keynesian — can make

conflicting predictions about the responses of macro variables to fiscal shocks. When the empirical evidence is imprecise little can be learned about the support for alternative theories and therefore the likely effects of fiscal policy. The results from this section provide improved inference by producing more precise estimates of impulse responses of output, government spending and receipts to a spending shock. We estimate a VAR for the vector of US variables $\mathbf{w}_t = (t_t, g_t, y_t)'$, where t_t is a measure of government revenue, g_t is government expenditures, and y_t is output. All variables are in logs of real per capita values. Government expenditure consists of government consumption and investment while government revenue is less government transfers. Following Blanchard and Perotti (2002), we use real per capita figures and deflate the variables using the GDP deflator.

As all of these variables display strong trending behaviors, and any test for a unit root will support this conclusion, the temptation is to model them in differences. However differencing removes information on the relationships among the levels of the variables. It is reasonable to conclude that government expenditure should not exceed GDP nor be negative. We can say something more than this as, in the US at least, government expenditure tends to remain fairly stable as a proportion of GDP. There is also reason to believe that government expenditures cannot wander too far from government receipts for too long. We therefore respecify the VAR into a VECM in which we impose stationarity of the difference $t_t - g_t$ and $g_t - y_t$. Imposing the first of these relations expresses a prior belief in Ricardian equivalence. This seems to be a reasonable restriction as evidence on these constraints suggests they are $I(0)$, although the exact form of the process may have structural breaks (see discussion in, for example, Martin, 2009).

Setting $\mathbf{z}_t = (t_t - g_t, g_t - y_t)'$, and assuming three autoregressive lags, the VECM has the form

$$\Delta \mathbf{w}_t = \boldsymbol{\beta}_{0,t} + \boldsymbol{\beta}_{1,t} \mathbf{z}_{t-1} + \sum_{l=1}^3 \boldsymbol{\beta}_{1+l,t} \Delta \mathbf{w}_{t-l} + \boldsymbol{\varepsilon}_t, \quad (16)$$

where $\boldsymbol{\beta}_{0,t}$ is a 3×1 vector of intercepts, $\boldsymbol{\beta}_{1,t}$ a 3×2 matrix of coefficients, and $\boldsymbol{\beta}_{2,t}, \dots, \boldsymbol{\beta}_{4,t}$ are 3×3 matrices of coefficients. In terms of the model in (1), we have

$$\mathbf{y}_t = \Delta \mathbf{w}_t = \begin{pmatrix} \Delta t_t \\ \Delta g_t \\ \Delta y_t \end{pmatrix}, \quad \mathbf{x}_t = \begin{pmatrix} 1 \\ \mathbf{z}_{t-1} \\ \Delta \mathbf{w}_{t-1} \\ \vdots \\ \Delta \mathbf{w}_{t-3} \end{pmatrix}, \quad \boldsymbol{\beta}_t = \text{vec}((\boldsymbol{\beta}_{0,t}, \dots, \boldsymbol{\beta}_{4,t})').$$

Furthermore, following the discussion in Section 4, we allow (16) to have time-varying volatility and the transition equation for β_t has a full error covariance matrix $\tilde{\Omega}$. More precisely, recall that $B_{0,t}$ is lower unittriangular and $\Sigma_t = \text{diag}(e^{h_{1,t}}, \dots, e^{h_{n,t}})$ is diagonal, where $h_{1,t}, \dots, h_{n,t}$ are the log-volatilities. Then, $\epsilon_t \sim \mathcal{N}(\mathbf{0}, Q_t^{-1})$, where $Q_t = B'_{0,t} \Sigma_t^{-1} B_{0,t}$. We refer to this specification as the TVP-SVECM model. This setup is further motivated by the well documented importance of allowing parameters to vary over time when analyzing macroeconomic data (see for example the survey in Koop and Korobilis, 2010). Indeed, Blanchard and Perotti (2002) allow for limited time-variation in the transmission mechanism through various trending specifications along with a quarterly dependence of the parameters. Nevertheless, while they acknowledge that allowing for more general time variation would be appropriate, they conclude that doing so “would have quickly exhausted all degrees of freedom.”

Accordingly, we utilize this framework to compare the inference obtained from the SMSS specification proposed in this article to the benchmark TVP-SVECM. In particular, we seek to ascertain whether the Tobit prior can provide the necessary degree of parsimony such as to yield sufficiently precise inference when the parameter-rich TVP-SVECM fails to do so, while allowing for certain parameters to vary over time with some (nonzero) probability. To this end, we complete the SMSS specification of the VECM outlined above by setting the following hyper-parameters on the priors discussed in Section 3:

$$\begin{aligned} \mathbf{a}_0 &= \mathbf{0}, & \mathbf{A}_0 &= \mathbf{I}_m, \\ \mathbf{h}_0 &= \mathbf{0}, & \mathbf{V}_0^h &= \mathbf{I}_n, \\ \nu_0 &= n + 11, & \mathbf{R}_0 &= 0.01^2(\nu_0 - n - 1)\mathbf{I}_n, \\ \lambda_{01} &= 0.1, & \lambda_{02} &= 0.1. \end{aligned}$$

For comparison, we also consider a version of the SMSS specification where the error covariance matrices in the transition equations for β_t and $\mathbf{h}_t$ are diagonal. In the latter case, $h_{i,0} \sim \mathcal{N}(0, 1)$ and the transition variance for the i th equation is distributed as

$$r_i^2 \sim IG\left(\frac{\nu_0 - n + 1}{2}, \frac{R_{0,i,i}}{2}\right),$$

which corresponds to the marginal distribution for $R_{i,i}$ implied by the inverse-wishart prior.

To estimate the standard TVP-SVECM, we employ the algorithm given in Chan and Jeliazkov (2009), with the priors chosen to match as closely as possible the SMSS prior

specification above. That is, we set $\beta_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_m)$, $\omega_h^2 \sim \mathcal{IG}(n + 3, 0.01)$ and $h_{i,0}, r_i^2$ as given above. Finally, our data essentially coincide with that used by Blanchard and Perotti (2002) but are extended to cover the sample period 1958Q3 to 2011Q4.

5.1 Identification Restrictions

Eliciting impulse response functions necessitates the estimation of *structural* shocks. Blanchard and Perotti (2002) provide the restrictions needed to identify the structural shocks in the VAR framework. It turns out that under these restrictions, structural parameters can be readily recovered from the reduced form covariance matrix. Consequently, we follow this strategy in simulating the posterior of the VECM given in (16) directly, then recovering draws of the necessary structural parameters by solving a system of nonlinear equations for each draw. We briefly discuss the utilized approach in this section, with the details provided in Appendix A.

Specifically, for a fixed time t , equations (2)-(4) in Blanchard and Perotti (2002) may be represented as

$$\begin{pmatrix} 1 & 0 & -a_1 \\ 0 & 1 & -b_1 \\ -c_1 & -c_2 & 1 \end{pmatrix} \begin{pmatrix} \varepsilon_t^t \\ \varepsilon_t^g \\ \varepsilon_t^y \end{pmatrix} = \begin{pmatrix} 1 & a_2 & 0 \\ b_2 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} s_1 & 0 & 0 \\ 0 & s_2 & 0 \\ 0 & 0 & s_3 \end{pmatrix} \begin{pmatrix} u_t^t \\ u_t^g \\ u_t^y \end{pmatrix}, \quad (17)$$

where $u_t^j \sim \mathcal{N}(0, 1)$ for $j = t, g, y$ and $\varepsilon_t = (\varepsilon_t^t, \varepsilon_t^g, \varepsilon_t^y)' \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_t^{-1})$ is the reduced form disturbance in (16). Given a draw of $\mathbf{Q}_t$ (constructed from draws of $\mathbf{B}_{0,t}$ and $\mathbf{\Sigma}_t$), therefore, it is possible to recover draws of the parameters in (17) by solving an appropriate system of equations. There are two key components in the identification strategy set forth by Blanchard and Perotti (2002) that make solving this system relatively straightforward:

1. a_1 and b_1 are given;
2. either (i) $a_2 = 0, b_2 \neq 0$, or (ii) $a_2 \neq 0, b_2 = 0$.

Blanchard and Perotti (2002) set by assumption $b_1 = 0$ and compute a value for $a_1 = 2.08$ (as the average *within-quarter elasticity of net taxes with respect to output*). Because we consider our data to be sufficiently similar, and since our central concern is to compare the precision in inferences obtained from the SMSS and TVP-SVECM specifications,

these values are retained for the purpose of the impulse response analysis reported below. Alternatively, one could specify non-degenerate distributions over a_1 and b_1 , to reflect less dogmatic prior information. Combining samples of a_1, b_1 from such distributions with posterior draws of the model parameters (in simulating impulse response functions) would essentially lead to an exercise in averaging over prior beliefs.

5.2 Results

We implement the Gibbs sampler described in Sections 3 and 4 to obtain 25,000 posterior draws (after a 2,500 burn-in) of the parameters in the VECM model (16). In doing so, we re-scale and re-center the three “first-differenced” series (i.e., $\Delta \mathbf{w}_t$) to match the naive priors. The two integration relations in $\mathbf{z}_t$ are left untouched. Accordingly, all impulse response functions reported below are appropriately adjusted such as to reverse these transformations and match the original series.⁴ Our implementation of the algorithm also utilizes the three Generalized Gibbs steps (Liu and Sabatti, 2000) discussed in Appendix B. The stability of the algorithm was verified by tracking the inefficiency factors of all sampled variables and replicating the simulation run several times.

To demonstrate the effect of the Tobit prior in controlling the time variation of the parameters, we report two types of posterior quantities: (i) the *time-invariance probability* (TIP), which represents the estimated probability that a particular $\beta_{j,t}$ is constant — i.e., $\Pr(\omega_{j,t} = 0 | \mathbf{y})$ — and (ii) the *maximum time variation* (MTV) in posterior means, which is computed as the difference $\max\{E(\beta_{j,t} | \mathbf{y})\}_{t=1}^T - \min\{E(\beta_{j,t} | \mathbf{y})\}_{t=1}^T$. As illustrated in Table 1, estimated TIPs under the SMSS specification with a full error covariance matrix $\tilde{\Omega}$ (SMSS-full) vary from 0.003 (effect of tax on GDP) to 0.923 (effect of third lag in tax on GDP). Consequently, both the magnitude of and variation in the probabilities of time invariance give strong support to this approach over the highly parameterized TVP-SVECM. This is further reinforced by comparing the MTVs reported in Table 2, across the two specifications. Note that these in general follow the same pattern for both the SMSS-full and the TVP-SVECM in the sense that parameter estimates that exhibit relatively more time variation under the TVP-SVECM also vary relatively more under the SMSS-full. Nevertheless, the SMSS-full specification evidently

⁴A similar approach to “standardizing” the series is undertaken by BKK and affects only the prior specification, as long as the transformations are appropriately accounted for in posterior computations. Alternatively, one may work with the original series and employ a *training sample* to specify the priors (for example, as in Primiceri, 2005), although this is operationally more involved. It is worthwhile noting that we apply the same standardization in both the SMSS and the TVP-SVECM specifications.

leads to an overall reduction in time variation across the parameters. It is also worth noting that the two versions of the SMSS specification — with a full or a diagonal error covariance matrix — give fairly similar results.

Table 1: Estimated Time-Invariance Probabilities under the two versions of the SMSS Specification

		SMSS-diag			SMSS-full		
		Tax	Spend	GDP	Tax	Spend	GDP
	Intercept	0.808	0.363	0.851	0.828	0.493	0.841
	$t_{t-1} - g_{t-1}$	0.566	0.569	0.617	0.588	0.605	0.679
	$g_{t-1} - y_{t-1}$	0.830	0.412	0.815	0.802	0.366	0.864
Contem	Tax	-	0.787	0.000	-	0.794	0.003
	Spend	-	-	0.854	-	-	0.873
	GDP	-	-	-	-	-	-
1st lag	Tax	0.733	0.732	0.768	0.727	0.782	0.843
	Spend	0.858	0.792	0.856	0.861	0.793	0.906
	GDP	0.687	0.153	0.819	0.652	0.234	0.879
2nd lag	Tax	0.838	0.845	0.398	0.839	0.873	0.495
	Spend	0.720	0.808	0.833	0.738	0.854	0.879
	GDP	0.783	0.709	0.774	0.763	0.756	0.838
3rd lag	Tax	0.252	0.850	0.873	0.345	0.873	0.923
	Spend	0.804	0.846	0.809	0.827	0.869	0.865
	GDP	0.629	0.856	0.870	0.682	0.900	0.914

There is likewise a strong link between the time-invariance probabilities and maximum time variations (e.g., under SMSS-full): as expected, higher TIPs generally correspond to lower MTVs. More interestingly, higher TIPs appear to be associated with more substantial reductions in MTVs. However, at the lower end of the TIPs, time variation actually increases. For example, the TIP related to the effect of the co-integration term $g_{t-1} - y_{t-1}$ on spending is estimated at 0.366, and the variation in its posterior mean is 0.832, while under the TVP-SVECM the maximum variation is only about 0.494.

The above results suggest that the Tobit prior does not only produce the effect of reducing “unnecessary” time variation, but it also reallocates time variation across the model parameters by reducing such variation in some parameters, while increasing it in others. As made evident in Tables 1 and 2, the SMSS specification (relative to the standard TVP-SVECM) tends to remove time variation from a majority of the model parameters and concentrate it on a select few. Indeed, this highlights the mechanism by which the Tobit prior induces parsimony in a time-varying parameter model.

To further illustrate the effectiveness of the SMSS approach, we compute impulse re-

Table 2: Comparison of Maximum Time Variation in Posterior Means under the SMSS and the TVP-SVECM Specifications

		SMSS-diag			SMSS-full			TVP-SVECM		
		Tax	Spend	GDP	Tax	Spend	GDP	Tax	Spend	GDP
	Intercept	0.047	0.718	0.024	0.041	0.637	0.057	0.268	0.504	0.190
	$t_{t-1} - g_{t-1}$	0.278	0.186	0.191	0.291	0.228	0.236	0.339	0.150	0.109
	$g_{t-1} - y_{t-1}$	0.039	0.609	0.050	0.065	0.832	0.043	0.358	0.494	0.332
Contem	Tax	-	0.052	1.337	-	0.057	1.342	-	0.324	1.256
	Spend	-	-	0.022	-	-	0.025	-	-	0.300
	GDP	-	-	-	-	-	-	-	-	-
1st lag	Tax	0.075	0.085	0.052	0.093	0.092	0.042	0.450	0.309	0.449
	Spend	0.014	0.058	0.013	0.019	0.074	0.012	0.298	0.274	0.220
	GDP	0.118	0.714	0.040	0.166	0.659	0.030	0.513	0.555	0.290
2nd lag	Tax	0.025	0.020	0.286	0.032	0.020	0.280	0.233	0.207	0.499
	Spend	0.081	0.036	0.026	0.096	0.033	0.020	0.416	0.286	0.288
	GDP	0.066	0.090	0.043	0.094	0.100	0.037	0.365	0.361	0.304
3rd lag	Tax	0.396	0.019	0.014	0.364	0.023	0.009	0.566	0.274	0.312
	Spend	0.035	0.027	0.039	0.045	0.034	0.035	0.256	0.253	0.204
	GDP	0.165	0.012	0.012	0.155	0.010	0.011	0.779	0.273	0.370

sponses to spending shocks for taxes, spending and GDP, in the spirit of Blanchard and Perotti (2002). In particular, we employ the identification restrictions discussed in Section 5.1, along with the procedure outlined in Appendix A, to decompose the reduced form covariance matrix $\mathbf{Q}_t^{-1}$. We then use the the resulting posterior draws of the structural parameters (together with draws of β) to compute the percent changes in taxes, spending and GDP to a one percent increase in spending. Therefore, in contrast to Blanchard and Perotti (2002), the impulse responses are expressed in terms of elasticities rather than absolute levels. Moreover, because we are working in a time-varying parameter context, we choose a specific (within-sample) time period — 2001Q1 to 2005Q4 (5 years/20 quarters) — to conduct the analysis.

The resulting impulse response functions are summarized in Figure 1: the the HPD intervals (shaded regions) are constructed as [16%, 84%] while the *estimated* impulse responses (solid lines) are posterior medians. In the figure, the top two rows contain the impulse responses generated by the SMSS-diag and the SMSS-full specifications, while in the bottom row are those corresponding to the standard TVP-SVECM. It is clear that that the TVP-SVECM generally yields wider HPD intervals relative to both versions of the SMSS. This is particularly evident for taxes and GDP (where incidentally SMSS tends to assign higher probabilities of time-invariance). On the other hand, the two

versions of the SMSS give very similar impulse responses. The more accurate inference delivered by the SMSS approach is important since, as pointed out by Perotti (2005), the effects of fiscal policy have become weaker over time and so it is important to have a method that can distinguish when the effect is different from zero.

Figure 1: Comparison of IRFs generated by the SMSS with a diagonal transition covariance, the SMSS with a full transition covariance and the benchmark TVP-SVECM for 20 quarters following a 1% Shock in Spending in 2000Q1

The TVP-SVECM yields somewhat peculiar (median) impulse responses for taxes and GDP, in contrast to the SMSS. For example, the TVP-SVECM predicts that both taxes and GDP responses are initially negative and require about three years to cross into positive territory, where they remain in the long run. This is counter-intuitive and appears to be an artifact of the imprecision with which the TVP-SVECM estimates the impulse responses. The reported HPD intervals contain zero for all of this initial period. Under the two SMSS specifications, on the other hand, the responses of taxes and GDP are estimated to be positive for the entire period in examination — a much

more plausible result and one consistent with the conclusions drawn in Blanchard and Perotti (2002).

Another interesting feature of the results is the clear effect of imposing of Ricardian equivalence in the responses. This feature is imposed upon the three models and the long run effect is obviously present in all models, but both versions of the SMSS give much nicer bounds. To highlight the comparison, Figure 2 plots the median responses for taxes and spending (without the HPD limits), as generated by the SMSS and standard TVP-SVECM specifications. Indeed, it seems that taxes are always chasing spending with a delayed response, until spending stabilizes and they converge. This is very clear in the SMSS, but noticeably more distorted with the TVP-SVECM, where taxes dip into the negative over the first five years. Based on the TVP-SVECM results alone, therefore, distinguishing this effect is made difficult by the imprecision of the estimates related to the parameter proliferation present in this specification. In contrast, employing the Tobit prior leads to a much clearer representation and greatly facilitates economic interpretation of the simulation results.

Figure 2: Median Impulse Response of Taxes (dotted lines) and Spending (solid lines) to a 1% Shock in Spending in 2000Q1

6 Concluding Remarks

Time-varying VARs are widely used for studying the dynamic effects of structural shocks on key economic variables through the estimation of impulse response functions. However, since these models are highly parameterized, inference is typically imprecise and conclusions are often difficult to draw. In this paper we present a new method that allows the data to decide whether parameters are time varying or time invariant in a VAR, thus allowing the model to automatically switch to a more parsimonious specification when the time-variation of the coefficient is small. By introducing a Tobit prior on the

variances in the state equations, the task of computing the many indicators is greatly simplified.

We apply the new methodology and the computation scheme to a brief study of responses in US government receipts, government expenditure and GDP to a fiscal policy shock. Compared to those results obtained under an unrestricted TVP-SVECM, we find some differences in the median response paths but the most significant improvement is in the precision of estimation of the responses. This is an important result since these effects have become weaker over time (Perotti, 2005) so greater precision in estimation is necessary for accurate inference.

Appendix A: Recovering Structural Parameters from the Precision Matrix

In terms of (17), for a fixed time t , define $\boldsymbol{\psi} = -\begin{pmatrix} a_1 \\ b_1 \end{pmatrix}$, $\boldsymbol{\gamma} = -\begin{pmatrix} c_1 & c_2 \end{pmatrix}$, $\boldsymbol{\Delta} = \begin{pmatrix} 1 & a_2 \\ b_2 & 1 \end{pmatrix}$, $\boldsymbol{\Xi}^{\frac{1}{2}} = \begin{pmatrix} s_1 & 0 \\ 0 & s_2 \end{pmatrix}$ and rewrite the system as

$$\begin{pmatrix} \mathbf{I} & \boldsymbol{\psi} \\ \boldsymbol{\gamma} & 1 \end{pmatrix} \boldsymbol{\varepsilon}_t = \begin{pmatrix} \boldsymbol{\Delta} & \\ & 1 \end{pmatrix} \begin{pmatrix} \boldsymbol{\Xi}^{\frac{1}{2}} & \\ & s_3 \end{pmatrix} \mathbf{u}_t, \quad \mathbf{u}_t \stackrel{iid}{\sim} \mathcal{N}(0, \mathbf{I}_3).$$

Hence,

$$\boldsymbol{\varepsilon}_t \sim \mathcal{N}\left(0, (\mathbf{K}'\mathbf{K})^{-1}\right),$$

where

$$\mathbf{K} = \begin{pmatrix} \boldsymbol{\Xi}^{-\frac{1}{2}} & \\ & \frac{1}{s_3} \end{pmatrix} \begin{pmatrix} \boldsymbol{\Delta}^{-1} & \\ & 1 \end{pmatrix} \begin{pmatrix} \mathbf{I} & \boldsymbol{\psi} \\ \boldsymbol{\gamma} & 1 \end{pmatrix}.$$

Given posterior draws of $\mathbf{B}_{0,t}$ and $\boldsymbol{\Sigma}_t$, therefore, posterior draws of the structural parameters and variances (e.g., the six free parameters in $\mathbf{K}$) can be obtained by solving $\mathbf{K}'\mathbf{K} = \mathbf{Q}_t \equiv \mathbf{B}_{0,t}'\boldsymbol{\Sigma}_t^{-1}\mathbf{B}_{0,t}$. More precisely, partitioning

$$\mathbf{Q}_t = \begin{pmatrix} \mathbf{Q}_{11} & \mathbf{q}_{12} \\ \mathbf{q}_{12}' & q_3 \end{pmatrix},$$

and denoting $\boldsymbol{\Theta} = (\boldsymbol{\Delta}\boldsymbol{\Xi}\boldsymbol{\Delta}')^{-1}$ leads to the system of nonlinear equations:

$$\boldsymbol{\Theta} + \frac{1}{s_3^2} \boldsymbol{\gamma}'\boldsymbol{\gamma} = \mathbf{Q}_{11} \tag{18}$$

$$\boldsymbol{\Theta}\boldsymbol{\psi} + \frac{1}{s_3^2} \boldsymbol{\gamma}' = \mathbf{q}_{12} \tag{19}$$

$$\boldsymbol{\psi}'\boldsymbol{\Theta}\boldsymbol{\psi} + \frac{1}{s_3^2} = q_3 \tag{20}$$

As discussed in the text, the following two conditions make this system easy to solve:

1. a_1 and b_1 are given (more specifically, $b_1 = 0$), and therefore, $\boldsymbol{\psi}$ is known
2. either (i) $a_2 = 0, b_2 \neq 0$, which implies that $\boldsymbol{\Delta}$ is *lower* triangular with 1's on the

diagonal, or (ii) $a_2 \neq 0, b_2 = 0$, which implies that Δ is *upper* triangular with 1's on the diagonal.

Given this, γ , Δ , $\Xi^{\frac{1}{2}}$, and s_3 can be solved for in three steps:

1. Solve for s_3 :

$$\begin{aligned}\psi' \Theta \psi + \frac{1}{s_3^2} \psi' \gamma' \gamma \psi &= \psi' \mathbf{Q}_{11} \psi \\ \psi' \Theta \psi + \frac{1}{s_3^2} \psi' \gamma' &= \psi' \mathbf{q}_{12} \\ \psi' \Theta \psi + \frac{1}{s_3^2} &= q_3,\end{aligned}$$

implies that

$$q_3 - \frac{1}{s_3^2} + \frac{1}{s_3^2} [s_3^2 (\psi' \mathbf{q}_{12} - q_3) - 1]^2 = \psi' \mathbf{Q}_{11} \psi,$$

and therefore,

$$s_3 = \frac{\sqrt{\psi' \mathbf{Q}_{11} \psi - 2\psi' \mathbf{q}_{12} + q_3}}{\psi' \mathbf{q}_{12} - q_3}. \quad (21)$$

2. Solve for γ , given s_3 :

Substituting (18) into (19) leads to

$$\gamma' \gamma \psi - \gamma' + s_3^2 (\mathbf{q}_{12} - \mathbf{Q}_{11} \psi) = 0. \quad (22)$$

It can be readily verified that the solution to (22) is of the form

$$\gamma' = \frac{1 \pm \sqrt{1 - 4s_3^2 \psi' (\mathbf{q}_{12} - \mathbf{Q}_{11} \psi)}}{2\psi' (\mathbf{q}_{12} - \mathbf{Q}_{11} \psi)} (\mathbf{q}_{12} - \mathbf{Q}_{11} \psi). \quad (23)$$

3. Solve for s_1 , s_2 and either a_2 or b_2 , given s_3 , γ :

From (18), we have

$$\Delta \Xi \Delta' = \left(\mathbf{Q}_{11} - \frac{1}{s_3^2} \gamma' \gamma \right)^{-1}. \quad (24)$$

Therefore,

(a) $a_2 = 0, b_2 \neq 0$: s_1, s_2, b_2 are found by the LDL decomposition of the 2×2 matrix $\left(\mathbf{Q}_{11} - \frac{1}{s_3^2} \gamma' \gamma \right)^{-1}$.

- (b) $a_2 \neq 0, b_2 = 0$: s_1, s_2, a_2 are found by the UDL decomposition of the 2×2 matrix $\left(\mathbf{Q}_{11} - \frac{1}{s_3^2} \boldsymbol{\gamma}' \boldsymbol{\gamma}\right)^{-1}$.

Of course, one caveat here is that the $\pm$ in (23) generates two possible values for $\boldsymbol{\gamma}$. Indeed, there are two solutions to (22). However, only one of these solutions leads to a correct decomposition of $\left(\mathbf{Q}_{11} - \frac{1}{s_3^2} \boldsymbol{\gamma}' \boldsymbol{\gamma}\right)^{-1}$ in Step 3. To solve this problem, we write a script that computes both solutions from (23) and checks which one leads to a correct solution of the entire system.

Appendix B: MCMC Sampler

In this appendix we provide the details of the MCMC sampler outlined in Section 3. First, we derive the expression for $\hat{\psi}_j$. Note that the joint conditional density for (ω_j^*, ω_j) is given by

$$p(\omega_j^*, \omega_j | \mathbf{y}, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\omega}_{\setminus j}^*, \boldsymbol{\Sigma}, \boldsymbol{\tau}, \lambda) \propto \phi(\mathbf{v}; \mathbf{g}_j \boldsymbol{\omega}_j, \mathbf{I}_T \otimes \boldsymbol{\Sigma}) \phi(\omega_j^*; \mu_j, \tau_j^2) \mathbf{1}(\omega_j^* \leq 0) \mathbf{1}(\omega_j = 0) \\ + \phi(\mathbf{v}; \mathbf{g}_j \boldsymbol{\omega}_j, \mathbf{I}_T \otimes \boldsymbol{\Sigma}) \phi(\omega_j^*; \mu_j, \tau_j^2) \mathbf{1}(\omega_j^* > 0) \mathbf{1}(\omega_j = \omega_j^*).$$

Hence, marginalizing over ω_j yields

$$p(\omega_j^* | \mathbf{y}, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\omega}_{\setminus j}^*, \boldsymbol{\Sigma}, \boldsymbol{\tau}, \lambda) \propto \phi(\mathbf{v}_j; \mathbf{0}, \mathbf{I}_T \otimes \boldsymbol{\Sigma}) \phi(\omega_j^*; \mu_j, \tau_j^2) \mathbf{1}(\omega_j^* \leq 0) \\ + \phi(\mathbf{v}_j; \mathbf{g}_j \boldsymbol{\omega}_j^*, \mathbf{I}_T \otimes \boldsymbol{\Sigma}) \phi(\omega_j^*; \mu_j, \tau_j^2) \mathbf{1}(\omega_j^* > 0) \\ \propto \Phi\left(-\frac{\mu_j}{\tau_j}\right) \phi(\mathbf{v}_j; \mathbf{0}, \mathbf{I}_T \otimes \boldsymbol{\Sigma}) \phi_{(-\infty, 0)}(\omega_j^*; \mu_j, \tau_j^2) \\ + \Phi\left(\frac{\hat{\mu}_j}{\hat{\tau}_j}\right) \frac{\phi(\mathbf{v}_j; \mathbf{g}_j \boldsymbol{\omega}_j^*, \mathbf{I}_T \otimes \boldsymbol{\Sigma}) \phi(\omega_j^*; \mu_j, \tau_j^2)}{\phi(\omega_j^*; \hat{\mu}_j, \hat{\tau}_j^2)} \phi_{(0, \infty)}(\omega_j^*; \hat{\mu}_j, \hat{\tau}_j^2) \\ \propto \phi_{(-\infty, 0)}(\omega_j^*; \mu_j, \tau_j^2) + \hat{\psi}_j \phi_{(0, \infty)}(\omega_j^*; \hat{\mu}_j, \hat{\tau}_j^2),$$

where

$$\hat{\psi}_j = \frac{\Phi(\hat{\mu}_j/\hat{\tau}_j)}{\Phi(-\mu_j/\tau_j)} \times \frac{\phi(\mathbf{v}_j; \mathbf{g}_j \boldsymbol{\omega}_j^*, \mathbf{I}_T \otimes \boldsymbol{\Sigma}) \phi(\omega_j^*; \mu_j, \tau_j^2)}{\phi(\mathbf{v}_j; \mathbf{0}, \mathbf{I}_T \otimes \boldsymbol{\Sigma}) \phi(\omega_j^*; \hat{\mu}_j, \hat{\tau}_j^2)} \\ = \frac{\Phi(\hat{\mu}_j/\hat{\tau}_j)}{\Phi(-\mu_j/\tau_j)} \times \frac{\hat{\tau}_j}{\tau_j} \times \exp \left\{ -\frac{1}{2} \left((\mathbf{v}_j - \mathbf{g}_j \boldsymbol{\omega}_j^*)' (\mathbf{I}_T \otimes \boldsymbol{\Sigma}^{-1}) (\mathbf{v}_j - \mathbf{g}_j \boldsymbol{\omega}_j^*) \right. \right. \\ \left. \left. - \mathbf{v}_j' (\mathbf{I}_T \otimes \boldsymbol{\Sigma}^{-1}) \mathbf{v}_j + \frac{1}{\tau_j^2} (\omega_j^* - \mu_j)^2 - \frac{1}{\hat{\tau}_j^2} (\omega_j^* - \hat{\mu}_j)^2 \right) \right\}. \quad (25)$$

However, since

$$\begin{aligned}\frac{\hat{\mu}_j}{\hat{\tau}_j^2} - \frac{\mu_j}{\tau_j^2} &= \mathbf{g}'_j (\mathbf{I}_T \otimes \Sigma^{-1}) \mathbf{v}_j \\ \frac{1}{\hat{\tau}_j^2} - \frac{1}{\tau_j^2} &= \mathbf{g}'_j (\mathbf{I}_T \otimes \Sigma^{-1}) \mathbf{g}_j,\end{aligned}$$

(25) simplifies to

$$\hat{\psi}_j = \frac{\Phi(\hat{\mu}_j/\hat{\tau}_j)}{\Phi(-\mu_j/\tau_j)} \frac{\hat{\tau}_j}{\tau_j} \exp \left\{ \frac{1}{2} \left(\frac{\hat{\mu}_j^2}{\hat{\tau}_j^2} - \frac{\mu_j^2}{\tau_j^2} \right) \right\}.$$

Next, we provide the details of Steps 1-3 in the MCMC sampler. For Step 1, sample

$$(\boldsymbol{\alpha} | \mathbf{y}, \boldsymbol{\gamma}, \boldsymbol{\omega}^*, \Sigma, \boldsymbol{\tau}, \lambda) \sim \mathcal{N}(\hat{\boldsymbol{\alpha}}, \hat{\mathbf{A}}^{-1}),$$

where

$$\hat{\mathbf{A}} = \mathbf{A}_0 + \mathbf{X}'(\mathbf{I}_T \otimes \Sigma^{-1})\mathbf{X}, \quad \hat{\boldsymbol{\alpha}} = \hat{\mathbf{A}}^{-1} (\mathbf{A}_0 \boldsymbol{\alpha}_0 + \mathbf{X}'(\mathbf{I}_T \otimes \Sigma^{-1})(\mathbf{y} - \mathbf{W}\boldsymbol{\gamma})),$$

and $\mathbf{W}$ is obtained by defining $\mathbf{W}_t = \mathbf{X}\boldsymbol{\Omega}^{\frac{1}{2}}$ and stacking

$$\mathbf{W} = \begin{pmatrix} \mathbf{W}_1 \\ \vdots \\ \mathbf{W}_T \end{pmatrix}.$$

For Step 2, we first define

$$\mathbf{H} = \begin{pmatrix} \mathbf{I}_m & & & \\ -\mathbf{I}_m & \mathbf{I}_m & & \\ & \ddots & \ddots & \\ & & -\mathbf{I}_m & \mathbf{I}_m \end{pmatrix}.$$

Then, sample

$$(\boldsymbol{\gamma} | \mathbf{y}, \boldsymbol{\alpha}, \boldsymbol{\omega}^*, \Sigma, \boldsymbol{\tau}, \lambda) \sim \mathcal{N}(\hat{\boldsymbol{\gamma}}, \hat{\mathbf{\Gamma}}^{-1})$$

using the precision-based sampler in Chan and Jeliazkov (2009), where

$$\hat{\mathbf{\Gamma}} = \mathbf{H}'\mathbf{H} + \mathbf{W}'(\mathbf{I}_T \otimes \Sigma^{-1})\mathbf{W}, \quad \hat{\boldsymbol{\gamma}} = \hat{\mathbf{\Gamma}}^{-1} \mathbf{W}'(\mathbf{I}_T \otimes \Sigma^{-1})(\mathbf{y} - \mathbf{X}\boldsymbol{\alpha}).$$

Finally, for a model with a time-invariant Σ , obtain

$$(\Sigma | \mathbf{y}, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\omega}^*, \boldsymbol{\tau}, \lambda) \sim \mathcal{IW}(\nu_0 + T, \hat{\Sigma}),$$

where

$$\hat{\Sigma} = \Sigma_0 + \sum_{t=1}^T \left(\mathbf{y}_t - \mathbf{X}_t \left(\boldsymbol{\alpha} + \Omega^{\frac{1}{2}} \boldsymbol{\gamma}_t \right) \right) \left(\mathbf{y}_t - \mathbf{X}_t \left(\boldsymbol{\alpha} + \Omega^{\frac{1}{2}} \boldsymbol{\gamma}_t \right) \right)'$$

In specifications with a time-varying Σ_t , this step is replaced by a routine stochastic volatility draw, e.g., Primiceri (2005). For the remaining parameters, recall that the details related to the conditional sampling of $\boldsymbol{\omega}^*$, $\boldsymbol{\tau}$, and λ (Steps 4-6) are provided in Section 3.

A few remarks regarding this algorithm are in order. First, one may readily observe that $\boldsymbol{\alpha}$ and $\boldsymbol{\gamma}$ have an analytically tractable joint distribution of the form

$$(\boldsymbol{\alpha}, \boldsymbol{\gamma} | \mathbf{y}, \boldsymbol{\omega}^*, \Sigma, \boldsymbol{\tau}, \lambda) \sim \mathcal{N}(\hat{\boldsymbol{\delta}}, \hat{\Delta}^{-1}), \quad (26)$$

where

$$\begin{aligned} \hat{\Delta} &= \begin{pmatrix} \mathbf{A}_0 & \\ & \mathbf{H}'\mathbf{H} \end{pmatrix} + \begin{pmatrix} \mathbf{X}' \\ \mathbf{W}' \end{pmatrix} (\mathbf{I}_T \otimes \Sigma^{-1}) \begin{pmatrix} \mathbf{X} & \mathbf{W} \end{pmatrix} \\ \hat{\boldsymbol{\delta}} &= \hat{\Delta}^{-1} \left(\begin{pmatrix} \mathbf{A}_0 \mathbf{a}_0 \\ \mathbf{0} \end{pmatrix} + \begin{pmatrix} \mathbf{X}' \\ \mathbf{W}' \end{pmatrix} (\mathbf{I}_T \otimes \Sigma^{-1}) \mathbf{y} \right). \end{aligned}$$

However, the $mT \times mT$ matrix $\hat{\Delta}$ presents a number of computational difficulties. In particular, while $\hat{\Delta}$ is indeed a sparse matrix, it does not inherit a very convenient sparsity structure; operations such as inversion and Cholesky decomposition (both crucial to sampling from (26)) are computationally intensive, and in fact, quite inhibitive in large dimensional settings.

On the other hand, sampling $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ sequentially — i.e., drawing $(\boldsymbol{\alpha} | \boldsymbol{\gamma}, \cdot)$ followed by $(\boldsymbol{\gamma} | \boldsymbol{\alpha}, \cdot)$ — does not appear to generate significant autocorrelation in the MCMC chain, to the extent that it would justify the additional computational burden necessary to sample $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ in a single block. In fact, a battery of tests applied under various settings and with various data consistently found that the computation time required to achieve the same level of efficiency was much lower under the sequential sampling scheme, in comparison to the blocking approach.

Moreover, the efficiency of the proposed algorithm may be improved by augmenting the six steps outlined in Section 3 with one or more *Generalized Gibbs* (GG) steps, as proposed in Liu and Sabatti (2000). The basic idea of such an augmentation is to transform the draws obtained by recursive, conditional sampling within a typical Gibbs loop in such a way as to preserve the invariant target distribution of the Markov chain. Because the transformation involves additional randomly sampled quantities, this has the potential to introduce randomness into otherwise highly autocorrelated MCMC chains, and hence, boost sampling efficiency.

To this end, we identify three potentially beneficial GG moves in our context. We present these transformations assuming that $\mu_j = 0$ — an assumption maintained in all practical implementations of the algorithm outlined in Section 3. Generalizing to cases with non-zero μ_j is conceptually straightforward, but requires more involved sampling procedures, as subsequently discussed. In particular, assuming for the moment a diagonal error covariance in the state equation, a combination of one or more of the following GG steps have been found to improve the mixing of the MCMC chain generated by the proposed Gibbs algorithm:

1. Given the current draws of $\boldsymbol{\gamma}, \tau^2, \lambda^2$, sample for each $j = 1, \dots, m$

$$(\vartheta_j \mid \boldsymbol{\gamma}_j, \tau_j^2, \lambda^2) \sim \mathcal{GIG} \left(\frac{T-3}{2}, \tau_j^2 \lambda^2, \boldsymbol{\gamma}_j' \widetilde{\mathbf{H}}' \widetilde{\mathbf{H}} \boldsymbol{\gamma}_j \right), \quad (27)$$

and apply the transformations

$$\begin{aligned} \tau_j^{2(\text{new})} &= \tau_j^2 / \vartheta_j \\ \omega_j^{*(\text{new})} &= \omega^* / \sqrt{\vartheta_j}, \quad \omega_j^{(\text{new})} = \omega_j^{*(\text{new})} \mathbf{1} \left(\omega_j^{*(\text{new})} > 0 \right) \\ \boldsymbol{\gamma}_j^{(\text{new})} &= \boldsymbol{\gamma}_j \sqrt{\vartheta_j}, \end{aligned}$$

where $\boldsymbol{\gamma}_j$ is a $T \times 1$ vector of the elements in $\boldsymbol{\gamma}$ that correspond to the j th covariate, and

$$\widetilde{\mathbf{H}} = \begin{pmatrix} 1 & & & & \\ -1 & 1 & & & \\ & \ddots & \ddots & & \\ & & & -1 & 1 \end{pmatrix}.$$

Likewise, $\mathcal{GIG}(\cdot)$ denotes the *Generalized Inverse Gamma* distribution and may be efficiently sampled from using the rejection method of Dagpunar (1989).

2. Given the current draws of $\boldsymbol{\omega}^*, \boldsymbol{\tau}_2^2, \lambda^2$, sample

$$(\vartheta \mid \boldsymbol{\omega}^*, \boldsymbol{\tau}_2^2, \lambda^2) \sim \mathcal{G} \left(\lambda_{01} + \frac{m}{2}, \lambda_{02} \lambda^2 + \frac{1}{2} \sum_{j=1}^m \frac{\omega_j^{*2}}{\tau_j^2} \right), \quad (28)$$

and apply the transformations

$$\lambda^{2(\text{new})} = \lambda^2 \vartheta, \quad \boldsymbol{\tau}^{2(\text{new})} = \boldsymbol{\tau}^2 / \vartheta.$$

3. Given the current draws of $\boldsymbol{\alpha}, \boldsymbol{\gamma}_1, \boldsymbol{\omega}$, and defining

$$\widehat{\boldsymbol{\Theta}} = \mathbf{I}_{nT} + \boldsymbol{\Omega}^{\frac{1}{2}} \mathbf{A}_0 \boldsymbol{\Omega}^{\frac{1}{2}}, \quad \widehat{\boldsymbol{\vartheta}} = \widehat{\boldsymbol{\Theta}}^{-1} \left(\boldsymbol{\gamma}_1 - \boldsymbol{\Omega}^{\frac{1}{2}} \mathbf{A}_0 (\boldsymbol{\alpha} - \boldsymbol{\alpha}_0) \right),$$

where $\boldsymbol{\gamma}_1$ a $m \times 1$ vector of the elements in $\boldsymbol{\gamma}$ that correspond to the period $t = 1$, sample

$$(\boldsymbol{\vartheta} \mid \boldsymbol{\alpha}, \boldsymbol{\gamma}_1, \boldsymbol{\omega}) \sim \mathcal{N} \left(\widehat{\boldsymbol{\vartheta}}, \widehat{\boldsymbol{\Theta}}^{-1} \right),$$

and apply the transformations

$$\boldsymbol{\alpha}^{(\text{new})} = \boldsymbol{\alpha} + \boldsymbol{\omega} \odot \boldsymbol{\vartheta}, \quad \boldsymbol{\gamma}^{(\text{new})} = \boldsymbol{\gamma} - \boldsymbol{\vartheta},$$

where $\odot$ represents element-by-element multiplication.

It is worthwhile to note that transformations 1 and 2 incorporate a straightforward extension of the *scale group* explicitly discussed in Liu and Sabatti (2000). The third GG move presented above is based on an extension of the *translation group*, and in fact, is nearly identical to the transformation applied in the context of state-space models in example 4.1 of Liu and Sabatti (2000). Therefore, their Theorem 1 guarantees that the target posterior is preserved under all three types of moves, given that $\mu_j = 0$ holds. To that end, GG move 1 can be similarly implemented in the case that $\mu_j \neq 0$, but the distributions (27) from which the scales $\{\vartheta_j\}$ are sampled would no longer be valid; the resulting distributions that ensure the invariance of the target posterior (with $\mu_j \neq 0$) are of nonstandard form, and hence, would require alternative implementations of the sampling algorithms. Moves 2 and 3, however, are not materially affected by the $\mu_j = 0$ assumption (one would only require replacing ω_j^* with $\omega_j^* - \mu_j$ in (28)).

In our experience, implementing all three moves into the Gibbs sampler of section 3, and to the extension with stochastic volatility but a diagonal error covariance in the

state equation, tends to increase execution time by approximately 10%. In exchange, we have observed a reduction in *inefficiency factors* — in particularly favorable cases — of up to 15 times. In other scenarios, however, the improvement in sampling efficiency is much less pronounced. Therefore, the decision regarding whether or not to include the GG moves into the sampler generally relies on trial and error. In the application of section 5, we incorporate all three moves outlined above; doing so leads to reductions in inefficiency factors of up to four times (depending on the variable) and noticeably more stable posterior estimates (given the same number of retained and discarded simulation draws).

It is important to point out, however, that Move 1 does not apply to the most general model with a full $\tilde{\Omega}$, and the presence of Φ as a result of the decomposition described in Section 4 precludes a straightforward GG step in this case (nevertheless, GG Moves 2 and 3 above remain valid). As an alternative, Kinney and Dunson (2007) propose a parameter expansion scheme by letting $\tilde{\eta}_{j,t} \sim \mathcal{N}(0, \zeta_j)$ and introducing an additional step to sample ζ_j conditional on other parameters. In our application, however, this does not appear to improve mixing. The reason for this is likely related to the discussion in Wu and Liu (1999) (see Remark 3 on p. 1268 regarding the conditional PX-DA algorithm) and the fact that ω_j is modeled with a hierarchical prior. More importantly, the Gibbs sampler with Lasso priors on ϕ exhibits satisfactory mixing properties without parameter expansion, and therefore, we do not pursue this issue further.

As a final remark, we emphasize that a certain amount of care should be taken in handling the large matrices $\mathbf{W}$ ($nT \times mT$) and $\mathbf{G}$ ($nT \times m$) (and for the full $\tilde{\Omega}$ case, the $mT \times m(m-1)/2$ matrix $\mathbf{F} = (\mathbf{F}'_1, \dots, \mathbf{F}'_T)'$). In particular, it should be noted that it is unnecessary — and impractical — to reconstruct these matrices entirely at each iteration of the Gibbs sampler. Rather, observing that both matrices are primarily sparse in nature, a reasonable algorithm would proceed by pre-allocating space prior to the commencement of the Gibbs loop and only updating the necessary elements in the course of execution.

In model (3)-(4), for example, a prototype $\mathbf{W}$ can be preconstructed (say, by letting $\mathbf{W} = \mathbf{I}_{nT} \otimes \boldsymbol{\iota}'_m$, where $\boldsymbol{\iota}_m$ denotes a $m \times 1$ vector of ones), with the indices of the mT nonzero elements stored in w_{nz} . Then, at each iteration, it is only required to update the nonzero elements as

$$\mathbf{W}(w_{\text{nz}}) = \text{vec} \left(\mathbf{X} \odot \begin{pmatrix} \boldsymbol{\omega} & \cdots & \boldsymbol{\omega} \end{pmatrix} \right), \quad (29)$$

where it is assumed that $\mathbf{W}(w_{\text{nz}})$ represents the $mT \times 1$ vector of the nonzero elements in $\mathbf{W}$ (e.g. this notation is consistent with the syntax of a variety of popular scripting software, such as Matlab).

A similar approach provides an efficient procedure for handling $\mathbf{G}$ as well, except in this case it is operationally easier to work directly with its transpose $\mathbf{G}'$. Accordingly, one may initialize the sampler by setting $\mathbf{G}' = \mathbf{X}'$. Storing the indices of the mT nonzero elements in g'_{nz} , these elements are updated at each iteration as

$$\mathbf{G}'(g'_{\text{nz}}) = \text{vec}(\mathbf{X}) \odot \boldsymbol{\gamma}. \quad (30)$$

References

- M. Banbura, D. Giannone, and L. Reichlin. Large Bayesian vector auto regressions. *Journal of Applied Econometrics*, 25(1):71–92, 2010.
- M. Belmonte, G. Koop, and D. Korobilis. Hierarchical shrinkage in time-varying parameter models. *Journal of Forecasting*, 33:80–94, 2014.
- O. Blanchard and R. Perotti. An empirical characterization of the dynamic effects of changes in government spending and taxes on output. *The Quarterly Journal of Economics*, 117:1329–1368, 2002.
- A. Carriero, T. Clark, and M. Marcellino. Bayesian VARs: specification choices and forecast accuracy. Federal Reserve Bank of Cleveland Working Paper 1112, 2011.
- J. C. C. Chan and I. Jeliazkov. Efficient Simulation and Integrated Likelihood Estimation in State Space Models. *International Journal of Mathematical Modelling and Numerical Optimisation*, 1:101-120, 2009.
- J. C. C. Chan, G. Koop, R. Leon-Gonzalez, and R. Strachan. Time varying dimension models. *Journal of Business and Economic Statistics*, 30(3):358–367, 2012.
- J. C. C. Chan and R. Strachan. Estimation in non-linear non-Gaussian state space models with precision-based methods, CAMA Working Paper, 2012.
- Z. Chen and D. B. Dunson. Random effects selection in linear mixed models. *Biometrics*, 59(4): 762–769, 2013.
- T. Cogley and T. J. Sargent. Drifts and volatilities: monetary policies and outcomes in the post WWII US. *Review of Economic Dynamics*, 8(2):262–302, 2005.
- T. Cogley, G. Primiceri, and T. Sargent. Inflation-gap persistence in the U.S. *American*

Economic Journal: Macroeconomics, 2:43–69, 2010.

J. S. Dagpunar. An easily implemented generalised inverse Gaussian generator. *Communications in Statistics - Simulation and Computation*, 18(2):703–710, 1989.

S. Frühwirth-Schnatter and H. Wagner. Stochastic model specification search for Gaussian and partial non-Gaussian state space models. *Journal of Econometrics*, 154:85–100, 2010.

J. Galí, J. Vallés and J. D. López-Salido. Understanding the effects of government spending on consumption. *Journal of the European Economic Association* March 5(1):227–270, 2007.

E. I. George and R. McCulloch. Variable Selection via Gibbs Sampling. *Journal of the American Statistical Association*, 88: 881–889, 1993.

E. I. George and R. McCulloch. Approaches for Bayesian variable selection. *Statistica Sinica*, 7:339–373, 1997.

S. Kim, N. Shephard and S. Chib. Stochastic volatility: likelihood inference and comparison with ARCH models. *The Review of Economic Studies*, 65(3): 361–393, 1998.

S. K. Kinney and D. B. Dunson. Fixed and random effects selection in linear and logistic models. *Biometrics*, 63(3): 690–698, 2007.

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

G. Koop and D. Korobilis. Bayesian multivariate time series methods for empirical macroeconomics. *Foundations and Trends in Econometrics*, 3(4): 267–358, 2010.

G. Koop and D. Korobilis. Large time-varying parameter VARs. *Journal of Econometrics*, 2013. DOI: 10.1016/j.jeconom.

G. Koop and S. M. Potter. Time varying VARs with inequality restrictions. *Journal of Economic Dynamics and Control*, 35:1126–1138, 2011.

G. Koop, R. Leon-Gonzalez, and R.W. Strachan. On the evolution of the monetary policy transmission mechanism. *Journal of Economic Dynamics and Control*, 33(4):997–1017, 2011.

D. Korobilis. VAR forecasting using Bayesian variable selection. *Journal of Applied Econometrics*, 2011. DOI: 10.1002/jae.1271.

D. P. Kroese and J. C. C. Chan. *Statistical Modeling and Computation*, Springer, New York, 2014.

J. S. Liu and C. Sabatti. Generalised Gibbs sampler and multigrid Monte Carlo for Bayesian computation. *Biometrika*, 87(2):353–369, 2000.

- J. S. Liu and Y. N. Wu. Parameter expansion for data augmentation. *Journal of the American Statistical Association*, 94(448): 1264–1274, 1999.
- J. Nakajima and M. West. Bayesian analysis of latent threshold dynamic models. *Journal of Business and Economic Statistics*, 31(2):151-164, 2013.
- R. Perotti. Estimating the effects of fiscal policy in OECD countries. *Proceedings, Federal Reserve Bank of San Francisco*, 2005.
- G. E. Primiceri. Time varying structural vector autoregressions and monetary policy. *Review of Economic Studies*, 72(3):821–852, 2005.
- G. A. F. Seber. *Linear regression analysis*. New York: Wiley, 1977.