

**Bayesian Modelling of VAR Precision
Matrices Using Stochastic Block
Networks**

Florian Huber
Gary Koop
Massimiliano Marcellino
Tobias Scheckel

Bayesian Modelling of VAR Precision Matrices Using Stochastic Block Networks

Florian Huber (University of Salzburg), Gary Koop (University of Strathclyde), Massimiliano Marcellino (Bocconi University), Tobias Scheckel (WIFO)

WIFO Working Papers 690/2024
November 2024

Abstract

Commonly used priors for Vector Autoregressions (VARs) induce shrinkage on the autoregressive coefficients. Introducing shrinkage on the error covariance matrix is sometimes done but, in the vast majority of cases, without considering the network structure of the shocks and by placing the prior on the lower Cholesky factor of the precision matrix. In this paper, we propose a prior on the VAR error precision matrix directly. Our prior, which resembles a standard spike and slab prior, models variable inclusion probabilities through a stochastic block model that clusters shocks into groups. Within groups, the probability of having relations across group members is higher (inducing less sparsity) whereas relations across groups imply a lower probability that members of each group are conditionally related. We show in simulations that our approach recovers the true network structure well. Using a US macroeconomic data set, we illustrate how our approach can be used to cluster shocks together and that this feature leads to improved density forecasts.

E-Mail: tobias.scheckel@wifo.ac.at

2024/1/W/23096

© 2024 Austrian Institute of Economic Research

Media owner (publisher), producer: Austrian Institute of Economic Research
1030 Vienna, Arsenal, Objekt 20 | Tel. (43 1) 798 26 01 0 | <https://www.wifo.ac.at>
Place of publishing and production: Vienna

WIFO Working Papers are not peer reviewed and are not necessarily based on a coordinated position of WIFO. The authors were informed about the Guidelines for Good Scientific Practice of the Austrian Agency for Research Integrity (ÖAWI), in particular with regard to the documentation of all elements necessary for the replicability of the results.

Free download: <https://www.wifo.ac.at/publication/pid/55753534>

Bayesian modelling of VAR precision matrices using stochastic block networks

Florian HUBER
University of Salzburg

Gary KOOP
University of Strathclyde

Massimiliano MARCELLINO
Bocconi University
CEPR, IGEIR, BAFFI, BIDS

Tobias SCHECKEL
University of Salzburg
WIFO

October 31, 2024

Abstract

Commonly used priors for Vector Autoregressions (VARs) induce shrinkage on the autoregressive coefficients. Introducing shrinkage on the error covariance matrix is sometimes done but, in the vast majority of cases, without considering the network structure of the shocks and by placing the prior on the lower Cholesky factor of the precision matrix. In this paper, we propose a prior on the VAR error precision matrix directly. Our prior, which resembles a standard spike and slab prior, models variable inclusion probabilities through a stochastic block model that clusters shocks into groups. Within groups, the probability of having relations across group members is higher (inducing less sparsity) whereas relations across groups imply a lower probability that members of each group are conditionally related. We show in simulations that our approach recovers the true network structure well. Using a US macroeconomic data set, we illustrate how our approach can be used to cluster shocks together and that this feature leads to improved density forecasts.

Keywords: Bayesian VAR; Stochastic Block Model; Stochastic search; Markov chain Monte Carlo.

JEL Classification: C11, C15, C30, C32.

We would like to thank Luca Barbaglia, Sune Karlsson, James Mitchell, Luca Onorante, Michael Smith, Mike West as well as participants of the 13th European Seminar on Bayesian Econometrics (Glasgow, 2023), the yearly meeting of the Austrian Economic Association 2023 (Salzburg, 2023), the 6th Annual Workshop on Financial Econometrics (Örebro, 2023) and the 17th International Conference on Computational and Financial Econometrics (Berlin, 2023) for helpful comments and suggestions. Huber and Scheckel gratefully acknowledge funding from the Austrian Science Fund (FWF, grant no. ZK-35) and the Jubiläumsfond of the Oesterreichische Nationalbank (OeNB, grant no. JF-18740 and JF-18763).

1 Introduction

Vector Autoregressions (VARs) are routinely used for structural analysis and predictive inference in academia and policy institutions. These models, however, suffer from overfitting issues if the number of time series is large. Bayesian solutions rely on shrinkage priors so as to force coefficients associated with irrelevant predictors towards zero and thus improve inference (see, among many others, [Carriero et al., 2009](#); [Bańbura et al., 2010](#); [Koop, 2013](#); [Korobilis, 2013](#); [Giannone et al., 2014](#); [Huber and Feldkircher, 2019](#)).

These papers propose shrinking the VAR coefficients towards a known location and by doing so reducing the number of effective variables. This makes sense in light of the huge number of regression coefficients. For moderately-sized panels one often overlooked problem is that the number of parameters in the covariance matrix grows quadratically in the number of time series. Except in a few papers working with noninformative priors in VARs of low dimension, Bayesians use informative priors such as the inverse-Wishart prior on the error covariance matrix Σ (or equivalently the Wishart prior on the error precision $\Omega := \Sigma^{-1}$) and, thus, induce shrinkage. However these priors do not induce parsimony (i.e. by restricting error covariances to be exactly or nearly zero).

Some papers that induce shrinkage on the error covariance matrix are [George et al. \(2008\)](#) and [Koop \(2013\)](#) which use priors such as the stochastic search variable selection (SSVS) prior. These priors assume that the variable selection parameters apply to each coefficient individually and independently (i.e. each error covariance is either shrunk to zero or not, independently of the other covariances). Thus, they do not capture situations where covariances cluster together. Moreover, these techniques introduce shrinkage on the lower Cholesky factor of Σ or Ω , implying that they are not order-invariant. In a recent paper, [Arias et al. \(2023\)](#), show that the (lack of) order invariance matters for predictive

likelihoods and hence, having techniques that are order-invariant might be preferable.¹

In this paper, we develop a new method that takes the network structure of the contemporaneous relations into account and is order-invariant. Our model assumes a time-varying covariance matrix that decomposes the covariance matrix into a time-varying and a time-invariant part. The time-invariant part establishes the contemporaneous relations across equations. On this part, we introduce our new shrinkage prior. Our prior draws on insights from the network literature and, in particular, is based on a stochastic block model ((SBM), see, e.g., [Legramanti et al., 2022](#)).

For many applications in macro and finance the covariance or precision matrices of the shocks can be interpreted as a network that encodes the contemporaneous relationships between different units (such as companies, countries or variable types) and these shocks might form clusters. For instance, shocks specific to a region might have an immediate effect only on variables within that given region whereas other regions are impacted only with a time lag through the VAR coefficients. These patterns have implications for the covariances of the shocks, leading to a possibly sparse structure. For example, if such a regional structure exists, a successful shrinkage prior should not introduce shrinkage on the blocks of the precision matrix within a given region but should force blocks between regions to zero. Standard shrinkage priors including SSVS and variants thereof do not take the presence of clusters into account. [Ahelegbey et al. \(2016b,a\)](#); [Billio et al. \(2019\)](#) use graphical models to shrink covariance matrices. However, their priors on the network still assume that the edge probabilities are independent a priori. In addition, to sample the network they require a Metropolis Hastings (MH) step. This can suffer from poor mixing, especially since the parameter space is extremely large.

This motivates the present paper. Our goal is to set up a standard VAR with heteroskedastic shocks. We use an SSVS prior ([George and McCulloch, 1993](#)) to shrink the elements of the precision matrix to zero. However, instead of assuming that the

¹For recent, order-invariant, approaches, see [Chan et al. \(2024\)](#) and [Huber and Koop \(2024\)](#).

indicators that control whether a given precision parameter should be forced to zero or not arise from a Bernoulli with a common (and often fixed) prior inclusion probability, we model the latter using a SBM (see, e.g., [Holland et al., 1983](#); [Nowicki and Snijders, 2001](#); [Legramanti et al., 2022](#)). The resulting model endogenously detects clusters and thus introduces shrinkage on the VAR precision matrix while taking possible within- and cross-cluster linkages into account. The SBM we use does not require the pre-selection of the number of clusters, rather it is inferred adaptively alongside the remaining model parameters.

To make sure that our resulting SBM-VAR model is scalable to large dimensions, we develop an efficient Markov Chain Monte Carlo (MCMC) sampler to simulate from the joint posterior distribution. This sampler builds on recent advances in MCMC estimation of large VARs (see [Carriero et al., 2022](#)).

To evaluate and illustrate our approach, we first carry out a thorough simulation exercise. Using synthetic data generated from a variety of alternative specifications, we show that our model accurately recovers the true network structure, and generally improves upon the standard SSVS prior.

We then move on to a forecasting exercise involving a large set of US economic variables. Similar to our synthetic data exercise, our model finds parsimonious network structures although it is interesting to note that these change somewhat over time both in terms of the number of clusters and their modularity. The forecasting exercise shows that this property tends to lead (with some exceptions) to better forecast performance than either an SSVS prior which does not incorporate a network structure or a non-informative prior.

The remainder of the paper is structured as follows. Section 2 introduces the basic VAR model, briefly describes commonly used priors on the VAR coefficients and the error covariances, provides a brief introduction to stochastic block models in the context of VAR error precision matrices, introduces the SBM-VAR, specifically focusing on the

VAR precision prior, and discusses the MCMC sampler. Section 3 provides simulation evidence that our model works well. Section 4 presents the empirical applications. Section 5 summarizes and concludes the paper. An Appendix contains additional details on the sampler and on the data used in the empirical application.

2 Stochastic Block Network VARs

2.1 VARs with Stochastic Volatility

Our goal is to model an M -dimensional vector of time series \mathbf{y}_t that is observed at time $t = 1, \dots, T$ using a VAR model:²

$$\mathbf{y}_t = \mathbf{A}_1 \mathbf{y}_{t-1} + \dots + \mathbf{A}_P \mathbf{y}_{t-P} + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_t), \quad (1)$$

where \mathbf{A}_j ($j = 1, \dots, M$) are $M \times M$ coefficient matrices, P is the maximum lag length and the shocks in $\boldsymbol{\varepsilon}_t$ are independent, zero mean and normally distributed with a time-varying variance-covariance matrix $\boldsymbol{\Sigma}_t$.

The existing literature has many treatments of the error covariance matrix for VARs based on different decompositions and the manner in which time-variation is allowed for. Perhaps the most common, see e.g. [Cogley and Sargent \(2005\)](#), is to take a Cholesky decomposition $\boldsymbol{\Sigma}_t = \mathbf{L} \mathbf{D}_t^2 \mathbf{L}'$ where \mathbf{L} is lower triangular with ones on the diagonal and $\mathbf{D}_t^2 = \text{diag}(e^{d_{1,t}}, \dots, e^{d_{M,t}})$ is a diagonal matrix of time-varying variances. It is common to assume, as we do in this paper, that $d_{j,t}$ evolves according to independent AR(1) models: $d_{j,t} = \rho_j d_{j,t-1} + \sigma_{\rho,j} v_{j,t}$ with ρ_j denoting the persistence parameter and $\sigma_{\rho,j}^2$ the error variance.

Use of the Cholesky decomposition has been criticized in papers such as [Arias et al.](#)

²In our theoretical discussion we ignore deterministic terms in the model. In our empirical application we include an intercept.

(2023) since it leads to a lack of order invariance.³ In this paper, we use a different decomposition of Σ_t which has attractive properties for our purposes in terms of ease of interpretation and computational efficiency.

To explain our specification, we begin with the following decomposition and define $\Omega_t := \Sigma_t^{-1}$ which is the error precision matrix:

$$\Sigma_t = \mathbf{D}_t \Sigma \mathbf{D}_t \quad \Leftrightarrow \quad \Omega_t = \mathbf{D}_t^{-1} \Omega \mathbf{D}_t^{-1},$$

where $\mathbf{D}_t = \text{diag}(e^{d_{1,t}/2}, \dots, e^{d_{M,t}/2})$ is a diagonal matrix of time-varying standard deviations so that $\mathbf{D}_t^2 = \mathbf{D}_t \odot \mathbf{D}_t$ with \odot meaning element-wise multiplication. For convenience, let $\mathbf{d}_t = (d_{1,t}, \dots, d_{M,t})'$ denote the M -dimensional vector of time-specific log-volatilities. Note that the decomposition involves two time-invariant matrices, Σ and Ω which, with some abuse of terminology we still refer to as covariance and precision matrices. That is, they are the covariance and precision matrices abstracting from the time-varying standard deviations. If elements of Ω (Σ) are zero then the corresponding elements of Ω_t (Σ_t) are zero.

Note that we face a choice as to whether to place our SBM prior on Ω or Σ . Either is possible, but we place it on Ω . This is for reasons outlined in papers such as Fan et al. (2016). This paper points out that covariance and precision matrices contain different types of information. In the context of sparse estimation (such as we are attempting to do using our SBM prior), working with the precision matrix may be more sensible since precision matrices are often sparser than covariance matrices. A sparse precision matrix does not necessarily imply a sparse covariance matrix (and vice versa). Fan et al. (2016) point out that the covariance matrix relates to marginal correlations between variables but the precision matrix relates to conditional correlations. That is, the precision

³Arias et al. (2023) also discusses a range of order invariant approaches and discusses their properties. Our method for decomposing the error covariance matrix is not the same as the order invariant specification of Archakov and Hansen (2021) but does share a similar structure to it.

matrix encodes the conditional dependence between $\boldsymbol{\varepsilon}_{i,\bullet}$ and $\boldsymbol{\varepsilon}_{j,\bullet}$ the i and j^{th} column of $\boldsymbol{E} = (\boldsymbol{\varepsilon}_1, \dots, \boldsymbol{\varepsilon}_T)'$. And these conditional correlations can be mapped into an undirected graph. In essence, if a sparse graphical or network structure exists it is better to model it via the precision matrix rather than the covariance matrix.

Decomposing $\boldsymbol{\Sigma}_t$ into $\boldsymbol{\Sigma}_t = \boldsymbol{D}_t \boldsymbol{\Sigma} \boldsymbol{D}_t$ implies that the $(i, j)^{\text{th}}$ element in $\boldsymbol{\Sigma}_t$ is scaled by $e^{d_{i,t}/2} \cdot e^{d_{j,t}/2}$ and hence the covariances depend on the volatility dynamics across the different equations. Moreover, and this constitutes another advantage of this decomposition, if the prior is placed on $\boldsymbol{\Omega}$, then the model is order-invariant. It would not be so if we used some sort of SBM prior in the Cholesky-decomposed model.

$\boldsymbol{\Sigma}$ can be further decomposed into $\boldsymbol{\Sigma} = \boldsymbol{D} \boldsymbol{R} \boldsymbol{D}$ with $\boldsymbol{D} = \text{diag}(e^{d_1/2}, \dots, e^{d_M/2})$ and \boldsymbol{R} denoting the correlation matrix. The matrix $\boldsymbol{D}^2 = \boldsymbol{D} \odot \boldsymbol{D}$ can be seen to introduce an intercept into the AR(1) models for the log-volatilities since $\boldsymbol{D}_t \odot \boldsymbol{D} = \text{diag}(e^{(d_1+d_{1,t})/2}, \dots, e^{(d_M+d_{M,t})/2})$. As we will discuss below, we use this further decomposition for computational reasons since, conditional on \boldsymbol{D}_t , the rows and columns in $\boldsymbol{\Sigma}$ can be sampled using a Gibbs sampler. Other related specifications, such as that of Archakov and Hansen (2021), which do not use this decomposition of $\boldsymbol{\Sigma}$ can be estimated using Bayesian MCMC methods. However, these MCMC algorithms are not straightforward Gibbs samplers and are not scalable to large dimensions.

2.2 Network-based priors for VARs

If the number of variables M (and/or the lag length P) is large, the model suffers from overfitting. Most notably, the number of dynamic coefficients in the VAR is $P \times M^2$ and thus grows quadratically with M . Bayesian shrinkage priors such as the Minnesota prior (e.g., Doan et al., 1984; Litterman, 1986; Giannone et al., 2015)) or global-local shrinkage priors (e.g., Huber and Feldkircher, 2019; Follett and Yu, 2019; Kastner and Huber, 2020) introduce shrinkage so as to reduce the effective number of elements in the VAR coefficients. However, as stated in the introduction, these priors either introduce

little shrinkage on the $M(M - 1)/2$ free elements of the error covariance matrix or do so in an unstructured manner. In the homoskedastic case, it is common to use an inverse Wishart prior on the covariance matrix of the VAR (reduced form) errors (e.g., [Kadiyala and Karlsson, 1997](#); [Giannone et al., 2015](#); [Chan, 2021](#)). Often relatively non-informative choices are made for the prior hyperparameters of the inverse Wishart prior. Many other papers, which allow for stochastic volatility, use a Gaussian prior for the lower triangular \mathbf{L} (or \mathbf{L}^{-1}), see, e.g., [Cogley and Sargent \(2005\)](#); [Carriero et al. \(2019\)](#). This strategy gives rise to three issues. First, it does not take a possible network structure into account. Second, it is not clear what shrinkage on \mathbf{L} implies for $\boldsymbol{\Sigma}_t$ or $\boldsymbol{\Omega}_t$. Third, placing a prior on \mathbf{L} implies that all empirical findings depend on how the elements in \mathbf{y}_t are ordered. The goal of this paper is to develop a shrinkage prior that surmounts these problems. For the reasons given in the preceding sub-section, we will do so by constructing a prior designed to capture a sparse network structure on $\boldsymbol{\Omega}$ and, thus, the error precision matrix $\boldsymbol{\Omega}_t$.

Let $\delta_{i,j}$ denote a binary indicator that equals 1 if there exists a relationship between $\varepsilon_{i,t}$ and $\varepsilon_{j,t}$ and 0 if they are assumed to be conditionally independent. The indicator $\delta_{i,j}$ will be used to set up a shrinkage prior on the elements in $\boldsymbol{\Omega}$. Following [George and McCulloch \(1993\)](#) and [George et al. \(2008\)](#), we assume that the off-diagonal elements of $\boldsymbol{\Omega}$, $\omega_{i,j}$ ($i = 2, \dots, M; j = 1, \dots, M - 1; i < j$) arise from a mixture of two Gaussian distributions:

$$\omega_{i,j} \sim \delta_{i,j} \mathcal{N}(0, \tau_{ij,1}^2) + (1 - \delta_{i,j}) \mathcal{N}(0, \tau_{ij,0}^2), \quad (2)$$

with $\tau_{ij,1}^2$ (called the 'slab' variance) and $\tau_{ij,0}^2$ (called the 'spike' variance) denoting prior scaling parameters such that $\tau_{ij,1} \gg \tau_{ij,0}$. These scaling parameters can be set using an empirical Bayes approach which involves the use of the sample standard deviation of the time series under consideration or using the OLS estimate of it from an auxiliary regression. Another approach is to follow [Ishwaran and Rao \(2005\)](#) who let $\tau_{ij,0}^2 = c \cdot \tau_{i,j}^2$ with c being a constant close to zero and $\tau_{i,j}^2$ is a common (over the two distributions in

the mixture) scaling parameter.

Given that the OLS estimates of $\mathbf{\Omega}$ are difficult to obtain due to the presence of SV, we follow the latter approach and estimate the hyperparameters of the prior. We follow [Ishwaran and Rao \(2005\)](#) and place an inverse gamma prior on the slab variance, such that:

$$\tau_{i,j}^2 \sim \text{InvGa}(a_\tau, b_\tau),$$

where $\text{InvGa}(\cdot, \cdot)$ denotes the inverse gamma distribution.⁴ We set $\alpha_\tau = 5$ and $\beta_\tau = 4$. We obtain the spike variance by multiplying $\tau_{i,j}^2$ with the constant $c = 2.5^{-5}$.

The key point to make here is that $\delta_{i,j}$ is a binary indicator that controls which component to use and prior independence is assumed over the elements of $\mathbf{\Omega}$. In other words, the following relationship holds:

$$\delta_{i,j} = 0 \quad \Leftrightarrow \quad \boldsymbol{\varepsilon}_{\bullet,i} \perp\!\!\!\perp \boldsymbol{\varepsilon}_{\bullet,j} | \{\boldsymbol{\varepsilon}_{\bullet,s}\}_{s \neq i,j} \quad \Leftrightarrow \quad \omega_{ij} \approx 0, \quad (3)$$

implying that $\boldsymbol{\varepsilon}_{\bullet,i}$ and $\boldsymbol{\varepsilon}_{\bullet,j}$ are approximately independent given the other shocks in the system. A typical assumption is that the indicators arise from a Bernoulli prior distribution with prior probability $\underline{\pi}$:

$$\delta_{i,j} \sim \text{Bernoulli}(\underline{\pi}).$$

Hence, the probability that $\delta_{i,j} = 1$ is $\underline{\pi}$ a priori which is the same for all i, j . This assumption does not take into account a possible network structure and thus useful statistical information may be lost. Network models such as the ones proposed in [Wang \(2015\)](#) or [Billio et al. \(2019\)](#) fix $\underline{\pi} = 2/(M - 1)$ or estimate $\underline{\pi}$.

⁴Here, a random variable z follows an inverse Gamma distribution with density $p(z|\alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} z^{-\alpha-1} \exp(-\frac{\beta}{z})$.

2.2.1 Combining SBMs and spike and slab priors

In this paper we embed a network structure into an SSVS prior through individual inclusion probabilities $\underline{\pi}_{i,j}$. Since estimating $N = M(M - 1)/2$ separate inclusion probabilities might lead to overfitting, we exploit the notion that shocks tend to cluster together. To achieve this, we assume that a binary adjacency matrix Δ with generic element $\delta_{i,j}$ arises from a stochastic block model (SBM, see [Holland et al., 1983](#); [Nowicki and Snijders, 2001](#)). Instead of assuming that the indicators $\delta_{i,j}$ are i.i.d. Bernoulli distributed with a common success probability $\underline{\pi}$, we allow for a latent network involving a SBM to govern the contemporaneous relationships among the shocks.

We start our discussion by introducing terminology necessary to describe networks. We assume that the time-invariant network encoding can be represented by a latent, random graph \mathcal{G} which in our case encodes possible networks that give rise to Ω . We furthermore assume that \mathcal{G} is undirected (i.e. $\delta_{i,j} = \delta_{j,i}$) and unweighted (all links are of the same magnitude). It is given by the duplet $(\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = 1, \dots, M$ denotes the vertex set (the members of the network) and $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$ the edge set (the connections among its members) which has cardinality N .

The graph is random in the sense that the edge set is generated by a Bernoulli process, but we assume the vertex set to be fixed. A binary representation of the graph can be obtained through the adjacency matrix Δ which has elements $\delta_{i,j} = 1$ if an edge exists between nodes i and j and $\delta_{i,j} = 0$ otherwise. It is worth stressing that, in contrast to papers such as [Legramanti et al. \(2022\)](#), our adjacency matrix is latent and controls whether the elements in Ω are shrunk to zero or not.

In the case of an SBM \mathcal{G} is driven by a community structure that clusters the nodes into H distinct groups. For typical macroeconomic applications, these clusters could represent country groups (e.g., emerging and developed economics), variable types (e.g., prices, labor market quantities or financial series) or types of shocks (e.g., demand versus




Figure 1: Realization of a stochastic block model with $H = 3$ and $M = 30$. Node and annotation colors indicate group membership. (a) Plot of the network (thicker lines indicate higher edge probabilities). (b) Heatmap representation of the adjacency matrix. (c) Heatmap of the edge probabilities (darker cells indicate higher edge probabilities).

supply shocks; common versus idiosyncratic shocks). We do not assume any such known structure, but rather estimate both the membership of each group and the number of groups.

Let $h, k \in \{1, \dots, H\}$ denote the group memberships of nodes i and j , respectively. An edge exists between i and j with probability $\pi_{h,h} \in [0, 1]$ if i and j are members of the same group h . If the groups of i and j differ, an edge exists with probability $\pi_{h,k}$. This implies that relations within as well as across groups are assumed to be homogeneous a priori. A typical assumption is that $\pi_{h,h} \gg \pi_{h,k}$, such that vertices within the same group are more likely to be connected with one another.

Figure 1 displays an example of such a network with $M = 30$ nodes and $H = 3$ groups alongside heatmaps of the implied adjacency matrix and edge probabilities. The key feature is that variables which display similar properties are grouped together in one of three groups. These groups, however, arise endogenously and we will use this feature to design our prior for the elements of Ω . Before discussing the mathematical structure of the model it is also worth emphasizing that the number of groups H is not fixed a priori. In what follows we let the model decide on the number of groups and estimate it alongside the remaining model parameters.

Our approach builds on the Extended Stochastic Block Model (ESBM) framework put

forward by [Legramanti et al. \(2022\)](#). It extends the original SBM established by [Holland et al. \(1983\)](#) and [Nowicki and Snijders \(2001\)](#) by providing a unified modeling framework and an efficient sampling algorithm for a range of different model formulations, with full uncertainty quantification and without the need to pre-specify the number of clusters H .

The SBM implies a conditional prior on the latent adjacency matrix Δ :

$$p(\Delta|\mathbf{\Pi}) = \sum_{h=1}^H \sum_{k=1}^h \pi_{h,k}^{m_{h,k}} (1 - \pi_{h,k})^{\bar{m}_{h,k}},$$

Here, $\mathbf{\Pi}$ is a $H \times H$ symmetric matrix with generic elements $\pi_{h,k}$. $m_{h,k}$ and $\bar{m}_{h,k}$ denote the number of edges and non-edges between groups h and k , respectively.

Notice that this prior only depends on the edge probabilities across groups and not across nodes. This is the main source of shrinkage in our model and requires the introduction of a set of group assignment indicators $\mathbf{z} = (z_1, z_2, \dots, z_M)$, which is an M -dimensional vector. Each element of \mathbf{z} , z_j , is a discrete random variable which takes values $1, \dots, H$. The node-specific probabilities can then be obtained as $\underline{\pi}_{ij} = \pi_{z_i, z_j}$.

We assume that the random discrete vector \mathbf{z} arises from a general Gibbs-type prior ([Legramanti et al., 2022](#)):⁵

$$p(\mathbf{z}) = \mathscr{W}_{M,H} \prod_{h=1}^H (1 - \sigma)_{n_h - 1}, \quad (4)$$

where $(a)_n$ for any $a > 0$ denotes the ascending factorial $a(a+1) + \dots + (a+n-1)$, n_h is the number of members in group h , $\sigma < 1$ a discount parameter and $\{\mathscr{W}_{M,H} : 1 \leq H \leq M\}$ a collection of non-negative weights satisfying $\mathscr{W}_{M,H} = (M - H\sigma)\mathscr{W}_{M+1,H} + \mathscr{W}_{M+1,H+1}$ and $\mathscr{W}_{1,1} = 1$. This general formulation nests a broad range of priors on the clustering structure that can accommodate different types of networks.⁶

⁵[Legramanti et al. \(2022\)](#) argue that previous SBMs (e.g., [Nowicki and Snijders, 2001](#); [Kemp et al., 2006](#); [Schmidt and Morup, 2013](#); [Geng et al., 2019](#)) can be interpreted as arising from the use of different Gibbs-type priors.

⁶See [Lijoi et al. \(2007a,b\)](#); [Blasi et al. \(2013\)](#); [De Blasi et al. \(2015\)](#); [Miller and Harrison \(2018\)](#) for

We consider various Gibbs-type priors in this paper. In particular, we use the Gnedin-process (GN; [Gnedin, 2010](#)), Dirichlet-Multinomial (DM; [Nowicki and Snijders, 2001](#)), Dirichlet-process (DP; [Kemp et al., 2006](#)) or Pitman-Yor-process (PY; [Pitman and Yor, 1997](#)) priors. Different choices of Gibbs-type priors lead to different grouping structures. As discussed in [Legramanti et al. \(2022\)](#), the DM prior typically generates modular, relatively stable clusters with a fixed number \bar{H} of population clusters. The GN prior can be seen as a generalization thereof, where \bar{H} is modeled as random but still assumed to be finite in population. The DP and PY priors on the other hand favor more fragmented networks and smaller groups with no limit on \bar{H} as M increases. They differ in the growth of H , which is much faster for the PY than for the DP prior. When specifying the hyperparameters for the Gibbs type priors, a sensible approach is to choose them such that their expected number of clusters matches the researcher’s prior beliefs.

The SBM prior is completed by specifying a prior on $\pi_{h,k}$. We follow [Legramanti et al. \(2022\)](#) and specify independent Beta-distributed priors on $\pi_{h,k} \sim \text{Beta}(a_\pi, b_\pi)$. Setting $a_\pi = b_\pi = 1$ induces a uniform prior bounded by zero and one.

2.2.2 Priors on the remaining parameters

In this sub-section, we briefly sketch the remaining priors on the model parameters. This includes the prior on the VAR coefficients and the prior on the parameters in the state equation of the log-volatilities.

The prior on the VAR coefficients in $\mathbf{a} = \text{vec}(\mathbf{A}_1, \dots, \mathbf{A}_p)$ matrices in (1) is a standard Horseshoe shrinkage prior (see, e.g., [Carvalho et al., 2010](#)). The Horseshoe implies a hierarchical Gaussian prior on each element in \mathbf{a} :

$$a_j \sim \mathcal{N}(0, c_j^2 d^2), \quad c_j \sim \mathcal{C}^+(0, 1), \quad d \sim \mathcal{C}^+(0, 1).$$

reviews and examples of various priors that fall into this category.

Here, we let c_j denote a local shrinkage parameter and d a global shrinkage factor. Both have half-Cauchy distributed prior distributions. We store these parameters in a $K(=M^2P) + 1$ vector $\boldsymbol{\xi} = (c_1^2, \dots, c_{M^2P}^2, d)'$. We rely on this prior given that it works well in large dimensions and its recent success in many macroeconomic forecasting applications (Follett and Yu, 2019; Feldkircher et al., 2022; Bai et al., 2022) but any other large VAR prior could be used without altering the main themes of this paper.

On the parameters of the state equation of the log-volatilities we use a truncated Gaussian distribution with mean equal to 0.7, variance of 0.1 and support over $[-0.99, 0.99]$ on the persistence parameter and a Gamma distributed prior with shape parameter 10 and rate parameter equal to 2 on the error precision of the state equation.

2.3 Posterior simulation

In this sub-section we discuss how to carry out posterior simulation in the SBM-VAR.

The joint posterior distribution of the parameters and latent states is given by:

$$p(\mathbf{a}, \boldsymbol{\xi}, \{\mathbf{d}_t\}_{t=1}^T, \boldsymbol{\Omega}, \boldsymbol{\Delta}, \{\tau_{i,j}\}_{i,j}, \boldsymbol{\Pi}, \mathbf{z} | \mathbf{Y})$$

where $\mathbf{Y} = (\mathbf{y}_1, \dots, \mathbf{y}_T)'$. This joint distribution is not available in closed form. However, for most parameters we have conditional distributions that take a well known form and hence are amenable to Gibbs sampling. Here, we provide a detailed overview of the algorithm. Additional details are provided in Appendix A.

Our MCMC algorithm cycles between the following steps:

1. **Sampling from** $p(\mathbf{a} | \boldsymbol{\xi}, \{\mathbf{d}_t\}_{t=1}^T, \boldsymbol{\Omega}, \mathbf{Y})$. We sample the rows of $\mathbf{A} = (\mathbf{A}_1, \dots, \mathbf{A}_p)$ using the triangularization algorithm developed in Carriero et al. (2022). Each of the M rows, labeled $\mathbf{a}_{j,\bullet}$ ($j = 1, \dots, M$) is normally distributed so that:

$$\mathbf{a}_{j,\bullet} | \mathbf{A}_{-j,\bullet}, \boldsymbol{\xi}, \{\mathbf{d}_t\}_{t=1}^T, \boldsymbol{\Omega}, \mathbf{Y} \sim \mathcal{N}(\bar{\mathbf{a}}_{j,\bullet}, \bar{\mathbf{V}}_{a,j}),$$

with $\mathbf{A}_{-j,\bullet}$ being the matrix \mathbf{A} with the j^{th} row excluded and $\overline{\mathbf{a}}_{j,\bullet}$ and $\overline{\mathbf{V}}_{a,j}$ denoting the posterior mean and variance, respectively. Precise forms of these can be found in Appendix A. Repeating this step for each row provides a valid draw from $p(\mathbf{a}|\boldsymbol{\xi}, \{\mathbf{d}_t\}_{t=1}^T, \boldsymbol{\Omega}, \mathbf{Y})$.

2. **Sampling from $p(\boldsymbol{\xi}|\mathbf{a}, \mathbf{Y})$.** The prior variances $\boldsymbol{\xi}$ are simulated based on the algorithm proposed in Makalic and Schmidt (2016). This involves introducing additional auxiliary parameters which are then sampled from inverse Gamma conditionals. Based on these, the conditional posteriors of $\{c_j^2\}_{j=1}^{M^2P}$ and d^2 are inverse Gamma as well. More details can be found in Appendix A.
3. **Sampling from $p(\{\mathbf{d}_t\}_{t=1}^T|\mathbf{a}, \boldsymbol{\Omega}, \mathbf{Y})$.** We update \mathbf{d}_t for $t = 1, \dots, T$ using the single-move algorithm outlined in Ishihara and Omori (2012). In our simulations and real data applications this simple sampler works well, yielding satisfactory mixing properties.
4. **Sampling from $p(\boldsymbol{\Omega}|\mathbf{a}, \{\mathbf{d}_t\}_{t=1}^T, \boldsymbol{\Delta}, \{\tau_{i,j}\}_{i,j}, \mathbf{Y})$.** Generating draws for $\boldsymbol{\Omega}$ is made difficult by the fact that it needs to be positive definite. Wang (2012, 2015) put forward a change-of-variable approach to sequentially sample the rows and columns of $\boldsymbol{\Omega}$. Let $\mathbf{S} = \tilde{\boldsymbol{\varepsilon}}'\tilde{\boldsymbol{\varepsilon}}$ where $\tilde{\boldsymbol{\varepsilon}} = (\mathbf{D}_1^{-1}\boldsymbol{\varepsilon}_1, \dots, \mathbf{D}_T^{-1}\boldsymbol{\varepsilon}_T)'$ are the VAR residuals rescaled using the time-varying part of the error volatility. Furthermore, let $[\underline{\mathbf{V}}_{\Omega}]_{ij} = \delta_{i,j}\tau_{i,j}^2 + (1 - \delta_{i,j})\tau_{i,j,0}^2$ be the prior variance matrix for $\boldsymbol{\Omega}$. To sample column j ($j = 1, \dots, M$) we permute $\boldsymbol{\Omega}$, \mathbf{S} and $\underline{\mathbf{V}}_{\Omega}$ such that j is ordered last. We denote the permuted matrices $\tilde{\boldsymbol{\Omega}}$, $\tilde{\mathbf{S}}$ and $\tilde{\underline{\mathbf{V}}}_{\Omega}$ and apply the following decomposition:

$$\tilde{\boldsymbol{\Omega}} = \begin{pmatrix} \boldsymbol{\Omega}_{(-j)(-j)} & \boldsymbol{\omega}_{(-j)j} \\ \boldsymbol{\omega}'_{(-j)j} & \omega_{jj} \end{pmatrix}, \quad \tilde{\mathbf{S}} = \begin{pmatrix} \mathbf{S}_{(-j)(-j)} & \mathbf{s}_{(-j)j} \\ \mathbf{s}'_{(-j)j} & s_{jj} \end{pmatrix}, \quad \tilde{\underline{\mathbf{V}}}_{\Omega} = \begin{pmatrix} \underline{\mathbf{V}}_{\Omega,(-j)(-j)} & \underline{\mathbf{v}}_{\Omega,(-j)j} \\ \underline{\mathbf{v}}'_{\Omega,(-j)j} & \underline{v}_{\Omega,jj} \end{pmatrix},$$

where $\boldsymbol{\Omega}_{(-j)(-j)}$ are all columns and rows of $\boldsymbol{\Omega}$ except for the j -th, $\boldsymbol{\omega}_{(-j)j}$ is the j -th column of $\boldsymbol{\Omega}$ without its j -th element and ω_{jj} is the j -th diagonal element of $\boldsymbol{\Omega}$.

As shown in Wang (2012, 2015) it is possible to introduce two auxiliary variables \mathbf{u} and v whose distributions have a well known form. In particular, we sample from the following conditional posterior distributions:

$$\begin{aligned} v &\sim \text{Ga}\left(\frac{T}{2} + 1, \frac{s_{jj}}{2}\right), \\ \mathbf{u} &\sim \mathcal{N}\left(-\mathbf{C}\mathbf{s}_{(-j)j}, \mathbf{C}\right), \end{aligned}$$

where the Gamma distribution is indicated by $\text{Ga}(\cdot, \cdot)$ and $\mathbf{C} = (s_{jj}\boldsymbol{\Omega}_{(-j)(-j)}^{-1} + \text{diag}(\mathbf{v}_{\Omega,(-j)j})^{-1})^{-1}$.⁷ Having obtained a draw for (\mathbf{u}, v) , we map them into $(\boldsymbol{\omega}_{(-j)j}, \omega_{jj})$ via:

$$\begin{aligned} \boldsymbol{\omega}_{(-j)j} &= \mathbf{u}, \\ \omega_{jj} &= v + \boldsymbol{\omega}'_{(-j)j}\boldsymbol{\Omega}_{(-j)(-j)}^{-1}\boldsymbol{\omega}_{(-j)j}. \end{aligned}$$

We perform these steps sequentially for $j = 1, \dots, M$ to obtain an update for each column and row of $\boldsymbol{\Omega}$.

5. **Sampling from** $p(\boldsymbol{\Delta}|\boldsymbol{\Omega}, \{\tau_{i,j}\}_{i,j}, \boldsymbol{\Pi}, \mathbf{z})$. Equation (2) can be used to infer δ_{ij} (Ahelegbey et al., 2016a,b). The resulting posterior distribution $p(\delta_{i,j}|\omega_{i,j}, \underline{\pi}_{i,j})$ is a Bernoulli distribution and $\boldsymbol{\Delta}$ can be constructed by sampling from $p(\delta_{i,j}|\omega_{i,j}, \underline{\pi}_{i,j}) \sim \text{Bernoulli}(\bar{\pi}_{i,j})$ with

$$\bar{\pi}_{i,j} = \frac{\mathcal{N}(\omega_{i,j}|0, \tau_{i,j}^2) \cdot \underline{\pi}_{i,j}}{\mathcal{N}(\omega_{i,j}|0, \tau_{i,j}^2) \cdot \underline{\pi}_{i,j} + \mathcal{N}(\omega_{i,j}|0, (c \cdot \tau_{i,j})^2) \cdot (1 - \underline{\pi}_{i,j})}.$$

6. **Sampling from** $p(\{\tau_{i,j}^2\}_{i,j}|\boldsymbol{\Omega}, \boldsymbol{\Delta})$. We update these by sampling from their inverse Gamma distributed conditional posterior distributions. More details are provided in Appendix A.

⁷A random variable z follows a Gamma distribution, with density $p(z|\alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)}(z)^{\alpha-1}\exp(-\beta z)$.

7. **Sampling from $p(\mathbf{\Pi}, \mathbf{z}|\mathbf{\Delta})$.** We sample the group-specific edge probabilities in the group assignments in two steps. First, we sample the group assignments marginally of $\mathbf{\Pi}$ from:

$$\mathbf{z} \sim p(\mathbf{z}|\mathbf{\Delta}).$$

This step is carried out using Algorithm 1 proposed in [Legramanti et al. \(2022\)](#). More information are provided in Appendix A.

Conditional on \mathbf{z} we sample the group-specific probabilities from Beta distributions:

$$\pi_{h,k} \sim \mathcal{B}(\pi_{h,k}|a_\pi + m_{h,k}, b_\pi + \bar{m}_{h,k}), \quad \text{for } h, k = 1, \dots, H,$$

where $h, k = 1, \dots, H$ are group indicators and $m_{h,k}(\bar{m}_{h,k})$ denotes the number of links (missing links) between clusters h and k .

Notice that this step implies that we first sample from the marginal (of $\mathbf{\Pi}$) and then from the conditional (on \mathbf{z}). Hence, the ordering of the steps matter to obtain a valid draw from $p(\mathbf{\Pi}, \mathbf{z}|\mathbf{\Delta})$.

After a sufficiently long burn-in period, this algorithm yields draws from $p(\mathbf{a}, \boldsymbol{\xi}, \{\mathbf{d}_t\}_{t=1}^T, \boldsymbol{\Omega}, \mathbf{\Delta}, \{\tau_{i,j}\}_{i,j}, \mathbf{\Pi}, \mathbf{z}|\mathbf{Y})$. In our simulation exercises and empirical work we repeat the sampler 15,000 times and discard the first 5,000 iterations as burn-in. Due to space constraints and to further reduce the autocorrelation of the retained draws we keep only every other draw.

3 Simulation-based evidence

In this section we analyze how our approach performs in a controlled environment. We generate data using a VAR(1) model:

$$\mathbf{y}_t = \mathbf{A}\mathbf{y}_{t-1} + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}_M, \boldsymbol{\Sigma}_t), \quad \boldsymbol{\Sigma}_t = \mathbf{D}_t\boldsymbol{\Sigma}\mathbf{D}_t, \quad t = 1, \dots, T$$

and $\mathbf{y}_0 = \mathbf{0}_M$. We assume that each element of $d_{j,t}$ ($j = 1, \dots, M$) evolves according to an AR(1) process $d_{j,t} = 0.9d_{j,t-1} + \sqrt{0.2}\nu_{j,t}$ with $\nu_{j,t} \sim \mathcal{N}(0, 1)$ and $d_{j,0} = 0$. This VAR(1) DGP is then set up so as to differ along several dimensions. First, we analyze model performance for differently sized models. The different sizes are $M \in \{5, 30, 50\}$, covering small, medium and large-sized datasets. Second, we consider the role of the length of the sample T . Here, we set $T \in \{200, 300, 400, 500\}$ to reflect situations commonly encountered when working with quarterly, post-war US data (i.e. so that T is between 200 and 300) as well as monthly series (with T being 400 or 500). Third, the DGPs assume a network structure so that the true matrix $\mathbf{\Omega}$ under the DGP has a particular structure which the SBM can uncover. The true $\mathbf{\Omega}$ is generated in the following way: Each of the variables is assigned to a cluster (with the number of true clusters depending on M). We draw the within-group edge probabilities from the Beta distribution $\text{Beta}(1, 100)$ and the cross-group edge probabilities from $\text{Beta}(100, 1)$. The elements of the adjacency matrix are generated by drawing from a Bernoulli distribution with the respective edge probability between variables i and j as success probability. Finally, we generate $\mathbf{\Omega}$ as a positive definite matrix with zero restrictions given by the adjacency matrix. We also consider the case of no network structure. In this case, the edge probabilities are set equal to 0.2.

This gives us a combination of 24 different DGPs. For each of these 24 DGPs we generate 25 datasets and evaluate the performance in terms of hit rates (HRs), i.e., correctly classified edges in the network, of our BVAR-SBM against a benchmark model with a standard SSVS prior on the covariances (henceforth labeled SSVS).

Our SBM-VARs differ in the clustering priors used. We consider each of the four Gibbs-type priors mentioned in Section 2.2.1: the GN, the DM, the DP and the PY prior. Their hyper parameters are specified such that the prior expected number of groups is equal to 1.5 times the true number of groups (rounded to whole numbers).

We start out by discussing the absolute HRs of the posterior median $\bar{\mathbf{\Delta}}$ of $p(\mathbf{\Delta}|\mathbf{Y})$

to the true value of Δ . Before we zoom into the DGP-specific performance, we consider the overall performance over different values of T, M and across the realizations from the DGP. These are, in the form of histograms, shown in Figure 2 for the DGPs where Ω features a network structure. In this figure, each of the histograms refers to one of the SBM priors (and the SSVS specification). The means of these histograms show the average quality of the network approximation (measured in HRs) over the space of DGPs.




Figure 2: Percentage of correctly classified edges of different SBM-VARs and the SSVS-VAR (red) over all draws from all DGPs with underlying network structure.




Figure 3: Percentage of correctly classified edges of different SBM-VARs and the SSVS-VAR (red) over all draws from all DGPs with no underlying network structure.

The figure reveals that the different SBM priors all do well in terms of recovering the true network structure. For all models considered, the mean of the distribution of HRs is around 95 percent. By contrast, the SSVS prior that assumes the prior inclusion probability as fixed across edges is performing significantly worse, with mean HRs that are

appreciably lower than the ones of the SBMs. Notice that the left-tail of the distributions coincide, implying that all methods (including the SSVS) produce estimates of the network that have the same lower bound on the HRs.

Next, we turn to the case where Ω features no clustering. This is shown in Figure 3. In this case, all SBM-based models produce hit rates that are close to the SSVS models. HRs are, however, markedly lower in all cases relative to the DGP that features clusters in Ω . Nevertheless, these results can be interpreted in a way that adding the SBM to the VAR does not harm network-detection accuracy if there is no particular clustering structure and, if there is clustering, it improves accuracy appreciably.

M	T	Clustered DGPs					Non-clustered DGPs				
		GN	DM	DP	PY	SSVS	GN	DM	DP	PY	SSVS
5	200	2.78	2.78	2.22	3.33	90.00	0.00	-1.00	-2.00	-1.00	93.50
	30	4.00	4.14	3.93	3.94	88.74	-0.60	-1.02	-0.76	-0.69	86.73
	50	4.06	5.71	5.84	4.23	85.03	-0.11	-0.26	-0.22	-0.01	83.99
5	300	2.22	2.22	2.22	2.22	92.78	-0.43	-0.43	0.00	-0.87	93.04
	30	3.69	4.04	3.95	3.61	90.24	-1.31	-1.66	-1.61	-1.10	86.84
	50	5.47	6.83	6.13	5.69	87.15	-0.50	-0.70	-0.63	-0.38	85.92
5	400	0.62	-0.62	0.00	-0.62	94.38	0.45	0.45	0.45	0.45	93.18
	30	2.79	2.99	2.98	2.81	91.74	-3.01	-2.47	-2.81	-2.72	85.43
	50	5.72	6.79	6.45	5.62	88.56	-1.20	-1.47	-0.91	-0.94	86.56
5	500	0.53	0.53	0.53	0.53	93.68	-0.43	0.00	0.00	0.00	93.48
	30	2.22	2.30	2.43	2.41	92.04	-2.79	-2.81	-2.72	-2.84	83.35
	50	5.61	6.75	6.04	5.29	89.64	-1.38	-1.24	-1.28	-1.33	86.42

Table 1: Differences in hit rates between the different SBM-VARs and the SSVS-VAR. All results are means over 25 simulations from each DGP. For the SSVS-VAR absolute hit rates are reported.

The histograms mask differences across different parameters of the DGPs. To understand under which circumstances our SBMs produce more favorable estimates of the network, we now zoom into the specific performance for each choice of M, T and whether the DGP has a network structure or not. All the results are provided in Table 1. The table shows the percentage point differences between the HR of a particular

SBM-VAR and the SSVS-VAR.

We start our discussion by focusing on the absolute performance of the SSVS-VAR first. For all combinations of T , M and for clustering and non-clustering DGPs, the SSVS-VAR produces hit rates well above 80 percent. For small models ($M = 5$), the model detects around 90 or more edges correctly. This is not surprising given the fact that the number of possible edges is 10 and thus potential clusters are (at best) characterized by including relatively few nodes. When we increase the dimensionality of the model (i.e. set $M \geq 15$), the performance of the SSVS-VAR deteriorates and drops below 90 percent. Notice that using longer time series ($T \geq 300$) improves network estimation accuracy across all model sizes.

The SBM-VARs with different priors on the clustering behavior mostly improve upon the SSVS-VAR if the DGP features clustering. In most cases, differences across priors are small and within one to two percentage points. When we consider the non-clustered DGP we find many values below zero. But these are small, suggesting that using the SBM does not hurt network detection accuracy if no clustering is present.

Zooming into the differences across priors for the non-clustered DGP we find that for some priors and DGPs, using an SBM produces (almost) the same HRs as in the case of the correctly-specified SSVS-VAR.

To sum up, our simulations indicate that the choice of the prior on the clustering process is not critical. Nevertheless, we find that the DP is very often producing the highest HR for both DGPs.

4 Macroeconomic forecasting using SBM-VARs

The previous section shows that, with artificial data, our model is capable of recovering the underlying network of shocks accurately. Next, we ask whether explicitly modeling a network among the shocks also translates into more accurate density forecasts. We do so

by using the different SBM-VARs to forecast US macroeconomic variables.

4.1 Data and design of the forecasting exercise

We rely on the quarterly version of the [McCracken and Ng \(2016\)](#) database from 1960:Q1 through 2023Q2. We consider $h = 1$ and $h = 4$ quarter ahead forecasts, calculated iteratively. Our forecast evaluation period starts in 1990:Q1.

We consider two model sizes: a medium-sized dataset with $M = 10$ and a larger one with $M = 19$ endogenous variables. Both datasets include a set of three focus variables: these are the unemployment rate (UNRATE), the GDP growth rate (GDPC1) and CPI inflation (CPIAUCSL). Then, depending on the dataset we add additional macroeconomic and financial series. Further information on the actual series included in each model and how they are transformed are provided in the Data Appendix.

To understand whether the prior on \mathbf{z} matters for forecast accuracy, we again consider all four SBM priors and use the acronym SBM-VAR-X, where $X \in \{\text{GN, DM, DP, PY}\}$, to denote them. We also include the SSVS prior (without a network structure). All forecasting results are benchmarked against the Bayesian VAR without shrinkage on $\mathbf{\Omega}$ (i.e. a model that sets $\delta_{i,j} = 1 \forall i, j$). This model is henceforth called the no shrinkage benchmark or baseline model. Thus we can separate out the gains from using some sort of SSVS prior shrinkage (by comparing SSVS results to the no shrinkage benchmark) from the gains from using our specific SSVS network prior (by comparing the SBM-VAR-X models to SSVS). We stress that all these models share the same prior for the VAR coefficients and other model parameters and differ only in the prior for $\mathbf{\Omega}$.

Forecast accuracy is measured through log predictive likelihoods (LPLs). Since our goal is to model a network, we do not only focus on the univariate LPLs of the three focus variables but on joint LPLs. These joint LPLs either measure the joint forecast density performance for the three focus variables (FOCUS) or for all variables in \mathbf{y}_t (ALL).

Before we discuss the empirical results, a brief word on how we set up the priors for

z . Note that the FRED-QD data divides the variables into groups (e.g. employment and unemployment, housing, etc.) and our data sets choose variables from some groups, but not others. The hyperparameters on the Gibbs priors are set so that the prior expectation on the number of groups equals the number of non-empty groups (for a particular model size). For instance, if the data set only includes labor market indicators and interest rates, we would tune the prior to expect only two clusters.

4.2 Forecasting results

4.2.1 Overall density forecasting performance

We start by discussing the overall forecast performance of the various VARs first. Table 2 shows the average LPLs of a particular model minus the one of the VAR without shrinkage on Ω . Positive numbers suggest that our model does better than the benchmark while negative numbers indicate the opposite. The column associated with 'BASE' shows the absolute LPL for the baseline specification.

We start by discussing the one-step-ahead marginal LPLs of the medium-scale VAR first. For two out of three focus variables (GDPC1 and CPIAUCSL), we find that SBM-VARs and the SSVS-VAR improve upon the no shrinkage benchmark. In these cases, differences between the SSVS-VAR (which assumes no network structure in the prior inclusion probabilities) and the SBMs are muted. In particular, we find that for GDP growth, all models produce a similar performance, with the SBM-VAR-DM and the SSVS-VAR producing identical LPL differences. For inflation, a similar pattern arises but here we find that the SBM-VAR-DP produces slightly more precise density forecasts than all competing models. Only for the unemployment rate we find that SBM-VAR-GN produces substantially more precise density forecasts than all other specifications. The unemployment rate is also the only focus variable where some of the SBM-VARs (and the SSVS-VAR) perform worse than the no shrinkage benchmark.

Variable Group	Medium VAR						Large VAR					
	GN	DM	DP	PY	SSVS	BASE	GN	DM	DP	PY	SSVS	BASE
1-quarter ahead												
GDPC1	0.11	0.12	0.10	0.11	0.12	3.00	0.06	0.01	0.00	0.03	0.00	2.98
UNRATE	0.24	0.12	-0.08	-0.20	-0.07	-7.62	0.73	-0.56	-0.68	-0.08	-1.57	-7.36
CPIAUCSL	0.11	0.11	0.14	0.08	0.11	3.78	0.03	0.08	0.08	0.08	0.08	3.78
FOCUS	0.73	0.45	0.29	0.16	0.35	-0.87	1.19	-0.12	-0.34	0.40	-1.22	-0.51
ALL	0.99	0.73	0.64	0.47	0.51		-0.29	-3.51	-0.93	-4.12	-4.52	
4-quarters ahead												
GDPC1	0.04	0.01	0.04	0.06	0.03	2.50	0.00	-0.01	0.01	0.00	0.00	2.57
UNRATE	0.11	0.29	0.10	0.15	-0.01	-6.15	0.53	0.21	0.44	0.44	1.07	-6.52
CPIAUCSL	0.12	0.09	0.11	0.11	0.08	3.54	0.16	0.10	0.14	0.13	0.14	3.49
FOCUS	0.37	0.52	0.34	0.38	0.18	0.92	0.85	0.52	0.80	0.74	1.52	0.36
ALL	0.58	0.63	0.63	0.63	0.38		2.42	2.45	3.53	2.81	1.91	

Table 2: Joint LPL scores across different variable groups of the medium-sized and large VARs relative to the baseline model with no shrinkage. Scores of the baseline are in absolute terms.

For the large model, a similar story arises. We find small gains for output and inflation. These, however, appear to be more pronounced when compared to the SSVS-VAR. For unemployment, we find that the only model that improves upon the no shrinkage VAR is the SBM-VAR-GN with all other specifications performing worse than the benchmark. Comparing the marginal LPLs for the medium and large baseline model shows almost no differences.

Intuitively, if our focus is on the marginal LPL for an individual variable, then the only way the SBM-VARs can show performance gains relative to SSVS is if they are better at modeling the contemporaneous spillovers between variables. Thus, it is unsurprising there are only small differences in marginal LPLs. Joint LPLs should profit even more given that they depend on the full predictive covariance matrix Σ_{t+h} .

To back this claim, we start by discussing the joint LPL over the three focus variables. For the medium model, we find gains vis-à-vis the benchmark that are more sizable than the ones for any of the three focus variables. Again, the best performing specification is

the SBM-VAR-GN with SBM-VAR-DM being second. In most cases (except for SBM-VAR-PY) we find that using the SBM to model Ω pays off if the target is the joint density of the three focus series. When we consider the larger models we find a somewhat different story. In this case, two out of four SBMs lose against the benchmark. But we also find that SBM-VAR-GN, again, substantially improves upon the benchmark and all other specifications (in particular the SSVS-VAR). This has to be considered in light of the absolute LPLs for the focus variables of the baseline model. Since these are slightly higher than the ones of the medium-scale VAR without shrinkage, we can say that if the researcher wishes to produce the most precise density forecasts for all three focus variables jointly, the single best performing specification is SBM-VAR-GN.

Next, we consider the joint performance for all series in \mathbf{y}_t . If all series are considered, a similar pattern to the one based on the joint LPLs over the focus variables arises. With the medium-sized data set, we find gains with respect to the benchmark with SBM-VAR-GN producing the most precise joint density predictions. For the large data set, interestingly, none of the models that induce shrinkage on Ω manages to outperform the no shrinkage benchmark. In this case, the SBM-VAR-GN produces similar density predictions (which are, nevertheless, slightly worse than the ones of the benchmark).

Focusing on the one-year-ahead forecast distributions, a similar picture to the one-step-ahead horizon emerges. For marginals, the gains are rather small for output and inflation whereas they appear to be more substantial for the unemployment rate. This pattern holds for both model sizes. When we consider the LPLs over the focus series we again find SBM-VAR-GN to produce a strong performance, improving upon the benchmark and the SSVS-VAR. Notice, however, that the other priors perform slightly better. But these differences are muted. This also holds for the joint LPLs and when the medium-scale dataset is considered.

For the large dataset, however, we find a different picture to the one observed for the one-quarter-ahead forecasts. In this case, we find that SBM-based models yield LPLs that




Figure 4: Recursive mean of joint differences in one-quarter ahead LPLS scores relative to a VAR with a diffuse prior on the precision matrix. Shaded areas denote NBER reference recessions.

are appreciably larger than the ones of the benchmark and the SBM-VAR-DP produces the most accurate density forecasts.

4.2.2 Density forecast performance over time

At this point, we have discussed only average forecast performance over all periods in the hold out. To better understand why particular models are doing well we turn to analyzing the forecasting performance over time.

We do so in Figure 4. This figure shows differences in recursive means between the different VARs that induce shrinkage on Ω and the no shrinkage benchmark. At each point in time, a number greater than zero implies that a particular model has been

outperforming the benchmark up to this point. For brevity, we focus on the joint density over the focus variables and the joint density over all elements in \mathbf{y}_t only.

Figure 4 strikingly shows that using models that induce shrinkage on $\mathbf{\Omega}$ helps forecast performance during volatile periods such as the global financial crisis or the pandemic. During both recessions, our models improve appreciably over the benchmark specification. In tranquil periods, this effect is much less pronounced and we even find that after the financial crisis, the relative differences somewhat decline up until 2020:Q1.

This general pattern holds for both datasets and if we consider LPLs over the focus series and over all variables in the system. Notice, however, that the evolution of the relative LPLs is very similar for panels (a) and (b) for the medium data set, but differences are much larger if we compare panels (a) and (b) for the large dataset. In this case, two of the four SBM models (GN, DP) and the SSVS-VAR perform poorly during the pandemic (see panel (a)) but all of them fail to improve upon the benchmark if we consider joint LPLs over all variables in the system.

This discussion (and the discussion in the previous section) shows that our SBM-VARs are capable of improving upon the model that induces no shrinkage on $\mathbf{\Omega}$ and the VAR which assumes that the prior inclusion probabilities are fixed and independent from each other. Within the class of SBM-VARs, we find that SBM-VAR-GN does best for the one-step-ahead horizon while SBM-VAR-DP works best for four-quarter-ahead forecasts. In most cases, however, differences across the priors on \mathbf{z} are rather small.

4.3 Inspecting network properties

In this section, we investigate the properties of the estimated networks of the SBM-VARs with the GN and DP priors and the SSVS-VAR. Our focus on SBM-BVAR-GN and SBM-BVAR-DP is predicated on their strong performance for one and four-step-ahead predictive likelihoods, respectively. We analyze the networks by considering a set of summary statistics related to the network structure. These are computed recursively

over the hold-out period. For brevity, we focus on the medium-scale model.

We consider three different summary statistics:

1. *Number of groups.* This measures the number of groups H detected by the SBM over the hold-out period. As a point estimate of H , we use the clustering which minimizes the posterior expectation of the variance of information loss function. [Wade and Ghahramani \(2018\)](#) discuss the advantageous properties of this approach.
2. *Average degree.* The degree for each node i is the number of edges between i and all other nodes:

$$d_i = \sum_{j \neq i} \delta_{i,j}.$$

The average degree is given by: $\bar{d} = 1/M \sum_{i=1}^M d_i$. The average degree measures the average number of links of a particular node to the other nodes in the network.

3. *Modularity.* Modularity is a measure of how strongly separated from one another the different clusters are. It is defined as:

$$Q = \frac{1}{2N} \sum_{i,j} \left(\delta_{i,j} - \frac{d_i d_j}{N} \right) \mathbb{1}(z_i = z_j),$$

where N is the overall number of edges and $\mathbb{1}(z_i = z_j)$ is the indicator function that takes on value 1 if nodes i and j are in the same groups, that is if $z_i = z_j$. If Q is large, most edges in the network are within a particular cluster whereas if Q is low, the community structure is weak and the edges are distributed uniformly over the network.

Figure 5 shows how these network summary metrics evolve over the hold-out period. We start by discussing how the number of groups evolves over time. Panel 5a suggests that the model with the GN prior produces only a single cluster up until around 2012. Afterwards, the number of clusters fluctuates between one and two. By contrast, when we

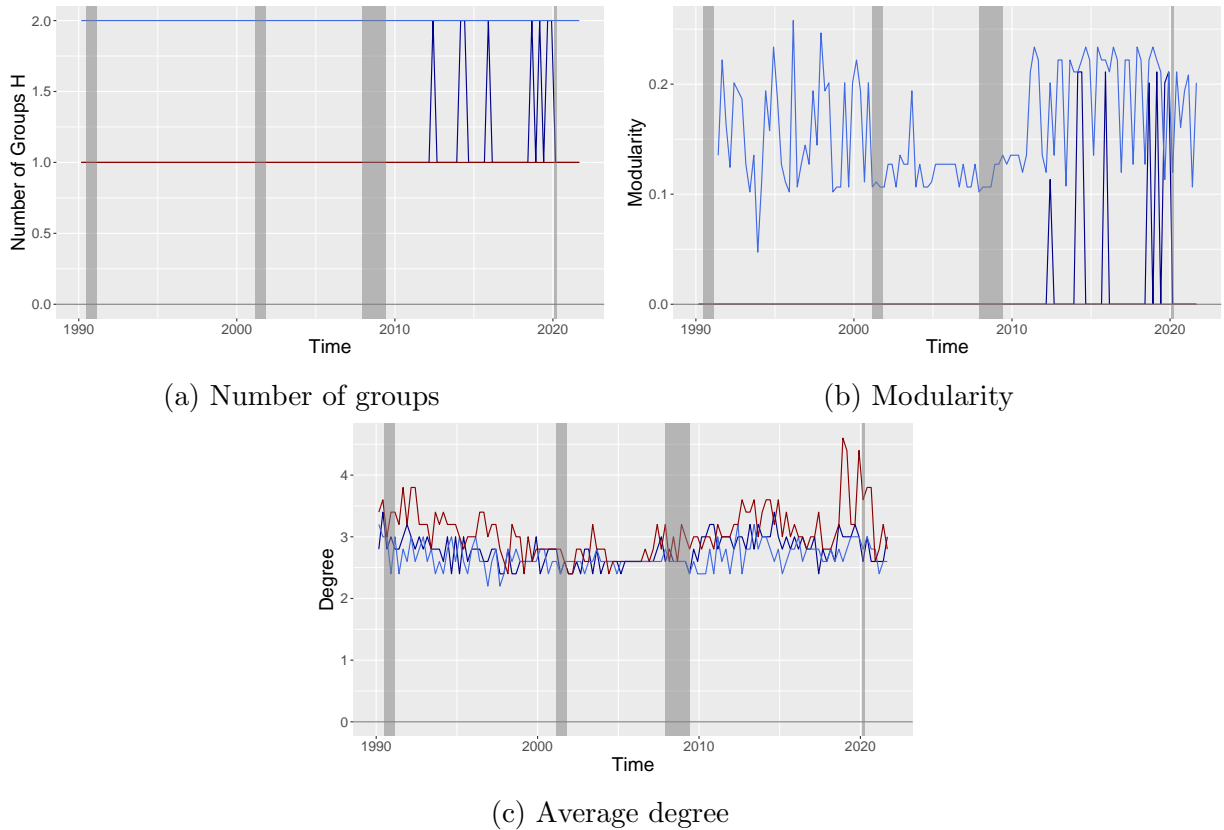


Figure 5: Network summary statistics over the holdout period. Dark blue is the SBM-GN prior, light blue the SBM-DP and red the SSVS prior. Shaded areas denote NBER reference recessions.

consider the DP prior, which favors more dispersed partitions and tends to open up new clusters more quickly than the GN prior, the figure suggests two clusters throughout. The first cluster is comprised of price series and short-term interest rates whereas the second cluster consists of the remaining series in the system.

Next, we consider network modularity, displayed in panel 5b. Again, considering the GN prior first shows that modularity increases in lockstep with the number of groups. It is noteworthy that modularity is appreciably lower in 2012 compared to the instances where the number of groups reach two in the later part of the hold-out period. This indicates that the structure of the groups has changed in the sense that community structure has been weaker in 2012 so that shocks within a group had a tendency to form more ties to shocks in the other group whereas for the cases where find two regimes later in the

sample the connections have been strong within a particular group but less so between the two groups. Under the DP prior, we find that during the 1990s up until the burst of the dotcom bubble modularity is higher and decreases in the period between the dotcom crisis and the global financial crisis. During and after the financial crisis we find that cross-group relations become stronger.

Finally, we consider the average degree over time. This is shown in panel 5c. We find that the network implied under the standard SSVS prior tends to be more dense than under the two SBM-type priors. For all three models, we observe that the degree is higher in the first third of the hold-out, before declining slightly during the early 2000s up to the financial crisis. In the aftermath of the financial crisis up to the pandemic the implied networks across all three priors tend to become more dense. This is particularly visible for the SSVS-VAR which gives rise to a rather dense Ω matrix.

5 Conclusions

Prior information can be important to ensure shrinkage and parsimony in large VARs. A substantial Bayesian VAR literature has developed proposing various priors, but few of them focus on the high-dimensional error covariance matrix. In this paper we develop such a prior based on the idea that the error precision matrix is likely to exhibit a parsimonious network structure. It does so by using an SSVS prior (to select the parsimonious specification) with a SBM (which ensures the specifications involve networks). We also develop a computationally efficient Bayesian MCMC algorithm which jointly estimates the number of clusters and the members in each cluster.

In simulated data we find that this algorithm does a fine job of uncovering the network structure whereas a standard SSVS prior, designed to ensure parsimony but not involving a network structure, does not perform as well. This ability to find a parsimonious network in a high dimensional error precision matrix is also found in our forecasting exercise

involving US macroeconomic data. In particular, we find a small number of clusters with sensible interpretations (e.g. grouping price and interest rate variables in one cluster and the remaining variables reflecting the real economy in another) but noticeable variation over time both in the membership of each cluster and in the strength of the links between members. Our forecasting exercise shows that the use of our SBM prior typically leads to some modest improvements over both an SSVS prior which ignores the network structure and a non-informative prior.

As possible avenues for further research we stress that our model is very general and can be applied in combination with any specification for the conditional mean. Moreover, the information we back out on the network can be used to inform shrinkage priors on the lagged coefficients as well.

From an applied perspective, we illustrate our approach using US macroeconomic data. But in principle, we can use our framework to model financial time series such as stock returns which display a more pronounced clustering or herding behavior (see, e.g. [Tsionas et al., 2022](#), for a multivariate SV model that dynamically clusters series so that they exhibit herding/non-herding behavior).

References

- Daniel Felix Ahelegbey, Monica Billio, and Roberto Casarin. Sparse graphical vector autoregression: A bayesian approach. *Annals of Economics and Statistics*, (123/124): 333–361, 2016a. doi: <https://doi.org/10.15609/annaeconstat2009.123-124.0333>.
- Daniel Felix Ahelegbey, Monica Billio, and Roberto Casarin. Bayesian graphical models for structural vector autoregressive processes. *Journal of Applied Econometrics*, 31(2): 357–386, 2016b. doi: <https://doi.org/10.1002/jae.2443>.
- Ilya Archakov and Peter Reinhard Hansen. A new parametrization of correlation matrices. *Econometrica*, 89(4):1699–1715, 2021. doi: <https://doi.org/10.3982/ECTA16910>.
- Jonas E. Arias, Juan F. Rubio-Ramírez, and Minchul Shin. Macroeconomic forecasting and variable ordering in multivariate stochastic volatility models. *Journal of Econometrics*, 235(2):1054–1086, 2023. doi: <https://doi.org/10.1016/j.jeconom.2022.04.013>.
- Yu Bai, Andrea Carriero, Todd E. Clark, and Massimiliano Marcellino. Macroeconomic forecasting in a multi-country context. *Journal of Applied Econometrics*, 37(6):1230–1255, 2022. doi: <https://doi.org/10.1002/jae.2923>.
- Marta Bańbura, Domenico Giannone, and Lucrezia Reichlin. Large bayesian vector autoregressions. *Journal of Applied Econometrics*, 25(1):71–92, 2010. doi: <https://doi.org/10.1002/jae.1137>.
- Monica Billio, Roberto Casarin, and Luca Rossini. Bayesian nonparametric sparse var models. *Journal of Econometrics*, 212(1):97–115, 2019. doi: <https://doi.org/10.1016/j.jeconom.2019.04.022>.
- Pierpaolo De Blasi, Antonio Lijoi, and Igor Prünster. An asymptotic analysis of a class

- of discrete nonparametric priors. *Statistica Sinica*, 23(3):1299–1321, 2013. doi: <https://doi.org/10.5705/ss.2012.047>.
- Andrea Carriero, George Kapetanios, and Massimiliano Marcellino. Forecasting exchange rates with a large bayesian var. *International Journal of Forecasting*, 25(2):400–417, 2009. doi: <https://doi.org/10.1016/j.ijforecast.2009.01.007>.
- Andrea Carriero, Todd E. Clark, and Massimiliano Marcellino. Large bayesian vector autoregressions with stochastic volatility and non-conjugate priors. *Journal of Econometrics*, 212(1):137–154, 2019. doi: <https://doi.org/10.1016/j.jeconom.2019.04.024>.
- Andrea Carriero, Joshua Chan, Todd E. Clark, and Massimiliano Marcellino. Corrigendum to “large bayesian vector autoregressions with stochastic volatility and non-conjugate priors” [j. econometrics 212 (1) (2019) 137–154]. *Journal of Econometrics*, 227(2):506–512, 2022. doi: <https://doi.org/10.1016/j.jeconom.2021.11.010>.
- Carlos M. Carvalho, Nicholas G. Polson, and James G. Scott. The horseshoe estimator for sparse signals. *Biometrika*, 97(2):465–480, 04 2010. doi: <https://doi.org/10.1093/biomet/asq017>.
- Joshua Chan. Minnesota-type adaptive hierarchical priors for large bayesian vars. *International Journal of Forecasting*, 37(3):1212–1226, 2021. ISSN 0169-2070. doi: <https://doi.org/10.1016/j.ijforecast.2021.01.002>.
- Joshua CC Chan, Gary Koop, and Xuewen Yu. Large order-invariant bayesian vars with stochastic volatility. *Journal of Business & Economic Statistics*, 42(2):825–837, 2024.
- Timothy Cogley and Thomas J. Sargent. Drifts and volatilities: monetary policies and outcomes in the post wwii us. *Review of Economic Dynamics*, 8(2):262–302, 2005.

ISSN 1094-2025. doi: <https://doi.org/10.1016/j.red.2004.10.009>. Monetary Policy and Learning.

Pierpaolo De Blasi, Stefano Favaro, Antonio Lijoi, Ramsés H. Mena, Igor Prünster, and Matteo Ruggiero. Are gibbs-type priors the most natural generalization of the dirichlet process? *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(2): 212–229, 2015. doi: <https://doi.org/10.1109/TPAMI.2013.217>.

Thomas Doan, Robert B. Litterman, and Christopher Sims. Forecasting and conditional projection using realistic prior distributions. *Econometric Reviews*, 3(1):1–100, 1984. doi: <https://doi.org/10.1080/07474938408800053>.

Jianqing Fan, Yuan Liao, and Han Liu. An overview of the estimation of large covariance and precision matrices. *The Econometrics Journal*, 19(1):C1–C32, 2016.

Martin Feldkircher, Florian Huber, Gary Koop, and Michael Pfarrhofer. Approximate bayesian inference and forecasting in huge-dimensional multicountry vars. *International Economic Review*, 63(4):1625–1658, 2022. doi: <https://doi.org/10.1111/iere.12577>.

Lendie Follett and Cindy Yu. Achieving parsimony in bayesian vector autoregressions with the horseshoe prior. *Econometrics and Statistics*, 11:130–144, 2019. ISSN 2452-3062. doi: <https://doi.org/10.1016/j.ecosta.2018.12.004>.

Junxian Geng, Anirban Bhattacharya, and Debdeep Pati. Probabilistic community detection with unknown number of communities. *Journal of the American Statistical Association*, 114(526):893–905, 2019. doi: <https://doi.org/10.1080/01621459.2018.1458618>.

Edward I. George and Robert E. McCulloch. Variable selection via gibbs sampling. *Journal of the American Statistical Association*, 88(423):881–889, 1993. doi: <https://doi.org/10.1080/01621459.1993.10476353>.

- Edward I. George, Dongchu Sun, and Shawn Ni. Bayesian stochastic search for var model restrictions. *Journal of Econometrics*, 142(1):553–580, 2008. doi: <https://doi.org/10.1016/j.eurocorev.2015.09.006>.
- Domenico Giannone, Michele Lenza, Daphne Momferatou, and Luca Onorante. Short-term inflation projections: A bayesian vector autoregressive approach. *International Journal of Forecasting*, 30(3):635–644, 2014. doi: <https://doi.org/10.1016/j.ijforecast.2013.01.012>.
- Domenico Giannone, Michele Lenza, and Giorgio E. Primiceri. Prior selection for vector autoregressions. *The Review of Economics and Statistics*, 97(2):436–451, 2015. doi: https://doi.org/10.1162/REST_a_00483.
- Alexander Gnedin. A species sampling model with finitely many types. *Electronic Communications in Probability*, 15:79 – 88, 2010. doi: <https://doi.org/10.1214/ECP.v15-1532>.
- Paul W. Holland, Kathryn Blackmond Laskey, and Samuel Leinhardt. Stochastic blockmodels: First steps. *Social Networks*, 5(2):109–137, 1983. doi: [https://doi.org/10.1016/0378-8733\(83\)90021-7](https://doi.org/10.1016/0378-8733(83)90021-7).
- Florian Huber and Martin Feldkircher. Adaptive shrinkage in bayesian vector autoregressive models. *Journal of Business & Economic Statistics*, 37(1):27–39, 2019. doi: <https://doi.org/10.1080/07350015.2016.1256217>.
- Florian Huber and Gary Koop. Fast and order-invariant inference in bayesian vars with non-parametric shocks. *Journal of Applied Econometrics*, forthcoming, 2024.
- Tsunehiro Ishihara and Yasuhiro Omori. Efficient bayesian estimation of a multivariate stochastic volatility model with cross leverage and heavy-tailed errors. *Computational Statistics & Data Analysis*, 56(11):3674–3689, 2012. doi: <https://doi.org/10.1016/j.>

csda.2010.07.015. 1st issue of the Annals of Computational and Financial Econometrics Sixth Special Issue on Computational Econometrics.

Hemant Ishwaran and J. Sunil Rao. Spike and slab variable selection: Frequentist and Bayesian strategies. *The Annals of Statistics*, 33(2):730 – 773, 2005. doi: <https://doi.org/10.1214/009053604000001147>.

K. Rao Kadiyala and Sune Karlsson. Numerical methods for estimation and inference in bayesian var models. *Journal of Applied Econometrics*, 12(2):99–132, 1997. doi: [https://doi.org/10.1002/\(SICI\)1099-1255\(199703\)12:2%3C99::AID-JAE429%3E3.0.CO;2-A](https://doi.org/10.1002/(SICI)1099-1255(199703)12:2%3C99::AID-JAE429%3E3.0.CO;2-A).

Gregor Kastner and Florian Huber. Sparse bayesian vector autoregressions in huge dimensions. *Journal of Forecasting*, 39(7):1142–1165, 2020. doi: <https://doi.org/10.1002/for.2680>.

Charles Kemp, Joshua B. Tenenbaum, Thomas L. Griffiths, Takeshi Yamada, and Naonori Ueda. Learning systems of concepts with an infinite relational model. In *Proceedings of the 21st National Conference on Artificial Intelligence and the 18th Innovative Applications of Artificial Intelligence Conference, AAAI-06/IAAI-06*, Proceedings of the National Conference on Artificial Intelligence, pages 381–388, 2006.

Gary Koop. Forecasting with medium and large bayesian vars. *Journal of Applied Econometrics*, 28(2):177–203, 2013. doi: <https://doi.org/10.1002/jae.1270>.

Dimitris Korobilis. Var forecasting using bayesian variable selection. *Journal of Applied Econometrics*, 28(2):204–230, 2013. doi: <https://doi.org/10.1002/jae.1271>.

Sirio Legramanti, Tommaso Rigon, Daniele Durante, and David B. Dunson. Extended stochastic block models with application to criminal networks. *The Annals of Applied Statistics*, 16(4):2369 – 2395, 2022. doi: <https://doi.org/10.1214/21-AOAS1595>.

- Antonio Lijoi, Ramsés H. Mena, and Igor Prünster. Controlling the reinforcement in bayesian non-parametric mixture models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69(4):715–740, 2007a. doi: <https://doi.org/10.1111/j.1467-9868.2007.00609.x>.
- Antonio Lijoi, Ramsés H. Mena, and Igor Prünster. Bayesian nonparametric estimation of the probability of discovering new species. *Biometrika*, 94(4):769–786, 2007b. doi: <https://doi.org/10.1093/biomet/asm061>.
- Robert B. Litterman. Forecasting with bayesian vector autoregressions—five years of experience. *Journal of Business & Economic Statistics*, 4(1):25–38, 1986. doi: <https://doi.org/10.1080/07350015.1986.10509491>.
- Enes Makalic and Daniel F. Schmidt. A simple sampler for the horseshoe estimator. *IEEE Signal Processing Letters*, 23(1):179–182, 2016. doi: [10.1109/LSP.2015.2503725](https://doi.org/10.1109/LSP.2015.2503725).
- Michael W. McCracken and Serena Ng. Fred-md: A monthly database for macroeconomic research. *Journal of Business & Economic Statistics*, 34(4):574–589, 2016. doi: <https://doi.org/10.1080/07350015.2015.1086655>.
- Jeffrey W. Miller and Matthew T. Harrison. Mixture models with a prior on the number of components. *Journal of the American Statistical Association*, 113(521):340–356, 2018. doi: <https://doi.org/10.1080/01621459.2016.1255636>.
- Krzysztof Nowicki and Tom A. B Snijders. Estimation and prediction for stochastic block structures. *Journal of the American Statistical Association*, 96(455):1077–1087, 2001. doi: <https://doi.org/10.1198/016214501753208735>.
- Jim Pitman and Marc Yor. The two-parameter Poisson-Dirichlet distribution derived from a stable subordinator. *The Annals of Probability*, 25(2):855 – 900, 1997. doi: [10.1214/aop/1024404422](https://doi.org/10.1214/aop/1024404422).

Mikkel N. Schmidt and Morten Morup. Nonparametric bayesian modeling of complex networks: an introduction. *IEEE Signal Processing Magazine*, 30(3):110–128, 2013. doi: <https://doi.org/10.1109/MSP.2012.2235191>.

Mike G Tsionas, Dionisis Philippas, and Nikolaos Philippas. Multivariate stochastic volatility for herding detection: Evidence from the energy sector. *Energy Economics*, 109:105964, 2022.

Sara Wade and Zoubin Ghahramani. Bayesian cluster analysis: Point estimation and credible balls (with discussion). *Bayesian Analysis*, 13(2):559 – 626, 2018. doi: 10.1214/17-BA1073.

Hao Wang. Bayesian Graphical Lasso Models and Efficient Posterior Computation. *Bayesian Analysis*, 7(4):867 – 886, 2012. doi: 10.1214/12-BA729.

Hao Wang. Scaling it up: Stochastic search structure learning in graphical models. *Bayesian Analysis*, 10(2):351 – 377, 2015. doi: <https://doi.org/10.1214/14-BA916>.

A Full Conditional Posterior Distributions

Sampling from $p(\mathbf{a}|\boldsymbol{\xi}, \{\mathbf{d}_t\}_{t=1}^T, \boldsymbol{\Omega}, \mathbf{Y})$. We let $\mathbf{x}_t = (\mathbf{y}_{t-1}, \dots, \mathbf{y}_{t-P})$ and stack the lagged observations over time to arrive at the $T \times (M P)$ -dimensional matrix $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_T)'$. Furthermore, \mathbf{B}_0 is the lower-Cholesky factor of $\boldsymbol{\Omega}$ and $\tilde{\mathbf{D}} = (\mathbf{d}_1, \dots, \mathbf{d}_T)'$ a $T \times M$ -matrix containing the sequence of time-varying observation error volatilities. [Carriero et al. \(2022\)](#) show that conditional on knowing \mathbf{B}_0 and $\tilde{\mathbf{D}}$ it is possible to factorize the model such that the VAR coefficients can be estimated equation-by-equation. In particular, for the j^{th} equation ($j = 1, \dots, M$) let:

$$\begin{aligned}\mathbf{Y}^{(j)} &= \text{vec}((\mathbf{Y} - \mathbf{X}\mathbf{A}^{[j=0]})\mathbf{B}'_{(0,j:M,\bullet)}) ./ \text{vec}(\tilde{\mathbf{D}}_{\bullet,j:M}^{0.5}), \\ \mathbf{X}^{(j)} &= (\mathbf{B}_{(0,j:M,j)} \otimes \mathbf{X}) ./ \text{vec}(\tilde{\mathbf{D}}_{\bullet,j:M}^{0.5}),\end{aligned}$$

where $\mathbf{A}^{[j=0]}$ refers to a modified version of \mathbf{A} where the j -th column is set equal to zero. The sub-matrix containing the rows with indices j through M of \mathbf{B}_0 is denoted as $\mathbf{B}_{0,j:M,\bullet}$ and the vector $\mathbf{B}_{0,j:M,j}$ contains the elements j to M of the j -th column of B_0 . Similarly, $\tilde{\mathbf{D}}_{\bullet,j:M}^{0.5}$ are the columns j through M of $\tilde{\mathbf{D}}^{0.5}$. The symbol $./$ indicates element-by-element division. Rewriting the variables in such a way, the moments of the conditional posterior arise from the standard Gaussian regression model:

$$\begin{aligned}\bar{\mathbf{V}}_{a,j} &= (\underline{\mathbf{V}}_{a,j}^{-1} + \mathbf{X}^{(j)'}\mathbf{X}^{(j)})^{-1}, \\ \bar{\mathbf{a}}_{j,\bullet} &= \bar{\mathbf{V}}_{a,j}(\underline{\mathbf{V}}_{a,j}^{-1}\underline{\mathbf{a}}_{j,\bullet} + \mathbf{X}^{(j)'}\mathbf{Y}^{(j)}),\end{aligned}$$

where $\underline{\mathbf{V}}_{a,j}$ denotes the prior variance (determined by the elements of $\boldsymbol{\xi}$) and $\underline{\mathbf{a}}_{j,\bullet}$ the prior mean. For the simulations and empirical application we use a prior centered around zero.

Sampling from $p(\boldsymbol{\xi}|\mathbf{a})$. Recall that $\mathbf{a} = \text{vec}(\mathbf{A}_1, \dots, \mathbf{A}_p)$. The prior variance of each generic element a_j ($j = 1, \dots, M^2P$), is given by the product of a local, variable-specific component c_j^2 and a global, shared component d^2 . [Makalic and Schmidt \(2016\)](#) develop

a simple sampling scheme for the case of the horseshoe prior involving two auxiliary variables ν_j and ζ . The conditional posterior of c_j^2 and d^2 takes the form:

$$c_j^2 | \bullet \sim \text{InvGa}\left(1, \frac{1}{\nu_j} + \frac{a_j^2}{2d^2}\right),$$

$$d^2 | \bullet \sim \text{InvGa}\left(\frac{M^2P + 1}{2}, \frac{1}{\zeta} + \frac{1}{2} \sum_{j=1}^{M^2P} \frac{a_j^2}{c_j^2}\right).$$

The auxiliary variables are in turn sampled from:

$$\nu_j | \bullet \sim \text{InvGa}\left(1, 1 + \frac{1}{c_j^2}\right),$$

$$\zeta | \bullet \sim \text{InvGa}\left(1, 1 + \frac{1}{d^2}\right).$$

Sampling from $p(\{\tau_{i,j}^2\}_{i,j} | \{\mathbf{d}_t\}_{t=1}^T, \mathbf{\Omega}, \mathbf{\Delta})$. Following [Ishwaran and Rao \(2005\)](#) we update the prior variances of the off-diagonal elements of the symmetric precision matrix $\tau_{ij}^2 = \tau_{ji}^2$ ($i = 1, \dots, M; j = 1, \dots, i - 1$) by drawing from an inverse gamma distribution. The conditional posterior is given by:

$$\tau_{i,j}^2 \sim \mathcal{G}^{-1}\left(a_\tau + \frac{1}{2}, b_\tau + \frac{\omega_{i,j}^2}{\delta_{i,j} + (1 - \delta_{i,j})c}\right).$$

Sampling from $p(\mathbf{\Pi}, \mathbf{z} | \mathbf{\Delta})$. Conditional on the adjacency matrix $\mathbf{\Delta}$ we employ the sampling algorithm by [Legramanti et al. \(2022\)](#) to obtain a draw for the estimated group membership. For each node $j = 1, \dots, M$ it performs the following steps:

1. Remove j from the network.
2. Check whether group z_j which previously included j is now empty. If necessary, eliminate it and relabel all cluster indicators such that each cluster $h = 1, \dots, H$ contains at least one node.

3. Sample z_j from the conditional posterior:

$$p(z_j = h | \mathbf{z}_{-j}, \mathbf{\Delta}) \propto \begin{cases} \mathcal{W}_{M, H^-} (n_h^- - \sigma) \prod_{k=1}^{H^-} \frac{\mathcal{B}(a_\pi + m_{hk}^- + r_{jk}, b_\pi + \bar{m}_{hk}^- + \bar{r}_{jk})}{\mathcal{B}(a_\pi + m_{hk}^-, b_\pi + \bar{m}_{hk}^-)} & \text{for } h \leq H^-, \\ \mathcal{W}_{M, H^- + 1} \prod_{k=1}^{H^-} \frac{\mathcal{B}(a_\pi + r_{jk}, b_\pi + \bar{r}_{jk})}{\mathcal{B}(a_\pi, b_\pi)} & \text{for } h = H^- + 1, \end{cases}$$

where \mathbf{z}_{-j} are the cluster membership indicators, n_h is the size of cluster h , m_{hk}^- (\bar{m}_{hk}^-) are the number of links (non-existing links) between clusters h and k and H^- is the number of non-empty clusters, all after removing j . We denote the number of links (non-existing links) between node j and the k -th cluster as r_{jk} (\bar{r}_{jk}). \mathcal{W}_{M, H^-} and $\mathcal{W}_{M, H^- + 1}$ are weights and σ is a discounting factor that controls how quickly the sampler opens up a new cluster. All of the latter are determined by the specification of the Gibbs-type prior.

B Data Appendix

An overview of the variables used in both models and their transformations can be found in table 3.

FRED Code	Variable Name	tcode	Group	Medium	Large
GDPC1	Real Gross Domestic Product	5	NIPA	x	x
PCECC96	Real Personal Consumption Expenditures	5	NIPA	x	x
GPDIC1	Real Private Domestic Investment	5	NIPA	x	x
PRFIx	Real Private Fixed Investment	5	NIPA	x	x
INDPRO	Industrial Production Index	5	IP		x
CUMFNS	Capacity Utilization: Manufacturing	2	IP		x
UNRATE	Civilian Unemployment Rate	2	Labor Market	x	x
SRVPRD	All Employees: Services	5	Labor Market		x
CE16OV	Civilian Employment Level	5	Labor Market		x
AWHMAN	Avg Weekly Hours of Production & Nonsupervisory Employees: Manufacturing	1	Labor Market		x
CES3000000008x	Avg Hourly Earnings : Manufacturing	5	Earnings	x	x
CPIAUCSL	Consumer Price Index for All Urban Consumers: All Items	6	Prices	x	x
PCECTPI	Personal Consumption Expenditures: Chain-type Price Index	6	Prices	x	x
GDPCTPI	Gross Domestic Product: Chain-type Price Index	6	Prices		x
GPDICTPI	Gross Private Domestic Investment: Chain-type Price Index	6	Prices		x
FEDFUNDS	Effective Federal Funds Rate	2	Interest Rates		x
GS1	1-Year Treasury Rate	2	Interest rates	x	
BAA10YM	Moody's Seasoned Baa Corporate Bond Yield Rel. to 10-Year Treasury Constant Maturity	2	Interest Rates		x
M2REAL	Real M2 Money Stock	5	Money & Credit		x
S&P 500	S&P 500 Index	5	Stock Market	x	x

Table 3: Macroeconomic variables alongside their assigned groups (as per [McCracken and Ng \(2016\)](#)) and their FRED identifier and tcode. tcode refers to the following transformations: 1) None 2) Δx_t 5) $\Delta \log(x_t)$ 6) $\Delta^2 \log(x_t)$