



Contents lists available at ScienceDirect

Journal of Econometrics

journal homepage: www.elsevier.com/locate/jeconomReducing the state space dimension in a large TVP-VAR[☆]Joshua C.C. Chan^{a,b}, Eric Eisenstat^c, Rodney W. Strachan^{c,*}^a Economics Department, Purdue University, Krannert Building, 403 W. State Street, West Lafayette, IN 47907-2056, USA^b Economics Discipline Group, University of Technology Sydney, Building 8, Dr. Chau Chak Building, 14-28 Ultimo Road, Ultimo NSW 2007, Australia^c School of Economics, University of Queensland, St Lucia QLD 4072, Australia

ARTICLE INFO

Article history:

Received 17 March 2018

Received in revised form 7 November 2019

Accepted 17 November 2019

Available online xxxx

JEL classification:

C11

C22

E31

Keywords:

Large VAR

Time varying parameter

Reduced rank covariance matrix

ABSTRACT

This paper proposes a new approach to estimating high dimensional time varying parameter structural vector autoregressive models (TVP-SVARs) by taking advantage of an empirical feature of TVP-(S)VARs. TVP-(S)VAR models are rarely used with more than 4–5 variables. However recent work has shown the advantages of modelling VARs with large numbers of variables and interest has naturally increased in modelling large dimensional TVP-VARs. A feature that has not yet been utilized is that the covariance matrix for the state equation, when estimated freely, is often near singular. We propose a specification that uses this singularity to develop a factor-like structure to estimate a TVP-SVAR for many variables. Using a generalization of the recenting approach, a rank reduced state covariance matrix and judicious parameter expansions, we obtain efficient and simple computation of a high dimensional TVP-SVAR. An advantage of our approach is that we retain a formal inferential framework such that we can propose formal inference on impulse responses, variance decompositions and, important for our model, the rank of the state equation covariance matrix. In a system with 15 variables, we show clear empirical evidence in favour of our model and improvements in estimates of impulse responses.

© 2020 Elsevier B.V. All rights reserved.

1. Introduction

Vector autoregressive models (VARs) have provided many valuable insights in applied macroeconometrics. The past decade has seen considerable interest in VARs with parameters that evolve over time – time varying parameter VARs (TVP-VARs) – particularly with heteroscedasticity, to better capture the evolving dynamics of the underlying variables. More recently researchers have been developing methods to estimate larger systems of variables in VARs to avoid limitations that arise when too few variables are modelled. The problems that motivate using both TVP-VARs and large VARs are compelling, but addressing both problems in one model leads to significant computational challenges. This paper proposes an approach to address these challenges.

Bañbura et al. (2010) argue for modelling many variables in a large VAR to avoid a number of problems that arise from modelling too few variables. They (and other authors such as Carriero et al. (2011), Giannone et al. (2014), Koop

[☆] This research was supported by funding from the Australian Research Council project DP180102373. We thank Sylvia Frühwirth-Schnatter, Roberto León-González, Silvia Kaufmann, Gregor Kastner, Gary Koop, Dimitris Korobilis, and Justin Tobias for helpful discussions. We also thank the editor and two anonymous referees for valuable comments that helped improve the paper. All errors are strictly our own.

* Corresponding author.

E-mail address: r.strachan@uq.edu.au (R.W. Strachan).

¹ Rodney Strachan is a Fellow of the Rimini Centre for Economic Analysis and the Centre for Applied Macroeconomic Analysis.

(2013) and Koop and Korobilis (2013)), point out that forecasts, policy advice and analysis of structure suffer problems resulting from omitted variable bias from using too few variables in the VAR. Typical sample sizes in the VAR literature, however, are not large and using large VARs leads to significant parameter proliferation making estimation and more general inference either difficult or infeasible. Bańbura et al. (2010) address this problem by employing the so-called Litterman prior to impose sufficient shrinkage to permit inference.

The time varying parameter vector autoregressive model (TVP-VAR) allows for the processes generating macroeconomic variables to evolve over time. These models, which are most commonly given a state space representation, have informed us on a range of questions of interest to policymakers with perhaps the most notable area of application being on the transmission of monetary policy (see, for example, Cogley and Sargent (2001, 2005), Primiceri (2005), and Koop et al. (2009)). Although the number of variables modelled using TVP-VARs has tended not to be very large, the arguments for using large VARs have quite naturally led to efforts to develop large TVP-VARs. As the number of states grows polynomially in the number of variables and time then, as in the large VARs, computational difficulties are encountered in these models when there are many variables. These difficulties tended to limit the number of variables modelled using the TVP-VAR.²

A range of approaches have been developed to estimating large TVP-VAR models to address the computational challenges. Some approaches aim to achieve parsimony by shrinking parameters towards zero (e.g., Belmonte et al. (2014)), others use time-varying FAVAR to incorporate the information from many sources. Since Bernanke et al. (2005), there has been a growing body of research on the use of factor-augmented VAR models as an alternative specification to TVP-VAR to modelling and forecasting with large systems. The factors and observables are given a VAR structure and, typically in these models, only a few factors drive the common variation in the system. More recently, these models have been extended to permit time-varying parameters. An interesting recent extension is Eickmeier et al. (2015) in which the states are observed and their dynamics are allowed to evolve over time. Such an approach allows a large amount of information to feed into the model and this information can also influence the variance. Eickmeier et al. (2015), for example, allow the factors to drive the volatility process as well as the mean. Unlike the factor structure, the approach in this paper preserves the original VAR structure and permits a full covariance matrix for all state shocks.

Koop and Korobilis (2013) (hereafter KK2013)) present an approach to estimating large TVP-VARs by altering the specification and reporting estimates that avoid an expensive exploration of the full posterior distribution. Using forgetting factors they replace the state equation covariance matrix with a matrix proportional to a filtered estimate of the posterior covariance matrix. They avoid running an MCMC investigation of the full posterior distribution by reporting estimates from the Kalman filter. As these estimates only use data up to the time of the forecast, t , and avoid a full MCMC investigation of the posterior, the approach in KK2013 is ideally suited to forecasting as demonstrated in that paper. The filter is appropriate as this approximates the information available to agents producing the forecasts at that time.³ Using the new specification with the filter delivers impressive computational speed in estimating a TVP-VAR for the purpose of forecasting.

Our paper, by contrast, undertakes an ex-post study of economic behaviour which is a very different purpose to replicating historical forecasts. The Kalman filter is not appropriate in this case (see Sims (2001) critique of Cogley and Sargent (2001) along these lines) as learning about the model at time t happens both before and after t . The Kalman smoother is better suited to our purpose as the resulting estimates use all available data. The model in this paper differs in that we preserve the full probability model and the dimension reducing restrictions imposed are suggested by empirical evidence. Coupling this specification with an MCMC approach to estimation, we are able to explore the full posterior distribution and so permit the full range of formal inferential opportunities.⁴

An issue that has been bubbling away in the background in the literature on TVP-VARs is the treatment of the state equation covariance matrix. This matrix is often specified as diagonal, although there is good reason to specify this as a full matrix. Primiceri (2005, p. 830) provides an argument that a full covariance matrix for the vector of all mean equation and structural parameter states would be most appropriate as the states are, and are expected to be, highly correlated. However, he does not adopt such a specification in order to avoid parameter proliferation and the attendant computational issues. Primiceri (2005) does maintain a full covariance matrix for the reduced form mean equation states and more papers are doing so (see for example, Eisenstat et al. (2016)). A full state equation covariance matrix poses significant computational challenges for large TVP-VARs. As the number of variables n grows, the number of mean parameters grows at order n^2 and the number of parameters in the state equation covariance matrix grows at n^4 . In this paper we present an reduced rank restriction on the state equation covariance matrix that results in a reduced number of state errors driving the time-varying parameters.

The first contribution of this paper is to present an alternative approach to estimating large TVP-VARs. We increase the number of variables we can model in a TVP-VAR by taking advantage of the strong correlations among the states. We

² A few papers, such as Carriero et al. (2016a,b) and Chan (2020), have developed large VARs with stochastic volatility. But these papers all restrict the VAR coefficients to be constant.

³ This does not account for the fact that the data are revised. However real-time forecasts could be readily incorporated.

⁴ There are a number of other differences with KK2013. In that paper they use a deterministic model for the evolution of the measurement error covariance matrix (an exponentially weighted moving average specification) whereas we use a stochastic specification. Further, KK2013 propose a model for a time-varying parameter VAR that allows the dimension of the model to change over time and this feature is shown to produce improvements in forecasting.

preserve the exact state space model but achieve parsimony by imposing a restriction suggested by the data; that the state equation covariance matrix has reduced rank. An early observation by [Cogley and Sargent \(2005\)](#) shows, using principal component analysis, that the posterior estimate of the covariance matrix for the state equation appears to have a very low rank. We formalize this observation into a model specification. [Primiceri \(2005\)](#) points out that small state equation error variances cause problems for frequentist computation. Our approach, by contrast, uses this feature to improve Bayesian estimation.

While the usual TVP-VAR will have the same number of time-varying parameters as states (and the terms *time-varying parameters* and *states* are interchangeable in this case), in our model the time-varying parameters are driven by a much smaller number of states (and the distinction between *time-varying parameters* and *states* becomes important). Note that no matter by how much we reduce the number of states, the number of time-varying parameters in the VAR does not change. This statement will become clearer when we develop the model in Section 2. Reducing the number of states driving the time varying parameters results in estimates of the time varying parameters that are far more precise.

We employ a range of strategies, in addition to the reduced rank structure, to mitigate the computational issues. Each makes a small contribution on their own, but collectively they allow us to estimate larger models. First, by estimating the structural form of the TVP-VAR directly, we remove one sampling step in the Gibbs sampler. This is particularly important as estimating the reduced form TVP-VAR involves drawing two blocks of parameters that are naturally highly correlated. We collapse these two blocks into one and draw that block in one step. Next, to achieve a readily computable specification we generalize the scalar non-centred specification of the state space model by [Frühwirth-Schnatter and Wagner \(2010\)](#) to the matrix non-centred specification. This removes another step from the sampler as we draw the initial states and the state covariance matrix together in a single step. Further, we avoid the Kalman filter and smoother and, instead, use the precision sampler of [Chan and Jeliazkov \(2009\)](#). This precision sampler uses a lower order of computations to draw from the same posterior as the Kalman smoother. [McCausland et al. \(2011\)](#) provide a useful discussion on the computational advantages of this approach and, in particular, point out that the gain is more significant for larger n . The main contributions to improving computational speed, however, are due to the dimension reduction (over 95%) that comes from the rank reduction of the state covariance matrix and the use of the precision sampler rather than the more computationally intensive Kalman filter and smoother.

The specification of the reduced rank model requires semi-orthogonal matrices and ordered positive elements. This specification induces nonstandard supports for the parameters and Bayesian computation on such supports is difficult. Another contribution of this paper, then, is to use a judicious selection of parameter expansions and priors for the expanding parameters to develop a specification that is fast, efficient and easy to compute. This expansion is part of the generalization of the recentering method of [Frühwirth-Schnatter and Wagner \(2010\)](#) to a multivariate setting mentioned above.

We apply the new specification and computation techniques to a study of the evolution of responses of a range of real and nominal macroeconomic variables to surprise productivity (non-news) and news shocks. Increasing the number of variables modelled in a TVP-VAR could also prove useful in many other settings. In a study of network spillovers among financial institutions, [Geraci and Gnabo \(2018\)](#) demonstrate the utility of TVP-VAR models for a system of four sectors. The approach in this paper permits the analysis of many more sectors or possibly disaggregated data. Similarly, [Ciccarelli and Rebucci \(2007\)](#) propose using a TVP-VAR to address simultaneity issues in studying contagion and interdependence among exchange rates. The approach could find uses outside of economics. In psychology, for example, the TVP-VAR is used to model emotion dynamics and has been proposed for the study of networks in psychopathology (see [Bringmann et al. \(2018\)](#) and references therein).

The structure of the paper is as follows. In Section 2 we present the idea with a general state space model. We outline the model specifications that result from different assumptions about the rank of the state equation covariance matrix. This section also contains a technical derivation of the reduced sources of errors model that results from a reduced rank state equation covariance matrix. In Section 3 we outline posterior computation. Section 4 presents an application using a TVP-VAR with 15 variables to demonstrate the proposed methodology. Section 5 concludes and gives some indication of directions for future research.

2. Reducing the sources of variation

2.1. Overview

We will apply the reduced sources of error approach to a structural form TVP-VAR (TVP-SVAR). In VAR analysis, the measurement equation is often specified on the reduced form parameters, although we can readily transform between the reduced form and structural form. We prefer the structural form as it reduces the number of blocks of parameters to be estimated and makes the dependence among the structural and reduced form parameters simpler (i.e., linear).

For the $n \times 1$ vector y_t for $t = 1, \dots, T$, the TVP-SVAR can be written as

$$B_{0,t}y_t = \mu_t + B_{1,t}y_{t-1} + \dots + B_{p,t}y_{t-p} + \varepsilon_t, \quad \varepsilon_t \sim N(0, \Sigma_t), \quad (1)$$

where $B_{0,t}, \dots, B_{p,t}$ are $n \times n$ and $\Sigma_t = \text{diag}(\exp(h_{1,t}), \dots, \exp(h_{n,t}))$. The first matrix $B_{0,t}$ is $n \times n$ with ones on the diagonal and is commonly specified as lower triangular.

Given the structure of $B_{0,t}$, we may write $B_{0,t} = I - B_t$ so that the matrix B_t has zeros on the diagonal. The TVP-SVAR can now be written as:

$$\begin{aligned} y_t &= \mu_t + B_t y_t + B_{1,t} y_{t-1} + \cdots + B_{p,t} y_{t-p} + \varepsilon_t \\ &= \mu_t + (y'_t \otimes I_n) D b_t + (y'_{t-1} \otimes I_n) b_{1,t} + \cdots + (y'_{t-p} \otimes I_n) b_{p,t} + \varepsilon_t, \end{aligned}$$

where $b_{l,t} = \text{vec}(B_{l,t})$, $l = 1, \dots, p$ and $D b_t = \text{vec}(B_t)$ where b_t contains all the $\frac{n(n-1)}{2}$ non-zero elements of B_t in a vector and D is an appropriately defined $n^2 \times \frac{n(n-1)}{2}$ selection matrix. If we define the $n \times k$ matrix

$$x_t = [I_n \quad (y'_t \otimes I_n) D \quad (y'_{t-1} \otimes I_n) \quad \cdots \quad (y'_{t-p} \otimes I_n)]$$

such that $k = (np + 1 + \frac{n-1}{2})n$ and the $(k \times 1)$ vector $\alpha_t = (\mu'_t \quad b'_t \quad b'_{1,t} \quad \cdots \quad b'_{p,t})'$, we can write the above model using a standard but reasonably general specification of the state space model for an observed $n \times 1$ vector of observations y_t with $n \times k$ matrix of regressors x_t :

$$y_t = x_t \alpha_t + \varepsilon_t, \quad \varepsilon_t \sim N(0, \Sigma_t), \quad (2)$$

$$\alpha_t = \alpha_{t-1} + \eta_t, \quad \eta_t \sim N(0, Q_\alpha), \quad \alpha_0 = \alpha \sim N(\alpha, V). \quad (3)$$

In the application in Section 4, it will be more convenient to transform from the VAR to the VECM form, but this again can be written in the general form in (2) and (3). We therefore continue with the general form of the model and delay giving specific details on the prior we use until Section 4. We can now present the idea of reducing the sources of errors in a general linear Gaussian state space model.

We have not imposed any restrictions on the above model at this point and all of the parameters in the VAR are able to vary over time. The dimension reduction occurs by applying a rank reduction to the covariance matrix for the state equation, Q_α . If we set the rank of Q_α to $r_\alpha = \text{rank}(Q_\alpha) \leq k$, then after applying the appropriate transformations (detailed in the next subsection below) we can write the model in (2) and (3) as follows:

$$y_t = x_t \alpha + x_t A_\alpha f_{\alpha,t} + \varepsilon_t, \quad \varepsilon_t \sim N(0, \Sigma_t), \quad (4)$$

$$f_{\alpha,t} = f_{\alpha,t-1} + z_{\alpha,t}, \quad z_{\alpha,t} \sim N(0, I_{r_\alpha}), \quad f_{\alpha,0} = 0, \quad (5)$$

where A_α is a $(k \times r_\alpha)$ matrix, $f_{\alpha,t}$ and $z_{\alpha,t}$ are $(r_\alpha \times 1)$ vectors and the errors ε_t and z_t are independent of one another. As r_α is generally much smaller than k , we call the model in (4) and (5) the reduced sources of error model.

The technical details on the link between the general form of the state space model in (2) and (3) and the final form in (4) and (5), including centring and parameter expansions, are presented in the following subsection for the interested reader. There are a number of choices in modelling the state space model and the correlation structure. In this paper we extend the above to reducing the rank of the covariance matrix for the volatility states, Q_h . We present two specifications, the second encompasses the first but there are significant differences in computation between the two specifications.

In the transformation from (2) to (3) we use $\alpha_t = \alpha + A_\alpha f_{\alpha,t}$ where $A_\alpha A'_\alpha = Q_\alpha$. This function implies that the k time varying parameters in α_t are driven by $r_\alpha \leq k$ states, $f_{\alpha,t}$, in a factor-like structure for the states. The elements of A_α and $f_{\alpha,t}$ are not identified and this results from the use of parameter expansions. These expansions relax the form of the model to improve estimation. In fact, we derive the above form starting from identified parameters but then introduce the parameter expansions that take away this identification.

To give an impression of the extent of dimension reduction that is typically achieved, consider our empirical application. We have $n = 15$ variables and $T = 250$ observations for a VAR with 2 lags. The dimension of the states α_t and the covariance matrix Q_α in the unrestricted model in (2) and (3) has dimension 305,805 (here we ignore the volatility). With rank of Q_α set to $r_\alpha = 4$, which is preferred in this application, then Specification 1 in (4) and (5) has dimension 3,844 representing a 98.7% reduction in model dimension. It is worth noting that the larger is n the greater is the dimension reduction.

Recall that with the full covariance matrix Q_α the dimension of this matrix grows at rate n^4 . Instead of using the specification of the state space model with a full covariance matrix Q_α , one might therefore use a diagonal specification of Q_α in the hope of reducing the dimension of the model. However, this does not result in as great a dimension reduction as using a reduced rank Q_α . In the case considered in our application, for example, the states α_t and the diagonal covariance matrix Q_α in the unrestricted model in (2) and (3) would have dimension 143,640. Thus the model in (4) and (5) with rank of Q_α being $r_\alpha = 4$ still has a dimension 97.3% smaller than if a diagonal form were chosen for Q_α .

2.2. Mapping to the reduced sources of errors model

In this subsection, we present the details of the transformations from (2) and (3) to (4) and (5). Important features of the transformed model are that there are no unknown parameters in the state equations and that the parameters to be estimated all appear in the mean equation. Further, all of the parameters in α , A_α and $f_{\alpha,t}$ have conditionally normal posteriors.

Frühwirth-Schnatter and Wagner (2010) develop a computationally efficient specification of the state space model that permits the time variation in individual parameters to be 'turned off'. This approach involves two transformations:

recentering (or non-centring) and parameter expansion. We leave for a subsequent paper consideration of turning off time variation. Rather we use the non-centred specification to develop a reduced rank model from which it is simpler to obtain draws of the parameters.

In recentering, the initial value is subtracted from all states and this is divided by the standard deviation of the state equation error. This transformation moves the initial state and the standard deviation into the mean equation leaving no unknown parameters in the state equation.

The Frühwirth-Schnatter and Wagner (2010) approach is developed for scalar or independent states. That is, Q_α is assumed to be scalar or a diagonal matrix. In our model the covariance matrix Q_α is a full symmetric matrix allowing correlation among the elements of η_t . We denote the initial state by α . Generalizing to this case, the recentering transforms from α_t to $\tilde{\alpha}_t$ via

$$\alpha_t = \alpha + Q_\alpha^{1/2} \tilde{\alpha}_t, \quad (6)$$

and the model subsequently becomes

$$y_t = x_t \alpha + x_t Q_\alpha^{1/2} \tilde{\alpha}_t + \varepsilon_t, \quad \varepsilon_t \sim N(0, \Sigma_t), \quad (7)$$

$$\tilde{\alpha}_t = \tilde{\alpha}_{t-1} + \tilde{z}_t, \quad \tilde{z}_t \sim N(0, I_k), \quad \tilde{\alpha}_0 = 0. \quad (8)$$

This more general specification requires a useful definition for $Q_\alpha^{1/2}$, the square root of the covariance matrix Q_α . There are several ways to define the square root of a full symmetric matrix, but for our purposes the definition must allow for Q_α to have reduced rank. Our preferred definition, which can readily accommodate rank reduction, uses the singular value decomposition.

The singular value decomposition of Q_α can be written as $Q_\alpha = U \Lambda U'$ where $\Lambda = \text{diag}\{\lambda_1, \lambda_2, \dots, \lambda_k\}$, $\lambda_i \geq \lambda_{i+1} \geq 0$ and $U \in O(k) \equiv \{U(k \times k) : U'U = I_k\}$ is an orthonormal matrix: $U'U = I_k$. Given Q_α , the elements of U are identified up to sign (which is trivially resolved). The matrix $Q_\alpha^{1/2}$ is defined simply as $Q_\alpha^{1/2} = U \Lambda^{1/2} U'$. In this paper we impose parsimony by letting the $k - r_\alpha$ smallest singular values of Λ to be zero. That is, we allow $\lambda_{r_\alpha+1} = \lambda_{r_\alpha+2} = \dots = \lambda_{k-1} = \lambda_k = 0$ and collect the nonzero singular values into $\Lambda_1 = \text{diag}\{\lambda_1, \lambda_2, \dots, \lambda_{r_\alpha}\}$. In this case, we can conformably decompose $U = [U_1 \ U_2]$ such that $U_1 \in V_{r_\alpha, k} \equiv \{U(k \times r_\alpha) : U'U = I_{r_\alpha}\}$ and $U_1' U_2 = 0$ an $r_\alpha \times (k - r_\alpha)$ matrix of zeros. Under this restriction

$$Q_\alpha^{1/2} = U \Lambda^{1/2} U' = U_1 \Lambda_1^{1/2} U_1'.$$

We introduce the square root of the reduced rank covariance matrix into the specification (6) to obtain the expression

$$\begin{aligned} \alpha_t &= \alpha + U_1 \Lambda_1^{1/2} U_1' \tilde{\alpha}_t \\ &= \alpha + U_1 \Lambda_1^{1/2} f_{-t}, \end{aligned} \quad (9)$$

where in the second line we have taken the linear combination $f_{-t} = U_1' \tilde{\alpha}_t$. The rank reduction implies a reduction in the number of states from k (in α_t) to r_α (in f_{-t}). Taking the linear combination f_{-t} in the state equation implies also taking the linear combinations of the $z_t = U_1' \tilde{z}_t$. Here we have used the result that a linear combination of standard normal random variables (\tilde{z}_t) in which the linear combinations are formed using a set of orthogonal unit vectors (U_1 in our case) results in a vector of standard normal variables (z_t). Thus the resulting state equation vector of errors, z_t , is an r_α -vector of standard normal variables. That is, the state equation is now

$$f_{-t} = f_{-t-1} + z_t, \quad z_t \sim N(0, I_{r_\alpha}), \quad f_{-0} = 0.$$

The specification in (9) involves the parameters U_1 and Λ_1 which have nonstandard supports. These nonstandard supports significantly complicate computation and it is difficult to obtain an efficient and simple algorithm. This issue is addressed by mapping to a less restrictive form by introducing unidentified parameters.

The second step in the approach of Frühwirth-Schnatter and Wagner (2010) is to introduce an unidentified parameter via an approach called *parameter expansion*, to map the parameters to more standard forms and supports. Used judiciously, transformation via parameter expansion can make computation much simpler and more efficient. This is achieved by the mapping to standard supports and employing standard distributions thereby simplifying computation and breaking down the dependency in the parameters (see discussion in, for example, Liu et al. (1998) and Liu and Wu (1999)). Importantly, this approach has proven useful in reduced rank models such as cointegrating vector error correction models (see Koop et al. (2010)), factor models (Chan et al., 2018), and simultaneous equations models (Koop et al., 2012).

Working in the scalar case, Frühwirth-Schnatter and Wagner (2010) introduce an indicator ι that randomly takes the values -1 or $+1$. The support for ι is therefore a one-dimensional orthogonal group, $O(1)$. Generalizing this, we expand the set of parameters by introducing the orthonormal matrix $C \in O(r_\alpha)$ where $O(r_\alpha)$ is the r_α -dimensional orthogonal group. Define the matrix $A_\alpha = U_1 \Lambda_1^{1/2} C'$. Note that the definition of A_α is just a standard singular value decomposition of a real matrix with singular values on the diagonal of $\Lambda_1^{1/2}$. Introducing this expanding parameter C into the model through (9) we obtain

$$\begin{aligned} \alpha_t &= \alpha + U_1 \Lambda_1^{1/2} C' C f_{-t} \\ &= \alpha + A_\alpha f_{\alpha, t}, \\ f_{\alpha, t} &= f_{\alpha, t-1} + z_{\alpha, t}, \quad z_{\alpha, t} \sim N(0, I_{r_\alpha}), \quad f_{\alpha, 0} = 0 \end{aligned}$$

in which $f_{\alpha,t} = C f_t$ and $z_{\alpha,t} = C z_t$. Introducing the above transformation into the measurement equation in (7) and replacing the state equation in (8) by the one above, we obtain the final form of the full state space model as that given in (4) and (5).

2.3. Two specifications for the variance

The standard model assumed in the literature specifies α_t and $h_t = (h_{1,t}, \dots, h_{n,t})'$ as *a priori* independent and that the covariance matrix in the state equation for h_t is full rank. For example, a standard specification is a random walk log-volatility

$$h_t = h_{t-1} + \eta_{h,t} \quad \eta_{h,t} \sim N(0, Q_h)$$

where $Q_h = \text{diag}(\sigma_{h1}^2, \dots, \sigma_{hn}^2)$ and the random walk is initialized with h_0 .

In this section we apply the dimension reduction to the log variances, h_t , in (1). That is, we generalize to permit Q_h to be a full, possibly reduced rank symmetric matrix. Much of the parameter proliferation in the TVP-SVAR occurs in the mean equations but we could just as reasonably wish to reduce the number of states driving the stochastic volatility. The volatility component of the models we propose here resembles that of [Carriero et al. \(2016a\)](#). Expanding upon the specification in Section 2.1, we consider two specifications of the log volatility h_t for reducing the dimensions of the TVP-SVAR. The first, Specification 1, assumes the mean equation and volatilities share common states while Specification 2 specifies them to be *a priori* independent. The rationale for the first specification is that structural change in the mean and variance could come from a common source. That is, structural change is driven by a common factor. Specification 2 adopts the more standard assumption that the mean and variance states are independent. Specification 1 of the process for α_t and h_t encompasses Specification 2.

It is not difficult to imagine that shocks can drive changes in the whole structure of the model such that changes in the mean and variance parameters are driven by the same states. To allow for this possibility, we allow for the mean equation and volatility to influence each other in the most general model specification. In this model, the most general form, we allow the mean equation states, α_t , to be correlated with the log volatilities in h_t . To permit this we specify a state equation for α_t and h_t jointly as:

$$\theta_t = \begin{pmatrix} \alpha_t \\ h_t \end{pmatrix}.$$

Specification 1 has state equation

$$\theta_t = \theta_{t-1} + \eta_{\theta,t}, \quad \eta_{\theta,t} \sim N(0, Q_{\theta}), \quad (10)$$

such that the mean and variance states are correlated. After applying the rank reduction to the above specification, the time varying parameters in the model are

$$\begin{aligned} \theta_t &= \theta + A f_{\theta,t}, & A &= \begin{pmatrix} A_{\alpha} \\ A_h \end{pmatrix}, \\ f_{\theta,t} &= f_{\theta,t-1} + z_t, & z_t &\sim N(0, I_r), \quad f_{\theta,0} = 0, \end{aligned}$$

where $r = r_{\alpha} + r_h$, A is $(n+k) \times r$ and $f_{\theta,t}$ is $r \times 1$.

It is more common to impose, usually for computational convenience, that the errors in the state equations for α_t and h_t are independent. However, we wish to retain dependence among the volatilities. The second model, Specification 2, assumes that α_t and h_t are independent such that

$$A = \begin{pmatrix} A_{\alpha} \\ A_h \end{pmatrix} = \begin{pmatrix} A_{\alpha,11} & 0 \\ 0 & A_{h,12} \end{pmatrix}.$$

In this case, we could rewrite the model for h_t as

$$\begin{aligned} h_t &= h + A_{h,11} f_{h,t}, \\ f_{h,t} &= f_{h,t-1} + z_{h,t}, \quad z_{h,t} \sim N(0, I_{r_h}), \quad z_{h,0} = 0, \end{aligned}$$

where $A_{h,11}$ is $n \times r_h$, $f_{h,t}$ is $r_h \times 1$ and, as we might reasonably expect that the volatilities can be modelled with common factors, then $r_h \leq n$.

3. Posterior estimation

The state space structure specifies the priors for the states – $f_{\alpha,t}$, $f_{h,t}$ and $f_{\theta,t}$ – so we now describe the priors for the initial conditions $\theta = (\alpha', h')'$ and covariance matrices $a = \text{vec}(A)$.

[Frühwirth-Schnatter and Wagner \(2010\)](#) provide evidence in support of using the Gamma prior, rather than the inverted Gamma prior, for their scalar state equation variance. In the generalization presented in this paper, this equates to using a Wishart prior for Q_{θ} . For the full rank ($r = k$) case, a zero mean normal prior for A implies a Wishart prior for

Q_θ (see, for example, Zellner (1971, pp. 389–392) and Muirhead (1982)). We therefore give the matrix A a normal prior distribution, $a = \text{vec}(A) \sim N(0, cI_{(n+k)r})$ for all three specifications. Through some experimentation, we find $c = 10^{-3}$ to be reasonable in a wide variety of settings.

For the initial state $\theta = \{\theta_j\}$ (which contains the elements of α and h), we note that in large models the dimension may be substantial, and hence, shrinkage priors may be desirable. This implies a choice of structure on the prior covariance matrix $\underline{V}_\theta = \{\underline{V}_{\theta_j}\}$. A number of options explored in the large Bayesian VAR literature may be applied here. We consider the stochastic search variable selection (SSVS) (e.g., George and McCulloch, 1993) prior of the form:

$$\begin{aligned}\theta_j | \delta_j &\sim \mathcal{N}(\underline{\theta}_j, c_{\delta_j} \underline{V}_{\theta_j}), \\ \delta_j &\sim q^{\delta_j} (1 - q)^{1 - \delta_j},\end{aligned}\tag{11}$$

where $\delta_j \in \{0, 1\}$, $c_1 = 1$ and c_0 is some small constant. Of course, this will collapse to a typical normal prior if either $q = 1$ or $c_0 = 1$. Further, we combine SSVS with Minnesota priors as suggested in Korobilis (2013). Having normal conjugate priors for the initial conditions (α, h), the covariances ($a = \text{vec}(A)$) and the states (the f_t), the resulting conditional posteriors are normal for Specifications 1 and 2.

For the purposes of this section, we collect the T states into the vectors $f_m = (f'_{m,1}, f'_{m,2}, \dots, f'_{m,T})'$ for $m = \alpha, h$ or θ . Further, let $a_\alpha = (\text{vec}(\alpha)', \text{vec}(A_\alpha)')'$ and $a_h = (h'_0, \text{vec}(A_h)')$. The description of the priors above implies that the vectors a_α, f_α, a_h and f_h have a normal form such as $N(\underline{\mu}_m, \underline{V}_m)$ for $\mu = a$ or f . Volatility Specification 2 leads to a straightforward sampler. For Specification 2, MCMC involves five blocks:

1. $(a_\alpha | s_h, f_\alpha, a_h, f_h, y) \sim \mathcal{N}(\bar{a}_\alpha, \bar{V}_\alpha);$
2. $(f_\alpha | s_h, a_\alpha, a_h, f_h, y) \sim \mathcal{N}(\bar{f}_\alpha, \bar{V}_{f,\alpha});$
3. $(s_h | a_\alpha, f_\alpha, a_h, f_h, y);$
4. $(a_h | s_h, a_\alpha, f_\alpha, f_h, y) \sim \mathcal{N}(\bar{a}_h, \bar{V}_h);$
5. $(f_h | s_h, a_\alpha, f_\alpha, a_h, y) \sim \mathcal{N}(\bar{f}_h, \bar{V}_{f,h});$

of which Steps 1, 2, 4 and 5 involve only analytically tractable conditional distributions, all of which are straightforward to sample from. The states, s_h , drawn in Step 3 are the states determining the normal mixture components when drawing the stochastic volatilities using the algorithm of Kim et al. (1998).

For Specification 1 the MCMC consists of sampling recursively from:

1. $(a_\alpha | f_\theta, h_0, A_h, y) \sim \mathcal{N}(\bar{a}_\alpha, \bar{V}_\alpha);$
2. $(f_\theta | a_\alpha, h_0, A_h, y);$
3. $(a_h | a_\alpha, f_\theta, y) \sim \mathcal{N}(\bar{a}_h, \bar{V}_\alpha).$

Under this specification, the measurement equation is nonlinear in f_θ (since it enters both the conditional mean and the volatility simultaneously), and therefore, $(f_\theta | \alpha, A_\alpha, h_0, A_h, y)$ is not analytically tractable. We therefore sample it using an accept–reject Metropolis–Hastings (ARMH) algorithm as described in Chan and Strachan (2012). Specifically, we use a normal proposal centred on the conditional posterior mode \hat{f}_θ with the variance \hat{V}_θ set to the negative inverse Hessian evaluated at the mode of $\ln p(f_\theta | \cdot, y)$. The derivation of \hat{f}_θ and \hat{V}_θ is given in Section 1 of the Online Appendix.

Once the mode \hat{f}_θ is obtained, the proposal precision \hat{V}_θ^{-1} is given by a by-product of the scoring algorithm and a matrix that can be easily evaluated at the mode upon convergence (See Online Appendix, Section 1). We then generate proposals as $f_\theta^c \sim \mathcal{N}(\hat{f}_\theta, \hat{V}_\theta)$ for the ARMH step as detailed in Chan and Strachan (2012). The use of ARMH as opposed to standard M–H appears to provide substantial gains in terms of acceptance rates (and hence sampling efficiency), particularly for larger models (i.e. as the size of f_θ increases). Intuitively, the normal proposal is symmetric, while $p(f_\theta | \cdot, y)$ will typically be skewed. This mismatch in shape will lead to higher rejection rates for a standard M–H approach as the dimension of f_θ increases. ARMH mitigates this by adjusting the shape of the proposal to better fit the skewness of the target distribution. As a result, acceptance rates are substantially increased. For example, in the macroeconomic application discussed below, the model with $n = 15$ and $r_\alpha = 10$ yields an acceptance rate of about 89.9%.

4. Application

4.1. Implementation

We use a data set containing a total of 15 variable to estimate the time-varying effects of surprise productivity (non-news) and news shocks. To understand the effects of dimension upon the results, we estimate the model with $n = 8$ variables and again with all $n = 15$ variables for contrast. The data consists of quarterly macroeconomic series covering the period 1954Q3–2008Q3, with each variable described in Table 1.⁵ Given a subset of these variables, we assume the

⁵ Following standard practice in the news shock literature, all series are de-measured.

Table 1

Variables used in each estimated model.

Core variables		Additional variables for the $n = 15$ model	
1	Log TFP	9	Log RPI
2	FED funds rate	10	Log real SEP500
3	GDP deflator inflation	11	Unemployment Rate
4	Log hours per capita	12	Vacancy rate
5	Log real GDP per capita	13	TB3MS Spread
6	Log real consumption per capita	14	GS10 Spread
7	Log real investment per capita	15	Log real dividends
8	GS5 Spread		

system admits a structural TVP-VAR representation of the form

$$y_t = B_{0,t}^{-1} \mu_t + \Pi_{1,t} y_{t-1} + \cdots + \Pi_{p,t} y_{t-p} + A_t \tilde{\varepsilon}_t, \quad \tilde{\varepsilon}_t \sim \mathcal{N}(0, I_n), \quad (12)$$

where $A_t = B_{0,t}^{-1} \Sigma_t^{1/2}$,

$$\Sigma_t = \text{diag}(\exp(h_{1,t}), \dots, \exp(h_{n,t})) \text{ and} \\ \Sigma_t^{1/2} = \text{diag}(\exp(h_{1,t}/2), \dots, \exp(h_{n,t}/2)).$$

Following Barsky and Sims (2011), non-news and news shocks in $\tilde{\varepsilon}_t$ are identified by the restrictions:

1. *non-news* is the only shock affecting TFP on impact;
2. *news* is the shock that, among all of the remaining shocks, explains the maximum fraction of the *forecast error variance* (FEV) of TFP at a long horizon (set to 20 years in our application).

To implement the methodology outlined in the previous sections in estimating (12), we begin with the structural form in (1)

$$y_t = \mu_t + B_t y_t + B_{1,t} y_{t-1} + \cdots + B_{p,t} y_{t-p} + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \Sigma_t)$$

where $\tilde{\varepsilon}_t = \Sigma_t^{-1/2} \varepsilon_t$. To more simply apply a prior that is more useful in large models, we respecify the model in VECM form

$$\Delta y_t = \mu_t + B_t y_t + \Pi_t y_{t-1} + \Gamma_{1,t} \Delta y_{t-1} + \cdots + \Gamma_{p-1,t} \Delta y_{t-p+1} + \varepsilon_t \quad (13)$$

where B_t is the same lower triangular matrix defined in (12). Next, define

$$x_t = (I_n \quad (y'_t \otimes I_n) D \quad (y'_{t-1} \otimes I_n) \quad (\Delta y'_{t-1} \otimes I_n) \quad \cdots \quad (\Delta y'_{t-p+1} \otimes I_n))$$

such that $k = (np + 1 + \frac{n-1}{2})n$ and α_t is the $(k \times 1)$ vector

$$\alpha_t = (\mu'_t \quad b'_t \quad \pi'_t \quad \gamma'_{1,t} \quad \cdots \quad \gamma'_{p-1,t})',$$

where μ_t and b_t are defined in Section 2 and $\gamma_{l,t} = \text{vec}(\Gamma_{l,t})$ $l = 1, \dots, p-1$ and $\pi_t = \text{vec}(\Pi_t)$. Consequently, we can now write (13) in the form of (2) and (3) as:

$$\Delta y_t = x_t \alpha_t + \varepsilon_t \quad \varepsilon_t \sim N(0, \Sigma_t), \\ \alpha_t = \alpha_{t-1} + \eta_t \quad \eta_t \sim N(0, Q_\alpha) \quad \alpha = \alpha_0 \sim N(\underline{\alpha}, \underline{V}).$$

The advantage of this VECM specification is that it facilitates specifying more flexible shrinkage priors for

$$\alpha = (\mu'_0 \quad b'_0 \quad \pi'_0 \quad \gamma'_{1,0} \quad \cdots \quad \gamma'_{p-1,0})',$$

which is useful in large dimensional settings.

In addition to the SSVS specification in (11) combined with the Minnesota prior for $\theta = (\alpha', h')' = \{\theta_j\}$, we also implement “inexact differencing” as advocated by Doan et al. (1984), Baríburá et al. (2010), and others. This is done by setting the prior mean to $\underline{\theta}_j = 0$ for all j and the prior variance as

$$\underline{V}_{\theta_j} = \begin{cases} 1 & \text{if } \theta_j \in h_0, \theta_j \in \mu_0, \text{ or } \theta_j \in b_0, \\ \frac{10^2}{2n} & \text{if } \theta_j \in \pi_0, \\ \frac{0.3}{2n^2} & \text{if } \theta_j \in \gamma_{l,0} \text{ for } l = 1, \dots, p-1. \end{cases}$$

For SSVS, we set $c_0 = 0.01$ and $q = 0.5$. Finally, we scale each Δy_i to have sample standard deviation one before commencing sampling, which facilitates the use of generic prior settings like the ones given above. However, the effect

Table 2

DICs for models specified with $n = 15$ and various combinations of r_α and r_h . All values are relative to the DIC of the constant coefficient model (i.e. $r_\alpha = r_h = 0$).

3 states			5 states			7 states			10 states			12 states		
r_α	r_h	DIC	r_α	r_h	DIC	r_α	r_h	DIC	r_α	r_h	DIC	r_α	r_h	DIC
3	0	-764	5	0	-766	7	0	-742	10	0	-366	12	0	-140
2	1	-771	4	1	-816	6	1	-688	8	2	-486	8	4	-573
1	2	-711	3	2	-887	4	3	-892	6	4	-697	7	5	-655
0	3	-562	2	3	-851	3	4	-888	5	5	-854	6	6	-800
			1	4	-756	1	6	-698	4	6	-876	5	7	-792
			0	5	-583	0	7	-565	2	8	-800	4	8	-840
									0	10	-545	0	12	-577
shared		-770	shared		-835	shared		-719	shared		-418	shared		199

of this scaling is reversed in the post-processing of draws such that all outputs such as impulse response functions are reported on the originally scaled data.

Once draws of B_t , Π_t , $\Gamma_{1,t}, \dots, \Gamma_{p-1,t}$, and Σ_t are obtained, they are transformed to draws of $\Pi_{1,t}, \dots, \Pi_{p,t}$ from (12) as

$$\begin{aligned}
 B_{0,t} &= I_n - B_t \\
 \Pi_{1,t} &= I_n + B_{0,t}^{-1} (\Pi_t + \Gamma_{1,t}), \\
 \Pi_{l,t} &= B_{0,t}^{-1} (\Gamma_{l,t} - \Gamma_{l-1,t}), \quad l = 2, \dots, p-1, \\
 \Pi_{p,t} &= -B_{0,t}^{-1} \Gamma_{p-1,t}.
 \end{aligned}$$

To recover A_t , we begin with $\tilde{A}_t = B_{0,t}^{-1} \Sigma_t^{-\frac{1}{2}}$. Note that by construction \tilde{A}_t is lower triangular and therefore the non-news shock is identified in accordance with restriction 1 above. However, the news shock generally does not satisfy Restriction 2. Following Barsky and Sims (2011), the desired restriction is implemented by constructing an orthogonal matrix Q_t using a spectral decomposition of impulse response functions.

Specifically, for each period t we compute the impulse responses of log TFP to all shocks excluding non-news for the periods $t, t+1, \dots, t+80$. Let R_s be the $(n-1) \times 1$ vector of impulse responses at time $t+s$ and take the spectral decomposition

$$\tilde{Q}_t D_t' \tilde{Q}_t = \sum_{s=0}^{80} R_s R_s',$$

where the eigenvalues in D_t are in descending order. Setting

$$Q_t = \begin{bmatrix} 1 & 0 \\ 0 & \tilde{Q}_t \end{bmatrix}$$

and $A_t = \tilde{A}_t Q_t$ achieves the desired identifying restriction, which is sufficient for computing forecast error variance decompositions. To derive impulse response functions, we further identify the sign of the news shock by requiring that the maximum impact of news on log TFP across all horizons is positive.⁶

4.2. Results

We begin by conducting an extensive empirical analysis on the choice of r_α (number of mean equation states) and r_h (number of states driving the volatility) using the Deviance Information Criterion (DIC) as the model comparison criterion. The DIC is based on the integrated likelihood – i.e., the joint density of the data marginal of all the latent states – and is computed using the method in Chan and Eisenstat (2018). The relative DICs for x are presented in Table 2. For the model with $n = 15$, the DIC selects Specification 2 with seven states: four states driving the mean equation coefficients in α_t and, again, three states driving the volatilities h_t . Specification 2 is also preferred by the DIC, with five states in total, for the $n = 8$ model (full results are provided in Table 1 of the Online Appendix).

In Figs. 1 to 4 we present short and long-run impulse response functions for models with $n = 8$ and $n = 15$ variables and the number of states selected by the smallest DIC in each case. Section 3 of the Online Appendix contains supplementary results, including a range of impulse response functions and variance decompositions quantifying the effects of news and non-news shocks on economic variables.

The impact shocks in Figs. 1 and 3 show significant variation over time in the impact of non-news shocks upon log TFP. Figs. 2 and 4 show that the long run effect of a news shock on the real variables – log TFP, log real per capita GDP, log real per capita consumption and log real per capita investment – has declined over time with the density of the response moving towards zero. This effect is particularly clear for the larger model with $n = 15$.

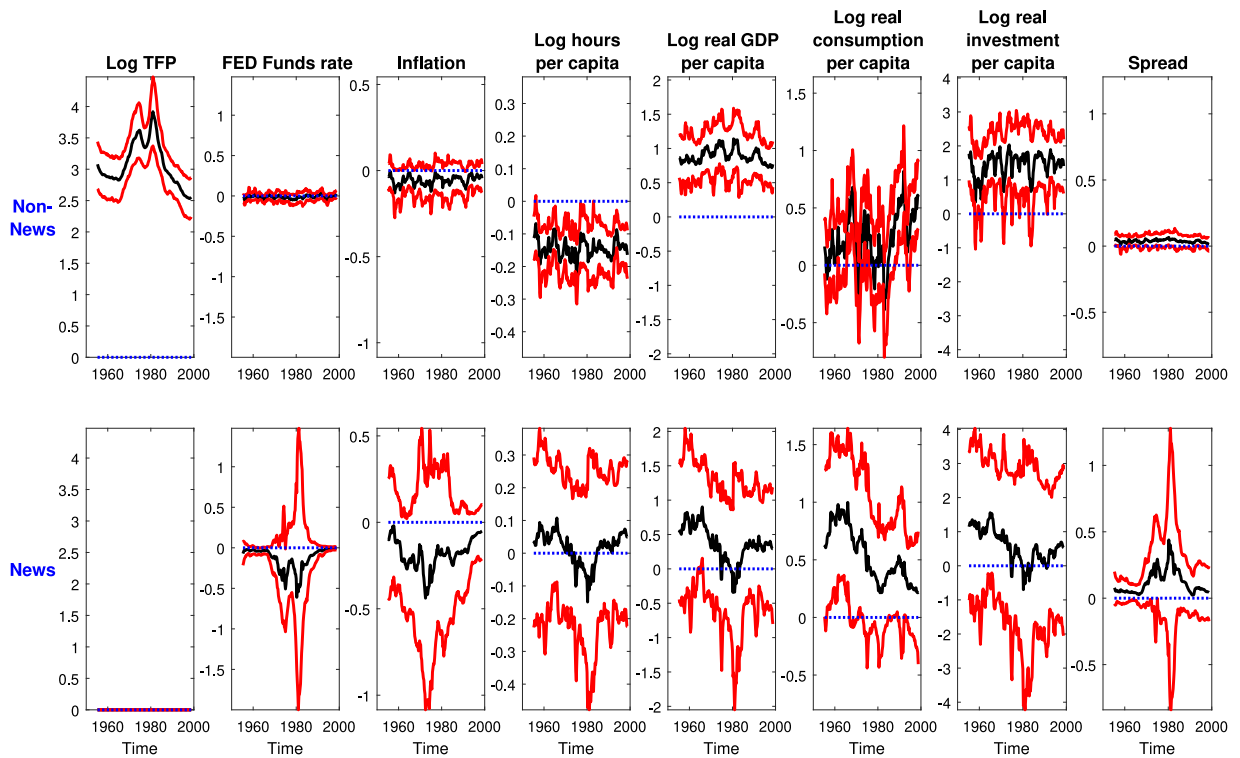


Fig. 1. Time-varying responses to non-news and news shocks on impact (mean, and 16–84 percentiles of the posterior distribution) for the $n = 8$ variables model.

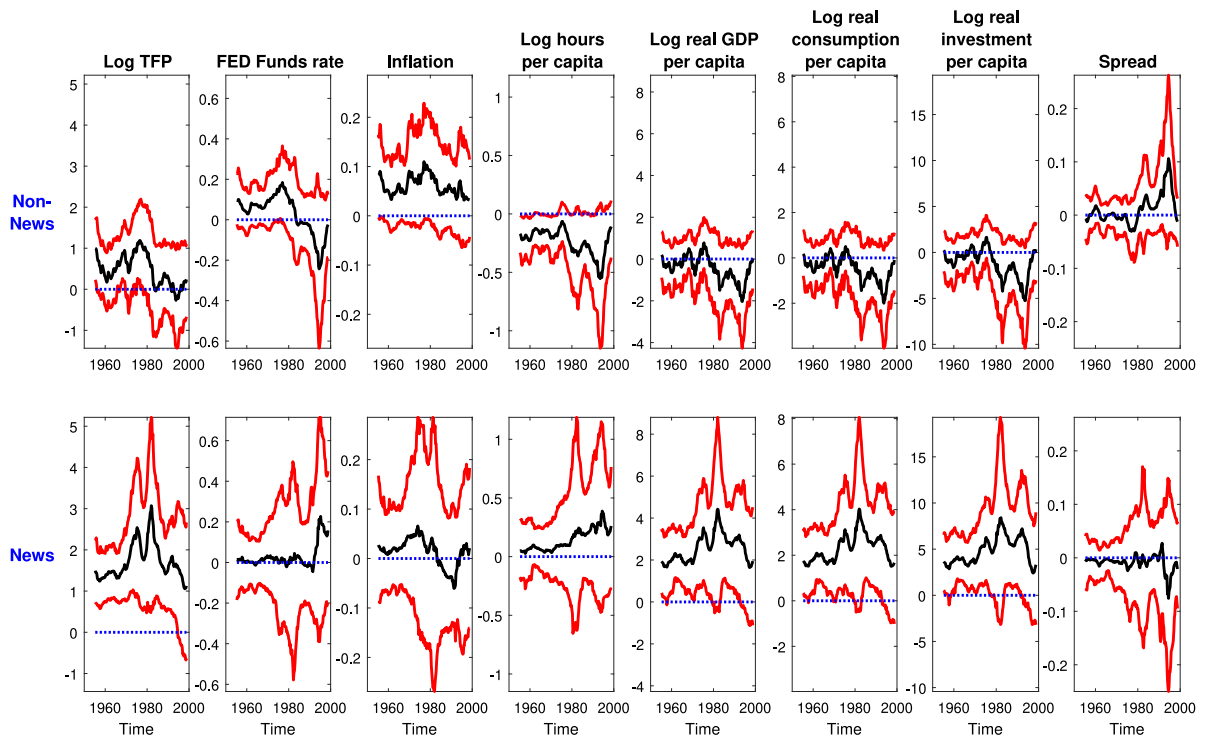


Fig. 2. Time-varying responses to non-news and news shocks at 40 quarters after impact (mean, and 16–84 percentiles of the posterior distribution) for the $n = 8$ variables model.

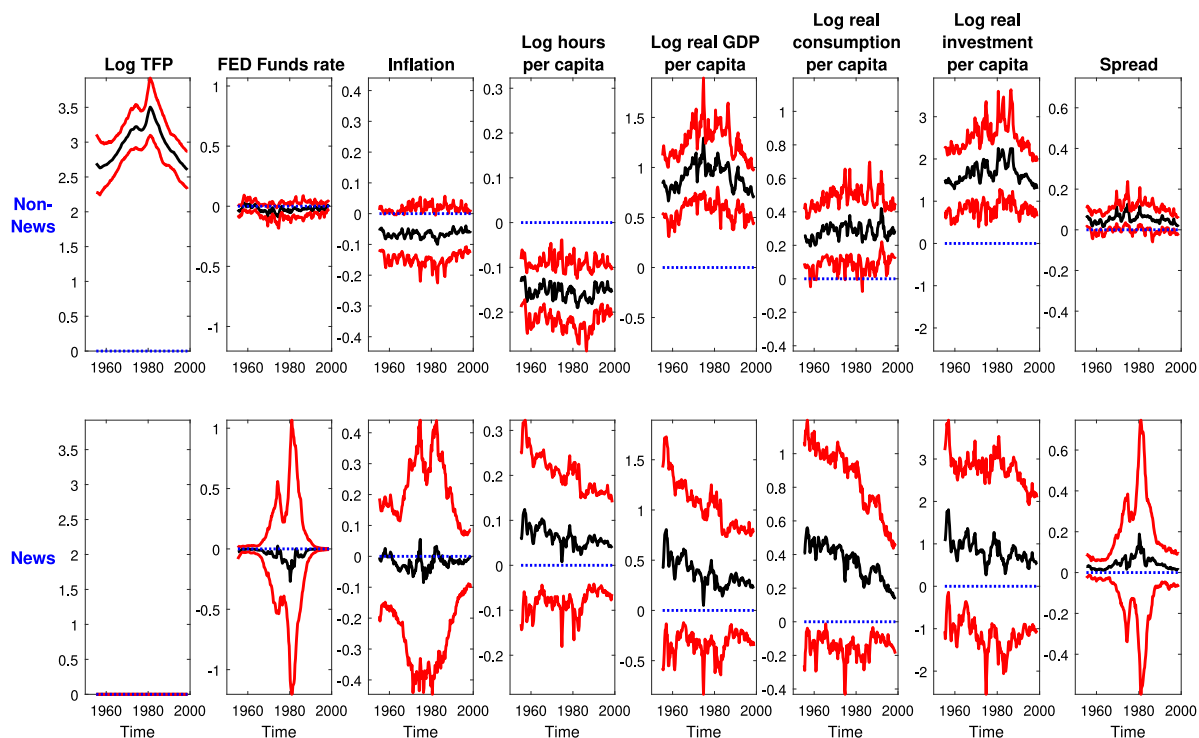


Fig. 3. Time-varying responses to non-news and news shocks on impact (mean, and 16–84 percentiles of the posterior distribution) for the $n = 15$ variables model.

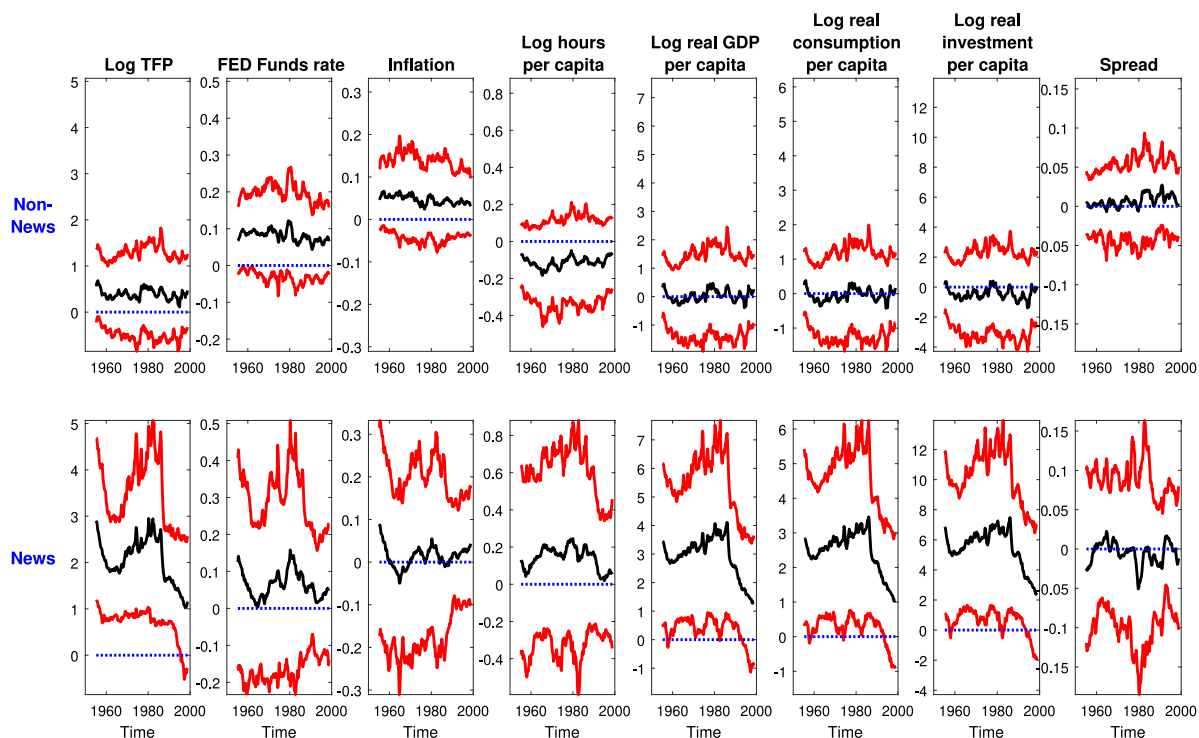


Fig. 4. Time-varying responses to non-news and news shocks at 40 quarters after impact (mean, and 16–84 percentiles of the posterior distribution) for the $n = 15$ variables model.

Table 3
Estimated distances $d(Q_{r_\alpha}, Q_{r_\alpha^*})$.

r_α^*	r_α			
	2	3	4	5
3	0.21			
4	0.31	0.10		
5	0.37	0.14	0.08	
6	0.41	0.18	0.10	0.08

There is a noticeable second order effect upon the estimated posterior impulse responses in both the $n = 8$ and the $n = 15$ models. Specifically, we see that the error bands suggest that there was a very large increase in uncertainty about the immediate effect of news shocks upon the Fed funds rate, the spread and to a lesser extent upon inflation around 1980. It is also in these second order effects upon the posterior that we see the effect of estimating a smaller model. Looking at Figs. 2 and 4, the error bands are much tighter for the larger model despite this model having many more parameters to be estimated. We also see that estimating the smaller model we have the impression that the posteriors for a number of impacts, particularly to non-news shocks, are skewed at particular points in time and have higher probability of producing outliers from one tail at these times. These effects largely disappear when we estimate the larger, less restricted model.

We have used DIC to select the rank r_α . DIC may not select the correct rank and so it is important to know how the variance–covariance matrix for the states, e.g., Q_α , differs for different dimensions r_α and r_α^* say. We end the Results section by investigating the effect on Q_α of changing the rank r_α .

Changing r_α changes the column dimension of A_α and thereby the rank of Q_α . We introduce a slight notation change to distinguish between results from different r_α . Denote the covariance matrix with rank r_α by $Q_\alpha = Q_{r_\alpha} = A_\alpha A_\alpha'$ where A_α is of dimension $k \times r_\alpha$. There are many norms that we could use to measure the distance between Q_{r_α} and $Q_{r_\alpha^*}$ but to help discern whether the distance between Q_{r_α} and $Q_{r_\alpha^*}$ is large, we choose a measure that is bounded on $[0, 1]$. A distance of 0 results if the matrices are the same and a distance of 1 occurs if they are orthogonal to each other. The measure from Herdin et al. (2005) is

$$d(Q_{r_\alpha}, Q_{r_\alpha^*}) = 1 - \frac{\text{tr} Q_{r_\alpha} Q_{r_\alpha^*}}{\|Q_{r_\alpha}\|_F \|Q_{r_\alpha^*}\|_F}$$

where $\|Q_{r_\alpha^*}\|_F$ is the Frobenius norm. For our application with $n = 15$, the DIC chose r_α equal to 4. We estimated models with a range of different ranks and Table 3 reports the various $d(Q_{r_\alpha}, Q_{r_\alpha^*})$ for $r_\alpha, r_\alpha^* = 2, 3, 4, 5, 6$. For example, $d(Q_4, Q_5) = 0.08$ which we take to indicate that overestimating r_α does not change Q_{r_α} much (see Table 3).

The evidence suggests that estimating Q_{r_α} with r_α below the value chosen by DIC (say at 2) does have a noticeable effect on Q_{r_α} while overestimating the rank (say at 6) has less of an impact. We take these results to suggest that the DIC is doing a good job of estimating r_α small enough to reduce the dimension without greatly impacting upon the estimated Q_α .

5. Conclusion

This paper presents an approach to reducing the dimension of the TVP-SVAR. We achieve this by reducing the number of states driving the time varying parameters, while preserving the full number of time varying parameters. The aim is to permit more efficient estimation of larger systems while preserving a full probability model and all formal inferential opportunities. The specification we employ is new and has a number of advantages. The dimension reduction is achieved by choosing a reduced rank of the state equation covariance matrix using empirical evidence. We employ DIC to select the rank of the covariance matrix. The specification is an exact one, allowing estimation of outputs, such as impulse responses and variance decompositions, and their full posterior distributions.

Computation remains a challenge in any large dimensional model, including the one presented in this paper. To mitigate this issue in this model we present a number of techniques that improve computation. These include careful specification of the model, judicious choice of computation algorithm, SSVS with a Minnesota prior to reduce the number of parameters, and use of parameter expansions to attain more readily computable forms for the final model. As a result, we present an approach that increases the range of models available to macroeconomists.

The application to a large system of 15 variables in a time varying VAR suggests that the estimates remain precise with sensible error bands. We find evidence of time variation in the impulse responses and differences between smaller and larger models. There are many directions in which this model could be extended. Subsequent work will consider automated selection of the rank of the state equation covariance matrix and inference on whether specific states vary over time or not (as per Frühwirth-Schnatter and Wagner (2010)).

⁶ In computing \tilde{Q}_t for $t > T - 80$ we set $\Pi_{t,t+s} = \Pi_{t,T}$ for all $t + s \geq T$.

Referees suggested a number of interesting extensions. One would be to allow the number of states to vary over time and change at unknown points such that this change is data driven. Such a model implies a time-varying dimension of the state space and covariance matrix and this extension could accommodate periods in which new latent factors appear in different economic environments. In a study of financial variables and estimating on sub-samples, Ando and Bai (2017) find more factors during the GFE than at other times using pre-selected periods among which the number of factors may change. Koop and Korobilis (2013) demonstrate the significant advantage in forecasting of allowing for the model dimension to evolve over time, not by changing the number of states but rather by changing the number of variables in the system. We are not aware of a model in which automatic or data driven changes in the number of factors at unknown points have been estimated, although one possible approach could follow the approach of Chan et al. (2012) and specify a time-varying dimension model using a dynamic mixture model. This would constitute an important advance but, as discussed in that paper, even that approach faces significant computational limitations.

A number of papers such as Cogley et al. (2010) (with $n = 3$) and Baumeister and Benati (2013) (with $n = 4$) allow for time variation in the state equation covariance matrix. This feature was important to allow for a change in the random-walk drift of macro variables. In each paper, the diagonal elements of the covariance matrix evolve while the off-diagonals are fixed. Unfortunately such a specification would not adapt for the model in this paper due to the dimension of the system and the reduced rank of the covariance matrix. Introducing stochastic volatility for the states in our model would imply letting the matrix A evolve over time as A_t . Stochastic volatility in factors (the states in our model may be seen as factors) and the challenges of estimation of such a model is discussed in Kastner et al. (2017) and references therein. The interweaving strategies discussed in their paper offer a potential approach to computing such a model.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jeconom.2019.11.006>.

References

- Ando, T., Bai, J., 2017. Clustering huge number of financial time series: a panel data approach with high-dimensional predictors and factor structures. *J. Amer. Statist. Assoc.* 112, 1182–1198.
- Bañbura, M., Giannone, D., Reichlin, L., 2010. Large Bayesian vector auto regressions. *J. Appl. Econometrics* 25 (1), 71–92.
- Barsky, R.B., Sims, E.R., 2011. News shocks and business cycles. *J. Monet. Econ.* 58 (3), 273–289.
- Baumeister, C., Benati, L., 2013. Unconventional monetary policy and the great recession: estimating the macroeconomic effects of a spread compression at the zero lower bound. *Int. J. Central Bank.* 165–212.
- Belmonte, M.A., Koop, G., Korobilis, D., 2014. Hierarchical shrinkage in time-varying parameter models. *J. Forecast.* 33, 80–94.
- Bernanke, B., Boivin, J., Elias, P., 2005. Measuring the effects of monetary policy: a factor-augmented vector autoregressive (FAVAR) approach. *Q. J. Econ.* 120, 387–422.
- Bringmann, L.F., Ferrer, E., Hamaker, E.L., Borsboom, D., Tuerlinckx, F., 2018. Modeling nonstationary emotion dynamics in dyads using a time-varying vector-autoregressive model. *Multivar. Behav. Res.* 53 (3), 293–314. <http://dx.doi.org/10.1080/00273171.2018.1439722>.
- Carriero, A., Clark, T.E., Marcellino, M., 2016a. Common drifting volatility in large Bayesian VARs. *J. Bus. Econom. Statist.* 34 (3), 375–390.
- Carriero, A., Clark, T., Marcellino, M., 2016b. Large Vector Autoregressions with stochastic volatility and flexible priors. Working paper 1617, Federal Reserve Bank of Cleveland.
- Carriero, A., Kapetanios, G., Marcellino, M., 2011. Forecasting large datasets with Bayesian reduced rank multivariate models. *J. Appl. Econometrics* 26 (5), 735–761.
- Chan, J.C.C., 2020. Large Bayesian VARs: a flexible Kronecker error covariance structure. *J. Bus. Econom. Statist.* 38 (1), 68–79. <http://dx.doi.org/10.1080/07350015.2018.1451336>.
- Chan, J.C.C., Eisenstat, E., 2018. Bayesian model comparison for time-varying parameter VARs with stochastic volatility. *J. Appl. Econometrics* 33 (4), 509–532.
- Chan, J.C.C., Jeliazkov, I., 2009. Efficient simulation and integrated likelihood estimation in state space models. *Int. J. Math. Modelling Numer. Optim.* 1, 101–120.
- Chan, J.C.C., Koop, G., León-González, R., Strachan, R.W., 2012. Time varying dimension models. *J. Bus. Econom. Statist.* 30 (3), 358–367.
- Chan, J.C.C., R. León-González, R., Strachan, R.W., 2018. Invariant inference and efficient computation in the static factor model. *J. Amer. Statist. Assoc.* 113 (522), 819–828. <http://dx.doi.org/10.1080/01621459.2017.1287080>.
- Chan, J.C.C., Strachan, R.W., 2012. Efficient estimation in non-linear non-Gaussian state-space models. Working paper, Research School of Economics, Australian National University.
- Ciccarelli, M., Rebucci, A., 2007. Measuring contagion and interdependence with a Bayesian time-varying coefficient model: an application to the Chilean FX market during the Argentine crisis. *J. Financ. Econ.* 5, 285–320.
- Cogley, T., Primiceri, G.E., Sargent, T.J., 2010. Inflation-gap persistence in the US. *Am. Econ. J.* 2, 43–69.
- Cogley, T., Sargent, T.J., 2001. Evolving post-world war II inflation dynamics. *NBER Macroecon. Annu.* 16, 331–373.
- Cogley, T., Sargent, T.J., 2005. Drifts and volatilities: Monetary policies and outcomes in the post WWII U.S. *Rev. Econ. Dyn.* 8, 262–302.
- Doan, T., Litterman, R., Sims, C., 1984. Forecasting and conditional projection using realistic prior distributions. *Econometric Rev.* 3, 1–100.
- Eickmeier, S., Lemke, W., Marcellino, M., 2015. A classical time varying FAVAR model: estimation, forecasting, and structural analysis. *J. Royal Stat. Soc. A* 178, 493–533.
- Eisenstat, E., Chan, J.C.C., Strachan, R.W., 2016. Stochastic model specification search for time-varying parameter VARs. *Econometric Rev.* 35, 1638–1665.
- Frühwirth-Schnatter, S., Wagner, H., 2010. Stochastic model specification search for Gaussian and partial non-Gaussian state space models. *J. Econometrics* 154 (1), 85–100.
- George, E.I., McCulloch, R.E., 1993. Variable selection via Gibbs sampling. *J. Amer. Statist. Assoc.* 88 (423), 881–889.
- Geraci, M.V., Gnabo, J.-Y., 2018. Measuring interconnectedness between financial institutions with Bayesian time-varying vector autoregressions. *J. Financ. Quant. Anal.* 53 (03), 1371–1390.

- Giannone, D., Lenza, M., Momferatou, D., Onorante, L., 2014. Short-term inflation projections: a Bayesian vector autoregressive approach. *Int. J. Forecast.* 30 (3), 635–644.
- Herdin, M., Czink, N., Özcelik, H., Bonek, E., 2005. Correlation matrix distance a meaningful measure for evaluation of non-stationary MIMO channels. In: *Proc. IEEE 61st VTC*, vol. 1, pp. 136–140.
- Kastner, G., Frühwirth-Schnatter, S., Lopes, H.F., 2017. Efficient bayesian inference for multivariate factor stochastic volatility models. *J. Comput. Graph. Statist.* 26 (4), 905–917.
- Kim, S., Shephard, N., Chib, S., 1998. Stochastic volatility: likelihood inference and comparison with ARCH models. *Rev. Econom. Stud.* 65, 361–393.
- Koop, G., 2013. Forecasting with medium and large Bayesian VARs. *J. Appl. Econometrics* 28, 177–203.
- Koop, G., Korobilis, D., 2013. Large time-varying parameter VARs. *J. Econometrics* 177, 185–198.
- Koop, G., Léon-González, R., Strachan, R.W., 2009. On the evolution of monetary policy. *J. Econom. Dynam. Control* 33, 997–1017.
- Koop, G., Léon-González, R., Strachan, R.W., 2010. Efficient posterior simulation for cointegrated models with priors on the cointegration space. *Econometric Rev.* 29 (2), 224–242.
- Koop, G., Léon-González, R., Strachan, R.W., 2012. Bayesian model averaging in the instrumental variable regression model. *J. Econom.* 171, 237–250.
- Korobilis, D., 2013. VAR forecasting using Bayesian variable selection. *J. Appl. Econometrics* 28, 204–230.
- Liu, C., Rubin, D.B., Wu, Y.N., 1998. Parameter expansion to accelerate EM: the PX-EM algorithm. *Biometrika* 85 (4), 755–770.
- Liu, J., Wu, Y., 1999. Parameter expansion for data augmentation. *J. Amer. Statist. Assoc.* 94, 1264–1274.
- McCausland, W., Miller, S., Pelletier, D., 2011. Simulation smoothing for state-space models: a computational efficiency analysis. *Comput. Statist. Data Anal.* 55, 199–212.
- Muirhead, R.J., 1982. *Aspects of Multivariate Statistical Theory*. Wiley, New York.
- Primiceri, G., 2005. Time varying structural vector autoregressions and monetary policy. *Rev. Econom. Stud.* 72, 821–852.
- Sims, C.A., 2001. Comment on Sargent and Cogley's 'Evolving US postwar inflation dynamics'. *NBER Macroeconomics Annual* 16, 373–379.
- Zellner, A., 1971. *An Introduction To Bayesian Inference in Econometrics*. John Wiley and Sons Inc, New York.