VU Research Portal

Bayesian nonparametric sparse VAR models

Billio, Monica; Casarin, Roberto; Rossini, Luca

published in

Journal of Econometrics
2019

DOI (link to publisher)

[10.1016/j.jeconom.2019.04.022](https://doi.org/10.1016/j.jeconom.2019.04.022)

document version

Publisher's PDF, also known as Version of record

document license

Article 25fa Dutch Copyright Act

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Billio, M., Casarin, R., & Rossini, L. (2019). Bayesian nonparametric sparse VAR models. *Journal of Econometrics*, 212(1), 97-115. <https://doi.org/10.1016/j.jeconom.2019.04.022>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

Bayesian nonparametric sparse VAR models[☆]

Monica Billio^{a,*}, Roberto Casarin^a, Luca Rossini^{a,b}

^a Ca' Foscari University of Venice, Italy

^b Free University of Bozen-Bolzano, Italy

ARTICLE INFO

Article history:

Available online 20 April 2019

Keywords:

Bayesian nonparametrics
Bayesian model selection
Connectedness
Large vector autoregression
Multilayer networks
Network communities
Shrinkage

ABSTRACT

High dimensional vector autoregressive (VAR) models require a large number of parameters to be estimated and may suffer of inferential problems. We propose a new Bayesian nonparametric (BNP) Lasso prior (BNP-Lasso) for high-dimensional VAR models that can improve estimation efficiency and prediction accuracy. Our hierarchical prior overcomes overparametrization and overfitting issues by clustering the VAR coefficients into groups and by shrinking the coefficients of each group toward a common location. Clustering and shrinking effects induced by the BNP-Lasso prior are well suited for the extraction of causal networks from time series, since they account for some stylized facts in real-world networks, which are sparsity, communities structures and heterogeneity in the edges intensity. In order to fully capture the richness of the data and to achieve a better understanding of financial and macroeconomic risk, it is therefore crucial that the model used to extract network accounts for these stylized facts.

© 2019 Elsevier B.V. All rights reserved.

1. Introduction

In the last decade, high dimensional models and large datasets have increased their importance in economics (e.g., see Scott and Varian, 2014) and finance. In macroeconomics, some authors investigate the use of large datasets (between 10 and 170 series) to improve forecasts (e.g., see Banbura et al., 2010; Stock and Watson, 2012; Koop, 2013; Carriero et al., 2015; McCracken and Ng, 2016; Kaufmann and Schumacher, 2017), while in finance, large datasets (between 10 and 200 series) have been used to analyse financial crises, contagion effects and their impact on the real economy (e.g., see Brownlees and Engle, 2016; Barigozzi and Brownlees, 2018).¹ Moreover, the level of interaction, or connectedness, between financial institutions represents a powerful tool in monitoring financial stability (e.g., see Diebold and Yilmaz, 2015; Scott, 2016), whereas measuring interdependence in business cycles and in financial markets (Diebold and Yilmaz, 2014; Demirel et al., 2018) is essential in pursuing economic stability. Making inference on dependence structures in large

[☆] We are grateful to the Guest Editor and the reviewers for their useful comments which significantly improved the quality of the paper. We would like to thank all the conference participants for helpful discussions at: “9th RCEA” and “12th RCEA” in Rimini; “Internal Seminar” at Ca' Foscari University; “Statistics Seminar” at University of Kent; “9th CFE” in London; “ISBA 2016” in Sardinia; “3rd BAYSM” at University of Florence; “7th ESOBE” in Venice; “10th CFE” in Seville; “7th ICEEE” in Messina; “Bomopav 2017” in Venice; “Big Data in Predictive Dynamic Econometric Modeling” at University of Pennsylvania and “11th BNP” in Paris. We benefited greatly from suggestions and discussions with Conception Ausin, Francis Diebold, Lorenzo Frattarolo, Pedro Galeano, Jim Griffin, Mark Jensen, Gary Koop, Dimitris Korobilis, Stefano Tonellato and Mike West. This research used the SCSCF multiprocessor cluster system at University Ca' Foscari and is part of the research activities at the Venice Center in Economic and Risk Analytics for public policies (VERA) at Ca' Foscari University of Venice.

* Corresponding author.

E-mail addresses: billio@unive.it (M. Billio), r.casarin@unive.it (R. Casarin), luca.rossini@unibz.it (L. Rossini).

¹ See Table S.1 and S.2 in Supplementary Material for a review.

sets of time series and representing them as graphs (or networks) becomes a relevant and challenging issue in financial and macroeconomic risk measurement (e.g., see [Billio et al., 2012](#); [Diebold and Yilmaz, 2015, 2016](#); [Bianchi et al., 2018](#)).

For forecasting economic variables and analysing their dynamic properties and causal structure, vector autoregressive (VAR) models have been widely used. In order to avoid overparametrization and overfitting issues, Bayesian inference and suitable classes of prior distributions have been successfully used for VARs (e.g., see [Litterman, 1986](#); [Doan et al., 1984](#); [Sims and Zha, 1998](#)) and for panel VARs ([Canova and Ciccarelli, 2004](#)).² Nevertheless, these prior distributions may be not effective when dealing with very large VAR models. Thus, new prior distributions have been proposed to deal with zero restrictions in VAR parameters. [George et al. \(2008\)](#) introduce Stochastic Search Variable Selection (SSVS) based on spike-and-slab prior distribution. SSVS has been extended in many directions (e.g., see [Korobilis, 2013](#); [Koop and Korobilis, 2016](#); [Korobilis, 2016](#)). [Wang \(2010\)](#), [Ahelgebey et al. \(2016a,b\)](#) propose graphical VAR relying on graphical prior distributions. [Gefang \(2014\)](#) proposes Bayesian Lasso based on doubly adaptive elastic-net prior. We thus propose a new Bayesian nonparametric (BNP) prior distributions for VARs, which can improve both estimation efficiency and prediction accuracy, thanks to the combination of two types of restrictions, called clustering and shrinking effects.

Our BNP Lasso prior distribution (BNP-Lasso) groups the VAR coefficients into clusters and shrinks the coefficients within a cluster toward a common location. The shrinking effects are due to the Lasso-type distribution at the first stage of the hierarchy. The Bayesian Lasso prior allows for reformulating the Bayesian VAR model as a penalized regression (see [Park and Casella, 2008](#)) and is now a standard choice for inducing sparsity. This assumption and our inference procedure can be easily modified to combine clustering effects with alternative shrinkage procedures such as adaptive Lasso ([Huber and Feldkircher, 2017](#)) and adaptive elastic-net ([Gefang, 2014](#)).

The clustering effects of our approach are due to a Dirichlet process (DP) hyperprior at the second stage, which induces a random partition of the parameter space. The attractive property of DP is that the number of clusters can be inferred without being specified in advance as it happens in finite mixture modelling. A similar approach for finite mixture models would be model averaging or model selection for the number of components, but it can be nontrivial to design efficient MCMC samplers (e.g. [Richardson and Green \(1997\)](#)). Meanwhile, for DP there are relatively simple and flexible samplers for posterior approximation (e.g. [Kalli et al. \(2011\)](#)). The DP prior is standard in BNP, but the assumption could be replaced by more flexible clustering-inducing priors, such as the Pitman–Yor process prior ([Pitman and Yor, 1997](#)). In the presentation of the model we introduce two further assumptions, namely a prior partition of the parameter space and the use of conditionally independent and DP priors on each subspace. We envisage two advantages in using this specification strategy: from a modelling perspective, it adds flexibility to the model since it allows for different degrees of prior information in the blocks of parameters; from a computational perspective, it allows for breaking the posterior simulation procedure into several steps overcoming the difficulties to deal with high dimensional parameter vectors. These assumptions are not essential for the clustering and shrinking effects to appear, since the effects are given by the two-stage structure of the hyper-prior and do not depend on the blocking strategy. Moreover, the independent DP assumptions can be easily removed to allow for dependent clustering features as in beta-dependent Pitman–Yor prior ([Bassetti et al., 2014](#); [Griffin and Steel, 2011](#)). We leave these issues for future research.

Up to our knowledge, our paper is the first to provide sparse Bayesian nonparametric VAR models, and the proposed prior distribution is general as it easily extends to other model classes, such as SUR models. In this sense, we contribute to the literature on Bayesian nonparametrics ([Ferguson \(1973\)](#) and [Lo \(1984\)](#)) and its applications to time series (e.g., see [Hirano, 2002](#); [Taddy and Kottas, 2009](#); [Jensen and Maheu, 2010](#); [Griffin and Steel, 2011](#); [Di Lucca et al., 2013](#); [Bassetti et al., 2014](#); [Nieto-Barajas and Quintana, 2016](#); [Griffin and Kalli, 2018](#)). See also [Hjort et al. \(2010\)](#) for a review. We substantially improve [Hirano \(2002\)](#), [Bassetti et al. \(2014\)](#) and [Griffin and Kalli \(2018\)](#) by allowing for sparsity in their nonparametric dynamic models and we extend [MacLehose and Dunson \(2010\)](#) by proposing vectors of dependent sparse Dirichlet process priors. As regards to the posterior approximation, we develop a MCMC algorithm building on the slice sampler introduced by [Walker \(2007\)](#), [Kalli et al. \(2011\)](#) and [Hatjispyros et al. \(2011\)](#) for vectors of dependent random measures.

We also contribute to the literature on the analysis of economic and financial fluctuations through the lens of networks (e.g., see [Acemoglu et al., 2012, 2015](#)) and the measures of risk connectedness (e.g., see [Billio et al., 2012](#); [Diebold and Yilmaz, 2014](#); [Demirer et al., 2018](#)). Our BNP-Lasso VAR model is particularly well suited for extracting causal networks from time series, since it allows for some features detected in many real-world networks, that are sparsity, communities or blocks, and edge heterogeneity. For the first time, we extract causal networks with clustering effects in the edge intensity. We use the posterior random partition induced by our BNP-Lasso prior to cluster the edge intensities into groups, which allows for multilayer network representation ([Kivela et al., 2014](#)) and thus for a better analysis of the network topology and consequently of financial and macroeconomic connectedness. In the last years, there is an emerging literature based on BNP for network analysis (see, e.g. [Caron and Fox, 2017](#), and reference therein). Moreover, as evidenced in many papers from the network literature, edge clustering is even more important to achieve a better description of the connectivity. Our BNP-Lasso VAR allows for extracting edge clustering in a well founded inference context.

We find empirical evidence of strong heterogeneity across network edges and centrality heterogeneity of nodes depending on their edge intensity levels. Finally, we also perform a comparison of our Bayesian nonparametric model with alternative shrinkage approaches and give evidence of the best forecasting performance of our proposed prior.

² See [Karlsson \(2013\)](#) for a comprehensive review.

The paper is organized as follows. Section 2 introduces our sparse Bayesian VAR model. Section 3 presents posterior approximation and network extraction methods. In Section 4, we present the applications to business cycle and realized volatility datasets. Section 5 concludes.

2. A sparse Bayesian VAR model

2.1. VAR models

Let N be the number of units (e.g. countries, regions or micro studies) in a panel dataset and $\mathbf{y}_{i,t} = (y_{i,1t}, \dots, y_{i,mt})'$ a vector of m variables available for the i th unit, with $i = 1, \dots, N$. A panel VAR is defined as the system of regression equations:

$$\mathbf{y}_{i,t} = \mathbf{b}_i + \sum_{j=1}^N \sum_{l=1}^p B_{ijl} \mathbf{y}_{j,t-l} + \boldsymbol{\varepsilon}_{i,t}, \quad t = 1, \dots, T \text{ and } i = 1, \dots, N, \quad (1)$$

where $\mathbf{b}_i = (b_{i,1}, \dots, b_{i,m})'$ the vector of constant terms and B_{ijl} the $(m \times m)$ matrix of unit- and lag-specific coefficients. We assume that $\boldsymbol{\varepsilon}_{i,t} = (\varepsilon_{i,1t}, \dots, \varepsilon_{i,mt})'$, are i.i.d. for $t = 1, \dots, T$, with Gaussian distribution $\mathcal{N}_m(\mathbf{0}, \Sigma_i)$ and that $\text{Cov}(\boldsymbol{\varepsilon}_{i,t}, \boldsymbol{\varepsilon}_{j,t}) = \Sigma_{ij}$. Eq. (1) can be written in the more compact form as

$$\mathbf{y}_{i,t} = B_i \mathbf{x}_t + \boldsymbol{\varepsilon}_{i,t}, \quad i = 1, \dots, N, \quad (2)$$

where $B_i = (\mathbf{b}_i, B_{i,11}, \dots, B_{i,1p}, \dots, B_{i,N1}, \dots, B_{i,Np})$ is a $(m \times (1 + Nmp))$ matrix of coefficients and $\mathbf{x}_t = (1, \mathbf{x}_{1,t}, \dots, \mathbf{x}_{N,t})'$ is the vector of lagged variables with $\mathbf{x}_{i,t} = (y_{i,t-1}, \dots, y_{i,t-p})$ for $i = 1, \dots, N$.

The system of equations in (2) can be written in the SUR regression form:

$$\mathbf{y}_t = (I_{Nm} \otimes \mathbf{x}_t') \boldsymbol{\beta} + \boldsymbol{\varepsilon}_t, \quad (3)$$

where $\boldsymbol{\beta} = \text{vec}(B)$, $B = (B'_1, \dots, B'_N)$, $\boldsymbol{\varepsilon}'_t = (\boldsymbol{\varepsilon}'_{1t}, \dots, \boldsymbol{\varepsilon}'_{Nt})$, $\otimes$ is the Kronecker product and $\text{vec}(\cdot)$ the column-wise vectorization operator that stacks the columns of a matrix into a column vector (Magnus and Neudecker, 1999, pp. 31–32).

2.2. A BNP-Lasso prior assumption

The number of coefficients in (3) is $n = Nm(1 + Nmp)$, which can be large for high dimensional panel of time series. In order to avoid overparameterization, unstable predictions and overfitting problem, we follow a hierarchical specification strategy for the prior distributions (e.g., Canova and Ciccarelli (2004), Kaufmann (2010), Bassetti et al. (2014)). Some classes of hierarchical prior distributions are used to incorporate interdependences across groups of parameters with various degrees of information pooling (see Karlsson, 2013, for a review). Other classes of priors (e.g., MacLehose and Dunson (2010), Wang (2010), Bianchi et al. (2018)) are used to induce sparsity. Our new class of hierarchical priors for the VAR coefficients $\boldsymbol{\beta}$ combines sparsity and information pooling within groups of coefficients.

We assume that $\boldsymbol{\beta}$ can be exogenously partitioned in M blocks, $\boldsymbol{\beta}_i = (\beta_{i1}, \dots, \beta_{in_i})$, $i = 1, \dots, M$. In our empirical applications, blocks correspond to the VAR coefficients at different lags. This assumption is not essential to appear in our prior for the sparsity and information pooling. Nevertheless, we envisage two advantages in using this specification strategy. From a modelling perspective, it adds flexibility to the model since it allows for different degrees of prior information in the blocks of parameters. For example, lag-specific shrinking effects as in a Minnesota type prior, or country-specific shrinking effects in panel VAR model can be easily introduced in our framework. From a computational perspective, it allows for breaking the posterior simulation procedure into several steps overcoming the difficulties of dealing with high dimensional parameter vectors.

For each block, we introduce shrinking effects by using a Bayesian-Lasso prior $f_i(\boldsymbol{\beta}_i) = \prod_{j=1}^{n_i} \mathcal{NG}(\beta_{ij} | \mu, \gamma, \tau)$, where:

$$\mathcal{NG}(\beta | \mu, \gamma, \tau) = \int_0^{+\infty} \mathcal{N}(\beta | \mu, \lambda) \mathcal{Ga}(\lambda | \gamma, \tau/2) d\lambda, \quad (4)$$

is the normal–gamma distribution with μ , γ and τ , the location, shape and scale parameter, respectively (see Appendix A). The normal–gamma distribution induces shrinkage toward the prior mean of μ . In Bayesian Lasso (Park and Casella, 2008), μ is set equal to zero. We extend the Lasso model by allowing for multiple location, shape and scale parameter such that: $f_i(\boldsymbol{\beta}_i) = \prod_{j=1}^{n_i} \mathcal{NG}(\beta_{ij} | \mu_{ij}^*, \gamma_{ij}^*, \tau_{ij}^*)$, and reduce curse of dimensionality and overfitting by employing a Dirichlet process (DP) as a prior for the normal–gamma parameters. More specifically, let $\boldsymbol{\theta}^* = (\mu^*, \gamma^*, \tau^*)$ be the normal–gamma parameter vector, our BNP-Lasso prior for $\boldsymbol{\theta}^*$ and $\boldsymbol{\beta}_i$ is

$$\beta_{ij} \stackrel{\text{ind}}{\sim} \mathcal{NG}(\beta_{ij} | \boldsymbol{\theta}_{ij}^*), \quad \text{and} \quad \boldsymbol{\theta}_{ij}^* | \mathbb{Q}_i \stackrel{\text{i.i.d.}}{\sim} \mathbb{Q}_i, \quad (5)$$

with $j = 1, \dots, n_i$, where $\mathbb{Q}_i$ is a random measure.

Following the strategy in Müller et al. (2004), Pennell and Dunson (2006), Hatjispyros et al. (2011), we assume $\mathbb{Q}_i$ is a convex combination of a common $\mathbb{P}_0$ and a block-specific $\mathbb{P}_i$ random measure, that is

$$\begin{aligned}\mathbb{Q}_1(d\theta_1) &= \pi_1 \mathbb{P}_0(d\theta_1) + (1 - \pi_1) \mathbb{P}_1(d\theta_1), \\ &\vdots \\ \mathbb{Q}_M(d\theta_M) &= \pi_M \mathbb{P}_0(d\theta_M) + (1 - \pi_M) \mathbb{P}_M(d\theta_M).\end{aligned}\quad (6)$$

The random measure $\mathbb{P}_0$ favours sparsity by shrinking coefficients toward zero, as in standard Bayesian Lasso, i.e.

$$\mathbb{P}_0(d\theta) \sim \delta_{\{(0, \gamma_0, \tau_0)\}}(d(\mu, \gamma, \tau)), \quad \text{with } (\gamma_0, \tau_0) \sim \mathcal{GS}(\gamma_0, \tau_0 | v_0, p_0, s_0, n_0), \quad (7)$$

where $\delta_{\{\theta^*\}}(\theta)$ denotes the Dirac measure indicating that the random vector θ has a degenerate distribution with mass at the location θ^* , and $\mathcal{GS}(\gamma, \tau | v, p, s, n)$ is a Gamma scale-shape distribution with parameters v, p, s and n (see Appendix A).

The random measure, $\mathbb{P}_i, i > 0$, is a DP prior (DPP), which induces a random partition of the parameter space and a parameter clustering effect (see, Hirano (2002), Griffin and Steel (2011), Bassetti et al. (2014)), and shrinks coefficients toward multiple non-zero locations, i.e.

$$\mathbb{P}_i(d\theta) \stackrel{i.i.d.}{\sim} \text{DPP}(\tilde{\alpha}, H), \quad \text{with } H \sim \mathcal{N}(\mu | c, d) \cdot \mathcal{GS}(\gamma, \tau | v_1, p_1, s_1, n_1), \quad (8)$$

where $\tilde{\alpha}$ and H are the DPP concentration parameter and base measure, respectively. See Appendix A for a definition of DPP. For the mixing parameter π_i , we assume a Beta distribution, $\pi_i \stackrel{i.i.d.}{\sim} \text{Be}(\pi_i | 1, \alpha_i)$.

The amount of shrinkage in $\mathbb{P}_0$ and $\mathbb{P}_i$ is determined by the hyperparameters of $\mathcal{GS}(v, p, s, n)$. In our empirical application, we assume the hyperparameter values $v_0 = 30, s_0 = 1/30, p_0 = 0.5$ and $n_0 = 18$ for the sparse component and $v_1 = 3, s_1 = 1/3, p_1 = 0.5$ and $n_1 = 10$ for the DPP component and $\alpha_i = 1$ for the mixing parameter, as in MacLehose and Dunson (2010).

For the variance-covariance matrix Σ , we assume a graphical prior distribution as in Carvalho et al. (2007) and Wang (2010), where the zero restrictions on the covariances are induced by a graph G , that is by an ordered pair of sets (N_G, E_G) , where N_G is a vertex set and E_G an edge set. Conditionally to a specified graph G , we assume a Hyper Inverse Wishart prior distribution for Σ , that is:

$$\Sigma \sim \text{HIW}_G(b, L), \quad (9)$$

where b and L are the degrees of freedom and scale hyperparameters, respectively. See Appendix A for further details on Hyper Inverse Wishart distributions.

The prior over the graph structure is defined as a product of Bernoulli distributions with parameter ψ , which is the probability of having an edge. That is, a m -node graph $G = (N_G, E_G)$, has a prior probability:

$$p(G) \propto \prod_{i,j} \psi^{e_{ij}} (1 - \psi)^{(1-e_{ij})} = \psi^{|E_G|} (1 - \psi)^{\kappa - |E_G|}, \quad (10)$$

with $e_{ij} = 1$ if $(i, j) \in E_G$, where $|N_G|$ and $|E_G|$ are the cardinalities of the vertex and edge sets, respectively, $\kappa = \binom{|N_G|}{2}$ the maximum number of edges. To induce sparsity we choose $\psi = 2/(p - 1)$ which would provide a prior mode at p edges.

In summary, our hierarchical prior in Eq. (5)–(10) is represented through the *Directed Acyclic Graph* (DAG) in Fig. 1. Shadow and empty circles indicate the observable and non-observable random variables, respectively. The directed arrows show the causal dependence structure of the model. The left panel shows the priors for β_i and Σ , which are the first stage of the hierarchy. The second stage (right panel) involves the sparse dependent DP prior for the shrinking parameters μ, γ and τ .

Our prior has the infinite mixture representation with a countably infinite number of clusters, which is one of the appealing feature of BNP:

$$f_i(\beta_i | \mathbb{Q}_i) = \sum_{k=0}^{\infty} \check{w}_{ik} \mathcal{G}(\beta_i | \check{\theta}_{ik}), \quad (11)$$

where

$$\check{w}_{ik} = \begin{cases} \pi_i, & k = 0, \\ (1 - \pi_i) w_{ik}, & k > 0, \end{cases} \quad \check{\theta}_{ik} = \begin{cases} (0, \gamma_0, \tau_0), & k = 0, \\ (\mu_{ik}, \gamma_{ik}, \tau_{ik}), & k > 0. \end{cases}$$

See Appendix B for a proof. This representation shows that our prior not only shrinks coefficients to zero as in Bayesian Lasso (Park and Casella, 2008) or Elastic-net (Zou and Hastie, 2005), but also induces a probabilistic clustering effect in the parameter set as in other Bayesian nonparametric models (Hirano, 2002) and Bassetti et al. (2014).

Our prior places no bounds on the number of mixture components and can fit arbitrary complex distribution. Nevertheless, since the mixture weights, w_{ik} 's, decrease exponentially quickly only a small number of clusters will be used to model the data a priori. The prior number of clusters is a random variable with distribution driven by the concentration parameter $\tilde{\alpha}$. In our applications, we set $\tilde{\alpha} = 1$.

Fig. 1. DAG of the Bayesian nonparametric model for inference on VARs. It exhibits the hierarchical structure of priors and hyperparameters. The directed arrows show the causal dependence structure of the model. Left panel is related to the first stage of the hierarchy and right panel to the second stage.

3. Posterior inference

3.1. Sampling method

Since the posterior distribution is not tractable, Bayesian estimator cannot be obtained analytically. In this paper, we rely on simulation based inference methods, and develop a Gibbs sampler algorithm for approximating the posterior distribution. For the sake of simplicity and without loss of generality, we shall describe the sampling strategy for the two-block case, i.e. $M = 2$.

In order to develop a more efficient MCMC procedure, we follow a data augmentation approach. For each block $i = 1, 2$, we introduce two sets of allocation variables, $\xi_{ij}, d_{ij}, j = 1, \dots, n_i$, a set of stick-breaking variables, $v_{ij}, j = 1, 2, \dots$ and a set of slice variables, $u_{ij}, j = 1, \dots, n_i$. The allocation variable, ξ_{ij} , selects the sparse component $\mathbb{P}_0$, when ξ_{ij} is equal to zero and the non-sparse component $\mathbb{P}_i$, when it is equal to one. The second allocation variable, d_{ij} , selects the component of the Dirichlet mixture $\mathbb{P}_i$ to which each single coefficient β_{ij} is allocated to. The sequence of stick-breaking variables defines the mixture weights, whereas the slice variable, u_{ij} , allows us to deal with the infinite number of mixture components by identifying a finite number of stick-breaking variables to be sampled and an upper bound for the allocation variables d_{ij} .

We demarginalize the Normal–Gamma distribution by introducing a latent variable λ_{ij} for each β_{ij} and obtain the joint posterior distribution

$$f(\Theta, \Sigma, \Lambda, U, D, V, \Xi | Y) \propto \prod_{t=1}^T (2\pi | \Sigma |)^{-1/2} \exp \left(-\frac{1}{2} (y_t - X_t' \beta)' \Sigma^{-1} (y_t - X_t' \beta) \right) \cdot \prod_{j=1}^{n_1} f_1(\beta_{1j}, \lambda_{1j}, u_{1j}, d_{1j}, \xi_{1j}) \prod_{j=1}^{n_2} f_2(\beta_{2j}, \lambda_{2j}, u_{2j}, d_{2j}, \xi_{2j}) \cdot \prod_{k>1} \mathcal{Be}(v_{1k} | 1, \alpha) \mathcal{Be}(v_{2k} | 1, \alpha) \mathcal{HTWG}(b, L) \mathcal{GS}(\gamma_0, \tau_0 | \nu_0, p_0, s_0, n_0) \cdot \prod_{k>1} \mathcal{N}(\mu_{1k} | c, d) \mathcal{GS}(\gamma_{1k}, \tau_{1k} | \nu_1, p_1, s_1, n_1) \mathcal{N}(\mu_{2k} | c, d) \mathcal{GS}(\gamma_{2k}, \tau_{2k} | \nu_1, p_1, s_1, n_1), \quad (12)$$

where $U = \{u_{ij} : j = 1, 2, \dots, n_i \text{ and } i = 1, 2\}$ and $V = \{v_{ij} : j = 1, 2, \dots \text{ and } i = 1, 2\}$ are the collections of slice variables and stick-breaking components, respectively; $D = \{d_{ij} : j = 1, 2, \dots, n_i \text{ and } i = 1, 2\}$ and $\Xi = \{\xi_{ij} : j = 1, 2, \dots, n_i \text{ and } i = 1, 2\}$ are the allocation variables; $\Theta = \{(\mu_0, \gamma_0, \tau_0), (\mu_{ik}, \gamma_{ik}, \tau_{ik}) : i = 1, 2 \text{ and } k = 1, 2, \dots\}$ are the atoms; $\pi = (\pi_1, \pi_2)$ are the block-specific probabilities of shrinking coefficients to zero and

$$f_i(\beta_{ij}, \lambda_{ij}, u_{ij}, d_{ij}, \xi_{ij}) = (\mathbb{I}(u_{ij} < \tilde{w}_{d_{ij}}) \mathcal{N}(\beta_{ij} | 0, \lambda_{ij}) \mathcal{Ga}(\lambda_{ij} | \gamma_0, \tau_0/2))^{1-\xi_{ij}} \cdot (\mathbb{I}(u_{ij} < \tilde{w}_{id_{ij}}) \mathcal{N}(\beta_{ij} | \mu_{id_{ij}}, \lambda_{ij}) \mathcal{Ga}(\lambda_{ij} | \gamma_{id_{ij}}, \tau_{id_{ij}}/2))^{\xi_{ij}} \pi_i^{1-\xi_{ij}} (1 - \pi_i)^{\xi_{ij}}. \quad (13)$$

See [Appendix B](#) for a derivation.

We obtain random samples from the posterior distributions by Gibbs sampling. The Gibbs sampler iterates over the following steps using the conditional independence between variables as described in [Appendix C](#):

1. The slice and stick-breaking variables U and V are updated given $[\Theta, \beta, \Sigma, G, \Lambda, D, \Xi, \pi, Y]$;
2. The latent scale variables Λ are updated given $[\Theta, \beta, \Sigma, G, U, V, D, \Xi, \pi, Y]$;
3. The parameters of the Normal–Gamma distribution Θ are updated given $[\beta, \Sigma, G, \Lambda, U, V, D, \Xi, \pi, Y]$;
4. The coefficients β of the VAR model are updated given $[\Theta, \Sigma, G, \Lambda, U, V, D, \Xi, \pi, Y]$;
5. The matrix of variance–covariance Σ is updated given $[\Theta, \beta, G, \Lambda, U, V, D, \Xi, \pi, Y]$;
6. The graph G is updated given $[\Theta, \beta, \Sigma, \Lambda, U, V, D, \Xi, \pi, Y]$;
7. The allocation variables D and Ξ are updated given $[\Theta, \beta, \Sigma, G, \Lambda, U, V, \pi, Y]$;
8. The probability, π , of shrinking-to-zero the coefficient is updated given $[\Theta, \beta, \Sigma, G, \Lambda, U, V, D, \Xi, Y]$.

3.2. Network extraction

Pairwise Granger causality has been used to extract linkages and networks describing relationships between variables of interest, e.g. financial and macroeconomic linkages (Billio et al., 2012; Barigozzi and Brownlees, 2018). The pairwise approach does not consider conditioning on relevant covariates thus generating spurious causality effects. The conditional Granger approach includes relevant covariates, however the large number of variables relative to the number of data points can lead to overparametrization and consequently to inefficiency in gauging the causal relationships (e.g., see Ahelgebey et al., 2016a,b). Our BNP-Lasso can be used to extract the network while reducing overfitting and curse of dimensionality problems. Also, it allows to define edge-coloured graphs, which account for various stylized facts recently investigated in financial networks, such as the presence of communities, hubs and linkage heterogeneity.

For sake of simplicity, we consider one lag, one block of coefficients and one unit, i.e. $M = 1$, $N = 1$ and $p = 1$. We denote with $G = (V, E)$ the directed graph with vertex set $V = \{1, \dots, m\}$ and edge set $E \subset V \times V$. A directed edge from i to j is the ordered set $\{i, j\} \in E$ with $i, j \in V$. The adjacency matrix A associated with G has (j, i) th element

$$a_{j,i} = \begin{cases} 1 & \text{if } \{i, j\} \in E, \\ 0 & \text{otherwise.} \end{cases}$$

See Newman (2010) for an introduction to graph theory and network analysis and Jackson (2008) for an economic perspective on networks.

A weighted graph is defined by the ordered triplet $G = (V, E, C)$, where C is the weight matrix with (j, i) th element $c_{i,j}$ representing the weight associated to the edge $\{i, j\}$. When $c_{i,j}$ belongs to a countable set, then $G = (V, E, C)$ is called edge coloured graph.

In time series graphs, causal networks encode the conditional independence structure of a multivariate process (Eichler, 2013). An edge $\{i, j\} \in E$ exists if and only if the VAR coefficient B_{ji} of the variable $y_{i,t-1}$ in the equation of $y_{j,t}$ is not null. An advantage in using BNP-Lasso VAR is that the inference procedure provides edge probabilities, edge weights and edge partition as a natural output and graph uncertainty can be easily included in network analysis.

More specifically, we assume an edge from i to j exists, i.e. $\{i, j\} \in E$, if $\hat{\xi}_{1,g(i,j)} = 0$, where $\hat{\xi}_{1,g(i,j)}$ is the Maximum a Posterior (MAP) estimator of $\xi_{1,g(i,j)}$. The posterior clustering of the VAR coefficients allows us to define an edge coloured graph $G = (V, E, C)$. Following a standard procedure in BNP (e.g., see Bassetti et al., 2014), we use the MCMC draws of the allocation variables d_{1i} and d_{1j} to approximate the posterior probabilities of joint classification of an edge $\phi_{ij} = P(d_{1i} = d_{1j} | Y, \xi_{1i} = 0, \xi_{1j} = 0)$,

$$\hat{\phi}_{ij} = \frac{1}{H} \sum_{h=1}^H \delta_{d_{1i}^h}(d_{1j}^h) \quad (14)$$

and to find the posterior partition $\hat{D}$, which minimizes the sum of squared deviations from the joint classification probabilities, i.e.

$$\hat{D} = \arg \min_{D \in \{D^1, \dots, D^H\}} \sum_{i=1}^{\tilde{m}} \sum_{j=1}^{\tilde{m}} \left(\delta_{d_{1i}^h}(d_{1j}^h) - \hat{\phi}_{ij} \right)^2. \quad (15)$$

The number of distinct elements $\hat{K}$ in the partition $\hat{D}$ provides an estimate of the number of edge intensity levels (colours). The weight c_{ji} indicating the intensity of the edge $\{i, j\}$ can be computed by using the location atom posterior mean, $\hat{\mu}_{kl}^*$,

$$c_{ji} = \begin{cases} \hat{\mu}_{11}^* & \text{if } \hat{d}_{1,g(i,j)} = 1 \text{ and } \hat{\xi}_{1,g(i,j)} = 0, \\ \vdots & \vdots \\ \hat{\mu}_{1\hat{K}}^* & \text{if } \hat{d}_{1,g(i,j)} = \hat{K} \text{ and } \hat{\xi}_{1,g(i,j)} = 0, \\ 0 & \text{if } \hat{\xi}_{1,g(i,j)} = 1. \end{cases} \quad (16)$$

Panel (a) in Fig. 2 provides an example of weighted network representation for a VAR(1) with $m = 4$ variables and $\hat{K} = 2$ intensity levels.

Fig. 2. Weighted graphs with $\hat{K} = 2$ intensity levels: $\hat{\mu}_1^* = 0.1$ (blue edges) and $\hat{\mu}_2^* = 0.3$ (red edges). In each graph the node v_i represents the variable i in the 4-dimensional VAR(1), a clockwise-oriented edge from node j to node i represents a non-null coefficient for the variable $y_{j,t-1}$ in the i th equation of the VAR. Panel (a): weighted graph $G = (V, E, C)$ (top) with vertex set $V = \{v_1, v_2, v_3, v_4\}$, edge set $E = \{e_1, e_2, e_3, e_4\}$, where $e_1 = \{v_1, v_2\}$, $e_2 = \{v_1, v_3\}$, $e_3 = \{v_1, v_4\}$, $e_4 = \{v_2, v_3\}$, $e_5 = \{v_3, v_1\}$, $e_6 = \{v_4, v_2\}$, $e_7 = \{v_4, v_3\}$ and the adjacency matrix A (middle) and the weights matrix C (bottom). Panel (b): the subgraph $G_1 = (V, E_1)$ with $E_1 = \{e_1, e_6, e_7\}$ induced by edges of intensity level $\hat{\mu}_1^* = 0.3$. Panel (c): the subgraph $G_2 = (V, E_2)$ with $E_2 = \{e_2, e_3, e_4, e_5\}$ induced by edges with intensity $\hat{\mu}_2^* = 0.1$. Panel (d): graph with two communities (red $V_1 = \{v_1, v_2\}$ and blue $V_2 = \{v_3, v_4\}$). Panel (e): graph with a hub (node v_1). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

A better understanding of the nodes centrality and of the risk connectedness can be achieved by exploiting the information in the edge weights (e.g., see [Bianchi et al., 2018](#)). One can view the number of edges as more important than the weights, so that the presence of many edges with any weight might be considered more important than the total sum of edge weights. However, edges with large weights might be considered to have a much greater impact than edges with only small weights (e.g., see [Opsahl et al., 2010](#)). BNP-Lasso VAR can be used to better analyse the contribution of the weights to nodes centrality. The posterior partition of the edges (VAR coefficients) into a finite number of groups allows us to represent $G = (V, E, C)$ as a $\hat{K}$ -multilayer graph $G = (V, \{E_k\}_{k=1}^{\hat{K}}, \{C_k\}_{k=1}^{\hat{K}})$, where the k th layer is the sub-graph $G_k = (V, E_k)$, with the set $E_k = \{\{i, j\} \in E; c_{ij} = \hat{\mu}_k^*\}$ of all edges with intensity $\hat{\mu}_k^*$ (see [Kivelä et al. \(2014\)](#)). Panels (b) and (c) of [Fig. 2](#) provide an example with two layers.

Node centrality and network connectivity can be now analysed exploiting the multilayer representation. The out-degree, ω_i^+ , and weighted out-degree ([Opsahl et al., 2010](#)), σ_i^+ , of the node i can be decomposed along the $\hat{K}$ layers

$$\omega_i^+ = \sum_{j=1}^m a_{ji} = \sum_{k=1}^{\hat{K}} \omega_{i,k}^+, \quad \text{where } \omega_{i,k}^+ = \sum_{j=1}^m a_{ji} \mathbb{I}(c_{ij} = \hat{\mu}_k^*), \quad (17)$$

$$\sigma_i^+ = \sum_{j=1}^m c_{ji} = \sum_{k=1}^{\hat{K}} \sigma_{i,k}^+ = \sum_{k=1}^{\hat{K}} \hat{\mu}_k^* \omega_{i,k}^+, \quad \text{where } \sigma_{i,k}^+ = \sum_{j=1}^m c_{ji} \mathbb{I}(c_{ij} = \hat{\mu}_k^*). \quad (18)$$

Similarly, it is possible to decompose the vertex in-degree and total-degree measures. These decompositions are useful in quantifying the contribution of each node to the system-wide connectedness. In our financial and macroeconomic applications, we find that nodes with large out-degree become less relevant when edge intensity is considered. To exemplify, node v_1 has the largest out-degree, $\omega_1^+ = 3$, in Panel (a), [Fig. 2](#), but is not central in the layer G_1 (Panel (b)). In this layer node v_1 has out-degree $\omega_{1,2}^+ = 1$ whereas node v_4 has the largest out-degree $\omega_{4,2}^+ = 2$. Following the weighted degree measure, node v_4 is central in G and in G_1 with out-degree, $\sigma_4^+ = 0.6$ and $\sigma_{4,1}^+ = 0.6$, respectively, but it is not central in G_2 , where its degree is $\sigma_{4,2}^+ = 0$.

The connectivity analysis presented above relies on the number of direct connections among nodes. In order to gain a broader idea of the network connectivity patterns, indirect connections need to be considered. The notion of path s_{ij} between two vertices i and j , called endvertices, accounts for the indirect connections between them. A path s_{ij} is defined by ordered sequences of distinct vertices $V(s_{ij}) = \{v_0, v_1, \dots, v_l\}$ and edges $E(s_{ij}) = \{e_1, \dots, e_l\} \subset V$, with $e_1 = (v_0, v_1)$, $e_l = (v_{l-1}, v_l)$, and $v_0 = i$ and $v_l = j$. The number of edges $|E(s_{ij})| = l$ in a path is called path length.

The distance between two nodes i and j is defined as the length of the shortest path, i.e. $d(i, j) = \min\{|E(s_{ij})|\}$, s.t. s_{ij} is a path between i and j . It can be used to define the d -step ego network of a node $i \in E$, that is a sub-graph of G induced by i and all nodes with distance less than d from i , $V(i) = \{j \in E, \text{ s.t. } d(i, j) \leq d\}$ and $E(i) = \{\{i, j\} \in E, \text{ s.t. } i, j \in V(i)\}$.

$V(i) \cup \{j, i\} \in E$, s.t. $i, j \in V(i)$. In our empirical applications, ego-networks will be used to show the heterogeneity in the linkages strength of a given node.

The distance is a local connectivity measure, which does not account for the whole interconnected topology, thus we consider the average path length (APL), that is a global connectivity measure defined as

$$APL(G) = \frac{1}{m(m-1)} \sum_{i=1}^m \sum_{j=1}^m d(i, j). \quad (19)$$

A low APL may indicate a random graph structure of either Erdős–Rényi type or Watts–Strogatz type (see [Newman \(2010\)](#)). APL assumes shocks follow the shortest possible path, but economic shocks are in general unlikely to be restricted to follow specific paths and are also likely to have feedback effects.

A measure which accounts for all possible paths connecting two nodes is the eigenvector centrality (modified Bonacich's beta centrality, see [Bianchi et al. \(2018\)](#))

$$h(G) = \left(I_m - \frac{1}{(m-1)^2} AA' \right)^{-1} \left(I_m - \frac{1}{(m-1)} A \right) \iota, \text{ with } \iota = (1, \dots, 1)'. \quad (20)$$

If a VAR of order $p > 1$ is used, then not only direct effects from $y_{i,t-p}$ to $y_{j,t}$ should be considered as in Granger causality, but also indirect causal effects from $y_{i,t-p}$ to $y_{j,t}$ through $y_{k,t-l}$, $1 < l < p$, as in Sims causality ([Eichler, 2013](#)). Moreover, VAR shocks can be correlated and their identification can be challenging in high dimension. In order to solve all these issues, [Diebold and Yilmaz \(2014\)](#) propose a global connectivity measure, which relies on a generalized identification framework and on the H -step ahead forecast error variance

$$\Theta(G) = \sum_{i,j=1, i \neq j}^m \sum_{h=0}^{H-1} \frac{(\mathbf{e}_i' \Phi_h \Sigma \mathbf{e}_j)^2}{m \sigma_{jj} c_i}, \quad H = 1, 2, \dots, \quad (21)$$

where c_i is a normalizing constant and Φ_h satisfies the recursion $\Phi_h = B_1 \Phi_{h-1} + B_2 \Phi_{h-2} + \dots + B_p \Phi_{h-p}$ with $\Phi_0 = I_m$ and $\Phi_h = O_m$ for $h < 0$. The second and the third term of the decomposition

$$(\mathbf{e}_i' \Phi_h \Sigma \mathbf{e}_j)^2 = \left(\sum_{l=1}^{h-1} \mathbf{e}_i' B_l \Phi_{h-l} \Sigma \mathbf{e}_j \right)^2 + 2 \left(\sum_{l=1}^{h-1} \mathbf{e}_i' B_l \Phi_{h-l} \Sigma \mathbf{e}_j \right) (\mathbf{e}_i' B_h \Sigma \mathbf{e}_j) + (\mathbf{e}_i' B_h \Sigma \mathbf{e}_j)^2 \quad (22)$$

provide the contribution of the h -lag connectivity structure to the global measure. See [Appendix B](#) for a derivation. The presence of nodes with large out-degree in the h th lag graph can impact significantly on the system-wide connectedness. Both global measures will be used in combination with multilayer decomposition of the coloured graph.

Finally, other two useful notions are the ones of community and hub. Assume the vertex set V is partitioned in a sequence of disjoint subsets $V_1, \dots, V_K$, then an element V_k of the partition is called community. An example of community structured graph is in Panel (d) of [Fig. 2](#). By construction, the partition of the nodes induces an edges partition (e.g., communities V_1 and V_2 in Panel (d)). Since our BNP-Lasso prior allows for any type of partition of the edge set, then typical community structures appearing in stochastic block models and modular graphs (see [Newman, 2010](#)) can be extracted with our model. Panel (e) of [Fig. 2](#) provides an example of hub, that is a node with a larger number of edges in comparison with other nodes in the graph. Node v_1 is an hub in the graph of Panel (e), since it has out-degree 3 whereas the other nodes have null out-degree.

3.3. Simulation experiments

We study the goodness of fit of our model and, following the standard practice in BNP analysis (e.g., see [Griffin and Steel \(2006\)](#), [Griffin and Steel \(2011\)](#), [Griffin and Kalli \(2018\)](#)), we simulate data from a parametric model, which is in the family of the likelihood of our model. We consider a VAR(1) model

$$\mathbf{y}_t = B \mathbf{y}_{t-1} + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \stackrel{i.i.d.}{\sim} \mathcal{N}_m(\mathbf{0}, \Sigma) \quad t = 1, \dots, 100,$$

where the matrix B represents the collection of pairwise relationships between the time series in the panel (i.e., countries or institutions). A non-zero coefficient indicates there is a linkage between two nodes in the network.

We consider three dimension settings for $\mathbf{y}_t$ ($m = 20$ (small), $m = 40$ (medium) and $m = 80$ (large)) and different parameter settings. For each setting, we conduct 50 experiments, generating 50 independent datasets and estimating our BNP-Lasso VAR model. In all the experiments, we have chosen the hyperparameters as in [Section 2.2](#), and set the degree of freedom parameter, $b_0 = 3$ and the scale matrix, $L = I_n$. We iterate 5000 times the Gibbs sampler described in [Section 3](#), and after discarding a burn-in sample and applying thinning, we use the MCMC samples of $\mathcal{E}$ to estimate the network adjacency matrix and the samples of D to estimate the number of colours and classify the edge intensity. We provide in [Fig. 3](#) the typical causality network estimated in four experiments. See Supplementary Material for further results.

As one can see from [Fig. 3](#), two experimental settings have been considered for B , which reflect two typical network configurations detected in many empirical studies. The first setting represents a system which is robust to random

Fig. 3. Weighted network estimated in four Monte Carlo experiments, for different model dimensions $m = 20$ (panel (a)), $m = 40$ (panel (b)) and $m = 80$, with block entries (panel (c)) and with random entries (panel (d)). Blue edges mean negative weights and red ones represent positive weights, while the edges are clockwise-oriented. In each graph the nodes represent the n variables of the VAR model, and a clockwise-oriented edge from node j to node i represents a non-null coefficient for the variable $y_{j,t-1}$ in the i th equation of the VAR. Blue edges represent negative coefficients, while red edges represent positive coefficients. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

failures, but vulnerable to targeted attacks. We consider many nodes (financial institutions) with low degree, and a small number of nodes, known as hubs (e.g. Barabási and Albert, 1999; Acemoglu et al., 2012), or systemically important institutions (e.g. Billio et al., 2012; Battiston et al., 2012), with a number of linkages that exceeds the average. The presence of hubs characterizes periods of high systemic risk and when the default of the largest institution (hubs) can lead to the breakdown of the whole system. Hubs have been detected in the S&P100 return networks (e.g., Bianchi et al., 2018) and in EuroStoxx-600 realized volatility networks (e.g., Ahelgebey et al., 2016b). Also, we assume nodes cluster into groups internally densely connected with no connections among them. This feature is known as community structure, or modularity (Newman, 2006; Leicht and Newman, 2008) and has been detected not only in financial-return networks, but also in bank-firm (Bargigli and Gallegati, 2013), regional banks (Puliga et al., 2016) and inter-bank market (De Souza et al., 2016), networks. The presence of communities structure is a relevant feature, since it can prevent the spread of contagion to the whole system. The setting described corresponds to a block-diagonal matrix B with (4×4) blocks B_j ($j = 1, \dots, m/4$) on the main diagonal representing the communities, and i.i.d. elements within each block, $b_{lk,j} \sim \mathcal{U}(-1.4, 1.4)$, $l, k = 1, \dots, 4$. The matrix B is checked for the weak stationarity condition of the VAR.

The second setting represents a stable financial or economic condition, which reflects in a network with few connections, randomly distributed across node pairs. This network configuration has been observed during periods of low systemic risk. More specifically, we consider a random matrix B of dimension (80×80) , select randomly 150 elements and draw their values from the uniform $\mathcal{U}(-1.4, 1.4)$. The remaining 6250 elements are set to zero. The matrix generated is then checked for weak stationarity. In addition to the two settings previously described, we consider a VAR(4) parametrized on the posterior mean of the macroeconomic application presented in Section 4. The network configuration at each lag, given in Fig. 5, exhibits edge heterogeneity and hubs.

The burn-in period and the thinning rate have been chosen following a graphical dissection of the posterior progressive averages and the standard convergence MCMC diagnostics. Table 1 shows the diagnostics averaged over 50 independent experiments.

Since in a data augmentation framework, the mixing of the MCMC may depend on the autocorrelation in both parameter and latent variable samples, we focus on λ_{ij} , u_{ij} and β_{ij} . The diagnostic for λ_{ij} , after removing 500 burn-in iterations, indicates the chain has converged. The inefficiency factor and the autocorrelation at lag 10 on the original sample suggest high levels of dependence in the sample, which can be mitigated by thinning the chains. We keep only every 5th sample and obtain substantial reduction of sample autocorrelation. See Supplementary Material for further details.

Table 1

Geweke (CD) and Kolmogorov–Smirnov (KS) convergence test; inefficiency factor (INEFF) and autocorrelation (ACF) computed at lag 10 for the ℓ^2 norm of λ_{ij} for each experiments on the whole (before) and thinned (after) MCMC samples. All quantities are averages over 50 independent experiments.

	CD	KS	INEFF		ACF(10)	
			before	after	before	after
$m = 20$	−0.629	0.3667	13.5092	7.0346	0.1724	0.0188
$m = 40$	−1.197	0.002	21.3259	8.5001	0.3366	0.1334
$m = 80$ block	−1.354	0.5806	20.1303	9.6897	0.3147	0.1229
$m = 80$ random	2.451	0.0158	17.6933	5.8095	0.267	0.0185

Fig. 4. Boxplot of Mean Square Deviation (MSD) of the estimated VAR coefficients from their true values based on 50 independent Monte Carlo exercises for our Bayesian nonparametric model (BNP), Elastic-net (EN), Bayesian Lasso (BLA) and Stochastic Search Variable Selection (SSVS). On each box, the central mark indicates the median, while the bottom and top edges of the box are the 25th and 75th percentiles, respectively. The whiskers extend to the most extreme points and the outliers are plotted using $\circ$.

We compare our prior (BNP) with the Bayesian Lasso (B-Lasso, [Park and Casella \(2008\)](#)), the Elastic-net (EN, [Zou and Hastie \(2005\)](#)) and Stochastic Search Variable Selection (SSVS) of [George et al. \(2008\)](#). For the SSVS, we assume the default hyperparameters values $\tau_1^2 = 0.0001$, $\tau_2^2 = 4$ and $\pi = 0.5$. Following [Korobilis \(2016\)](#), we use the Mean Square Deviation (MSD) for measuring the performance of the four different priors. See Supplementary Material for a more detailed description. For each parameter setting we generate 50 independent datasets and estimate the models under comparison. [Fig. 4](#) shows the quartiles and median of the MSD statistics based on 50 Monte Carlo experiments by means of box plots. In all settings, it is likely that the BNP-Lasso works better compared to other shrinkage methods in recovering the true VAR coefficients since the interquartile ranges (i.e. the boxes) do not overlap. However, one can better appreciate a superior relative performance of the BNP-Lasso when model dimension increases (from $m = 20$ (left) to $m = 80$ (right)).

In our simulation results, the BNP-Lasso shrinking effect is a midway between SSVS and Bayesian Lasso. It is smaller with respect to B-Lasso and Elastic-Net and comparable to SSVS when the true β_{ij} is nonnull and it is larger than SSVS when $\beta_{ij} = 0$. The reduction in MSD is largely a result of decreased bias. The BNP-Lasso clustering effect tends to identify the correct coefficient clustering reducing the bias. In the Supplementary Material, the comparison in other settings ([Fig. S.8–S.11](#)) confirms our findings. Moreover, [Fig. S.12](#) shows that all shrinkage methods are over-performing the Minnesota prior, thus confirming the results given in [Korobilis and Pettenuzzo \(2019\)](#) for a VAR of dimension $m = 20$. Our findings extend their results to higher dimensions ($m = 40$ and $m = 80$) and to other shrinkage methods (BNP-Lasso, Elastic-Net). Also, we find that the performance of the Minnesota prior deteriorates when the model dimension increases.

4. Empirical applications

4.1. Measuring business cycle connectedness

Following the literature on international business cycles we consider a multi-country macroeconomic dataset (e.g., see [Kose et al., 2003](#); [Francis et al., 2017](#); [Kaufmann and Schumacher, 2017](#)) and extract a network of linkages between the cycles of the OECD countries, by applying BNP-Lasso VAR.

We consider the quarterly GDP growth rate (logarithmic first differences) from 1961:Q1 to 2015:Q2, for a total of $T = 215$ observations. Due to missing values in some series, we choose a subset of high industrialized countries: Rest of the World (Australia, Canada, Japan, Mexico, South Africa, Turkey, United States) and Europe (Austria, Belgium, Denmark,

Fig. 5. Weighted Granger networks for the GDP growth rates of 25 OECD countries for the sample period 1961Q1 to 2015Q2. We consider different lags: (a) $t - 1$, (b) $t - 2$, (c) $t - 3$, (d) $t - 4$. Node size indicates node degree, blue edges represent negative weights and red ones positive weights. At the lag l , clockwise-oriented edge from node j to node i represents a non-null coefficient for the variable $y_{j,t-l}$ in the i th equation of the VAR. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Finland, France, Germany, Greece, Ireland, Iceland, Italy, Luxembourg, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, United Kingdom).

The posterior probability of a macroeconomic linkage is 0.13, which provides evidence of sparsity in the network and indicates that a small proportion of linkages is responsible for connectedness in business cycle risk. The posterior number of clusters and the co-clustering matrix (see Supplementary Material) reveals the existence of three levels of linkage intensity, customarily called “negative”, “positive” and “strong positive”. Fig. 5 draws the coloured networks at different time lags with negative (blue) and positive (red) weights. In a multilayer graph, each intensity level (or edge colour) identifies the set of nodes and edges belonging to a specific layer. The network connectivity and nodes centrality can be investigated either at the global level, or for each single layer.

The network topology analysis in Table 2 reveals a significant connectivity (graph density and average degree) and complexity (average path length) at all lags. Consequently, each lag-specific network contributes to the build up of the global connectedness as shown in Eq. (22). It becomes important to study the dynamical structure of the directional connectedness received from other countries (in-degree) or transmitted to other countries (out-degree). The multilayer analysis (coloured edges in Fig. 5 and Table 2) will be used to identify the countries, who mostly contribute at each lag to the system-wide connectedness (Diebold and Yilmaz, 2014).

At the first two lags, core European countries (Austria, Belgium, Finland, France, Germany, Netherlands) are mainly receivers, whereas the periphery European countries (Greece, Ireland, Italy, Portugal, Spain) transmit the highest percentage of shocks. In the other lags, core countries receive and transmit a high percentage of shocks.

Decomposing the degree by the intensity levels (see Eq. (17)), we see that the average degree at the first lag (2.92) is mainly driven by positive and strong positive edge intensities (2.64). The 88% of the linkages provides a positive contribution to the connectedness and originates from Japan (the central country) and European countries. The most central country in the blue layer is the Netherlands, those negative linkages could contribute to reduce the total connectedness.

At the second lag, the red and blue layers have similar average degree. Austria is central in the overall network exhibiting both negative and positive linkages. Nevertheless, if one considers only positive linkages, the central country is Greece.

At the third and the fourth lag, negative linkages prevail over positives. European countries have positive linkage strength within them, thus contributing to a marginal increase of the total connectedness. A reduction of connectedness comes from the negative linkages within rest-of-the-world countries and between them and European countries.

Table 2

Statistics for single layers networks of the GDP growth networks of 25 OECD countries for the sample period 1961Q1 to 2015Q2. The sub-graphs (layers) are induced by positive (red) and negative (blue) edge intensity. The average path length represents the average graph-distance between all pairs of nodes. Connected nodes have graph distance 1.

	Links	Avg degree	Density	Avg path length
$t - 1$	73	2.92	0.122	3.423
$t - 1$ blue	7	0.28	0.012	1.143
$t - 1$ red	66	2.64	0.11	2.587
$t - 2$	45	1.80	0.075	3.211
$t - 2$ blue	22	0.88	0.037	2.634
$t - 2$ red	23	0.92	0.038	2.033
$t - 3$	41	1.64	0.069	2.479
$t - 3$ blue	25	1.00	0.042	2.268
$t - 3$ red	16	0.64	0.027	1.667
$t - 4$	52	2.08	0.086	2.718
$t - 4$ blue	32	1.28	0.053	1.435
$t - 4$ red	20	0.80	0.033	1.791

In order to validate our model, we study its forecasting abilities. In the comparison in Supplementary Material, we find that BNP-Lasso on this specific dataset over-performs Elastic-Net (EN), Bayesian Lasso (B-Lasso) and SSVS. These findings suggest that the BNP-Lasso is useful not only for a better understanding of the connectedness but could also be employed for forecasting purposes in macroeconomics (e.g. [McCracken and Ng, 2016](#)). Since forecasting is behind the scope of the paper, we leave it for further research.

4.2. Risk connectedness in European financial markets

Based on the literature on risk connectedness among financial institutions and markets (see [Hautsch et al. \(2015\)](#) and [Diebold and Yilmaz \(2014\)](#)), we construct daily realized volatilities using intraday high–low–close price indexes of 118 institutions of the Euro Stoxx 600 obtained from Datastream, from January 3, 2005 to September 19, 2014. The dataset consists of 42 banks, 31 financial services, 31 insurance companies and 22 real estates from Austria, Belgium, Finland, France, Germany, Greece, Ireland, Italy, Luxembourg, Netherlands, Portugal and Spain. As in [Ahelgebey et al. \(2016b\)](#), we build the realized volatility as $RV_t = 0.5 (\log H_t - \log L_t)^2 - (2 \log 2 - 1) (\log C_t - \log C_{t-1})^2$, where H_t , L_t and C_t denote the high, low and closing price of a given stock on day t , respectively and consider a VAR(1) model.

There is a strong evidence of sparsity with a probability of 0.16 of edge existence ([Fig. 6](#)). Based on the posterior number of clusters, we identify five levels of linkage strengths, customarily called "negative" (blue), "weak positive" (green), "positive" (orange) and "strong positive" (red). Negative linkages imply some institutions react with a volatility decrease to positive volatility shocks, thus reducing systemic risk. Top-left plot in [Fig. 6](#) shows the coloured realized-volatility network. The node size increases with the node eigenvector centrality, which is evaluated on the overall network without accounting for the linkage intensity. We find that insurance companies and banks (e.g., AGEAS and Bank of Ireland) are central and present different type of linkages (e.g., see AGEAS two-step ego network in the bottom plot). Also, a community structure can be detected with some small communities of same countries banks, such as the one related to the Italy (Monte dei Paschi and other Popular Banks) and Greece (Alpha Bank, National Bank of Greece, Bank of Piraeus).

Evaluating the node eigenvector centrality in the subgraphs of positive and strong positive linkages, we found that also real estates sector is central (e.g. Immofiz, top-right plot). Our approach thus helps in better understanding the role of financial institutions in the build-up of the systemic risk, revealing different aspects of their activity and in particular for insurance and real estate industries.

5. Conclusions

This paper introduces a novel Bayesian nonparametric Lasso prior (BNP-Lasso) for VAR models, which combines Dirichlet process and Bayesian Lasso priors. The two-stage hierarchical prior allows for clustering the VAR coefficients and for shrinking them either toward zero or random locations, thus inducing sparsity in the VAR coefficients. The simulation studies illustrate the good performance of our model in high dimension settings, compared to some existing priors, such as the Stochastic Search Variable Selection, Elastic-net and simple Bayesian Lasso.

The BNP-Lasso VAR is well suited for extracting networks which account for some stylized facts observed in financial and macroeconomic networks such as community structures and linkage heterogeneity. The linkages partition can be easily obtained with the BNP-Lasso framework and the extraction of a coloured network allows for analysing the centrality of the nodes in the intensity-specific subgraphs.

In the macroeconomic (financial) application, our BNP-Lasso VAR helps in understanding the centrality of countries (institutions) and their role in the build-up of risk. In all applications, there is evidence of network sparsity and of different

Fig. 6. Weighted financial networks of realized volatility for 118 financial institutions (nodes) of the EuroStoxx 600 from January 3, 2005 to September 19, 2014. We show the pairwise directed edges extracted with BNP-Lasso VAR(1). Edge colours indicate five levels of linkage strengths: "negative" (blue), "weak positive" (green), "positive" (orange) and "strong positive" (red). The edges are clockwise-oriented and the node size is based on the node degree centrality (top-left) or eigenvector centrality (top-right). Bottom: AGEAS two-step ego network. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

levels of linkage intensity, meaning that few linkages could be responsible for the connectedness. We show evidence that some nodes are central when the intensity level is considered, even if they are not central in the network. Also, few nodes possess linkages with different intensities whereas most of the nodes have only one type of linkages.

Appendix A. Further details on prior specification

A.1. Normal–Gamma distribution

A normal–gamma random variable $X \sim \mathcal{NG}(\mu, \gamma, \tau)$ has probability density function

$$f(x|\mu, \gamma, \tau) = \frac{\tau^{\frac{2\gamma+1}{4}} |x - \mu|^{\gamma - \frac{1}{2}}}{2^{\gamma - \frac{1}{2}} \sqrt{\pi} \Gamma(\gamma)} K_{\gamma - \frac{1}{2}}(\sqrt{\tau} |x - \mu|),$$

where $K_{\gamma}(\cdot)$ represents the modified Bessel function of the second kind with index γ , $\mu \in \mathbb{R}$ is the location parameter, $\gamma > 0$ is the shape parameter and $\tau > 0$ is the scale parameter. The normal–gamma distribution has the double exponential distribution as a special case for $\gamma = 1$ and can be represented as a scale mixture of normals,

$$\mathcal{NG}(x|\mu, \gamma, \tau) = \int_0^{+\infty} \mathcal{N}(x|\mu, \lambda) \mathcal{Ga}(\lambda|\gamma, \tau/2) d\lambda, \quad (\text{A.1})$$

where $\mathcal{Ga}(\cdot|a, b)$ denotes a gamma distribution with mean a/b and variance a/b^2 .

A.2. Gamma scale-shape distribution

A Gamma scale-shape random vector $(X, Y) \sim \mathcal{GS}(\nu, p, s, n)$ has pdf

$$g(x, y|\nu, p, s, n) \propto \tau^{\nu x - 1} p^{x - 1} \exp\{-sy\} \frac{1}{\Gamma(x)^n}, \quad (\text{A.2})$$

Fig. A.1. Probability density function, $g(x)$, for sparse (dashed line) and non-sparse (solid line) case, respectively.

with parameters $\nu > 0$, $p > 0$, $s > 0$ and $n > 0$ (see [Miller \(1980\)](#)). The pdf in Eq. (A.2) factorizes as $g(x, y) = g(y|x)g(x)$, where

$$g(y|x) = \frac{g(x, y)}{g(x)} = \frac{\tau^{\nu x - 1} e^{-sy}}{\Gamma(\nu x)} s^{\nu x}$$

that is a density of a $\mathcal{Ga}(\nu x, s)$, and

$$g(x) = \int_0^\infty g(x, y) dy = C \frac{\Gamma(\nu x) p^{x-1}}{\Gamma(x)^n s^{\nu x}}$$

is the marginal density with normalizing constant C such that $\int_0^\infty g(x) dx = 1$. We show in [Fig. A.1](#) the density function, $g(x)$, for the two parameter settings used in the empirical application: $\nu = 30$, $s = 1/30$, $p = 0.5$ and $n = 18$ (dashed line) and $\nu = 3$, $s = 1/3$, $p = 0.5$ and $n = 10$ (solid line).

A.3. Hyper-inverse Wishart distribution

A Hyper-inverse Wishart random variable $\Sigma \sim \mathcal{HIW}_G(b, L)$ has pdf

$$p(\Sigma) = \prod_{P \in \mathcal{P}} p(\Sigma_P) \left(\prod_{S \in \mathcal{S}} p(\Sigma_S) \right)^{-1},$$

where $\mathcal{S} = \{S_1, \dots, S_{n_S}\}$ and $\mathcal{P} = \{P_1, \dots, P_{n_P}\}$ are the set of separators and of prime components, respectively, of the graph G . In particular, b is the degree of freedom, L is the scale matrix and

$$p(\Sigma_P) \propto |\Sigma_P|^{-(b+2\text{Card}(P))/2} \exp \left\{ -\frac{1}{2} \text{tr}(\Sigma_P^{-1} L_P) \right\},$$

with L_P the positive-definite symmetric diagonal block of L corresponding to Σ_P .

A.4. Dirichlet process prior

The Dirichlet Process, $\text{DP}(\tilde{\alpha}, H)$, can be defined by using the stick-breaking representation ([Sethuraman, 1994](#)) given by:

$$\mathbb{P}_i(\cdot) = \sum_{j=1}^{\infty} w_{ij} \delta_{\{\theta_{ij}\}}(\cdot), \quad i = 1, \dots, M.$$

Following the definition of the dependent stick-breaking processes (see [MacEachern \(1999, 2001\)](#)), the atoms θ_{ij} are i.i.d. sequences of random elements with probability distribution H ($\theta_{ij} \stackrel{\text{i.i.d.}}{\sim} H$); and the weights w_{ij} are determined through the stick-breaking construction, for $j = 1$, $w_{i1} = v_{i1}$ and for $j > 1$

$$w_{ij} = v_{ij} \prod_{k=1}^{j-1} (1 - v_{ik}), \quad i = 1, \dots, M$$

with $v_j = (v_{1j}, \dots, v_{Mj})$ independent random variables taking values in $[0, 1]^M$ distributed as a $\mathcal{Be}(1, \tilde{\alpha})$ such that $\sum_{j \geq 1} w_{ij} = 1$ almost surely for every i .

Appendix B. Proof of the results in the paper

Proof of result in Eq. (11). Let $K(\beta_i|\theta_i)$ be the joint density kernel of the sequence $\beta_{ij} \stackrel{\text{ind}}{\sim} \mathcal{NG}(\beta_{ij}|\theta_{ij}^*)$, $j = 1, \dots, n_i$ and $\theta_{ij}^*|\mathbb{Q}_i \stackrel{\text{i.i.d.}}{\sim} \mathbb{Q}_i$, the atoms of the random probability measure $\mathbb{Q}_i$. Let $f_i(\beta_i|\mathbb{Q}_i)$ be the random probability density function obtained by integrating out the random measure $\mathbb{Q}_i$ (Lo, 1984):

$$f_i(\beta_i|\mathbb{Q}_i) = \int K(\beta_i|\theta_i)\mathbb{Q}_i(d\theta_i). \quad (\text{B.1})$$

Using the definition of $\mathbb{Q}_i$ as a convex combination of a sparse component and a DPP in Eq. (B.1), we obtain

$$f_i(\beta_i|\mathbb{Q}_i) = \pi_i \int K(\beta_i|\theta_i)\mathbb{P}_0(d\theta_i) + (1 - \pi_i) \int K(\beta_i|\theta_i)\mathbb{P}_i(d\theta_i) \quad (\text{B.2})$$

We use the stick-breaking representation of the DPP (Appendix A) and get

$$\begin{aligned} f_i(\beta_i|\mathbb{Q}_i) &= \pi_i \int \mathcal{NG}(\beta_i|\mu, \gamma, \tau)\mathbb{P}_0(d(\mu, \gamma, \tau)) + (1 - \pi_i) \int \mathcal{NG}(\beta_i|\mu, \gamma, \tau)\mathbb{P}_i(d(\mu, \gamma, \tau)) \\ &= \pi_i \mathcal{NG}(\beta_i|0, \gamma_0, \tau_0) + (1 - \pi_i) \sum_{k=1}^{\infty} w_{ik} \mathcal{NG}(\beta_i|\mu_{ik}, \gamma_{ik}, \tau_{ik}) \end{aligned} \quad (\text{B.3})$$

By defining,

$$\check{w}_{ik} = \begin{cases} \pi_i, & k = 0, \\ (1 - \pi_i)w_{ik}, & k > 0, \end{cases} \quad \check{\theta}_{ik} = \begin{cases} (0, \gamma_0, \tau_0), & k = 0, \\ (\mu_{ik}, \gamma_{ik}, \tau_{ik}), & k > 0. \end{cases}$$

we have the infinite mixture representation

$$f_i(\beta_i|\mathbb{Q}_i) = \sum_{k=0}^{\infty} \check{w}_{ik} \mathcal{NG}(\beta_i|\check{\theta}_{ik}). \quad \square$$

Proof of result in Eq. (13). We introduce a set of slice latent variables, u_{ij} , $j = 1, \dots, n_i$, which allows us to represent the full conditional of β_{ij} as follows,

$$\begin{aligned} f_i(\beta_{ij}, u_{ij}|\mu_i, \gamma_i, \tau_i, w_i) &= \pi_i \sum_{k=0}^{\infty} \mathbb{I}(u_{ij} < \check{w}_{ik}) \mathcal{NG}(\beta_{ij}|\mu_{ik}, \gamma_{ik}, \tau_{ik}) + \\ &+ (1 - \pi_i) \sum_{k=1}^{\infty} \mathbb{I}(u_{ij} < w_{ik}) \mathcal{NG}(\beta_{ij}|\mu_{ik}, \gamma_{ik}, \tau_{ik}) \\ &= \pi_i \mathbb{I}(u_{ij} < \check{w}_0) \mathcal{NG}(\beta_{ij}|\mu_0, \gamma_0, \tau_0) + (1 - \pi_i) \sum_{k=1}^{\infty} \mathbb{I}(u_{ij} < w_{ik}) \mathcal{NG}(\beta_{ij}|\mu_{ik}, \gamma_{ik}, \tau_{ik}), \end{aligned}$$

where $\check{w}_{ik} = \check{w}_0 = 1$ if $k = 0$ and $\check{w}_{ik} = 0$ for $k > 0$ and, for simplicity of notations, we denote $(0, \gamma_{i0}, \tau_{i0}) = (0, \gamma_0, \tau_0)$.

The slice variables allow us to represent the infinite mixture as a finite mixture conditionally on a random number of components. More specifically, define the set

$$\mathcal{A}_{w_i}(u_{ij}) = \{k : u_{ij} < w_{ik}\}, \quad j = 1, \dots, n_i,$$

then it can be proved that the cardinality of $\mathcal{A}_{w_i}$ is almost surely finite. Posterior draws for u_{ij} are easily obtained by slice sampling.

The conditional joint pdf for β_{ij} and u_{ij} , $f_i(\beta_{ij}, u_{ij}|\mu_i, \gamma_i, \tau_i, w_i)$, is equal to

$$\pi_i \mathbb{I}(u_{ij} < \check{w}_0) \mathcal{NG}(\beta_{ij}|\mu_0, \gamma_0, \tau_0) + (1 - \pi_i) \sum_{k \in \mathcal{A}_{w_i}(u_{ij})} \mathcal{NG}(\beta_{ij}|\mu_{ik}, \gamma_{ik}, \tau_{ik}).$$

We iterate the data augmentation principle and for each f_i we introduce two allocation variables ξ_{ij} and d_{ij} , associated with the sparse and non-sparse components, respectively, of the random measure $\mathbb{Q}_i$. The joint pdf is

$$\begin{aligned} f_i(\beta_{ij}, u_{ij}, d_{ij}, \xi_{ij}) &= (\pi_i \mathbb{I}(u_{ij} < \check{w}_{d_{ij}}) \mathcal{NG}(\beta_{ij}|\mu_0, \gamma_0, \tau_0))^{1-\xi_{ij}} \cdot \\ &((1 - \pi_i) \mathbb{I}(u_{ij} < w_{id_{ij}}) \mathcal{NG}(\beta_{ij}|\mu_{id_{ij}}, \gamma_{id_{ij}}, \tau_{id_{ij}}))^{\xi_{ij}}. \end{aligned}$$

From (4), we de-marginalize the Normal–Gamma distribution by introducing a latent variable λ_{ij} for each β_{ij} and conclude that the joint distribution is

$$f_i(\beta_{ij}, \lambda_{ij}, u_{ij}, d_{ij}, \xi_{ij}) = \left(\mathbb{I}(u_{ij} < \tilde{w}_{d_{ij}}) \mathcal{N}(\beta_{ij} | \mathbf{0}, \lambda_{ij}) \mathcal{G}a(\lambda_{ij} | \gamma_0, \tau_0/2) \right)^{1-\xi_{ij}} \cdot \\ \left(\mathbb{I}(u_{ij} < w_{id_{ij}}) \mathcal{N}(\beta_{ij} | \mu_{id_{ij}}, \lambda_{ij}) \mathcal{G}a(\lambda_{ij} | \gamma_{id_{ij}}, \tau_{id_{ij}}/2) \right)^{\xi_{ij}} \pi_i^{1-\xi_{ij}} (1 - \pi_i)^{\xi_{ij}}. \quad \square$$

Proof of result in Eq. (22). For $h \leq p$, Φ_h satisfies $\Phi_h = R_h + B_h$, where $R_h = \sum_{l=1}^{h-1} B_l \Phi_{h-l}$. By the properties of the Hadamard product $\circ$, it follows

$$(\mathbf{e}'_i \Phi_h \Sigma \mathbf{e}_j)^2 = \mathbf{e}'_i ((R_h + B_h) \Sigma) \circ ((R_h + B_h) \Sigma) \mathbf{e}_j = \mathbf{e}'_i ((R_h \Sigma) \circ (R_h \Sigma)) \mathbf{e}_j \\ + \mathbf{e}'_i ((R_h \Sigma) \circ (B_h \Sigma)) \mathbf{e}_j + \mathbf{e}'_i ((B_h \Sigma) \circ (R_h \Sigma)) \mathbf{e}_j + \mathbf{e}'_i ((B_h \Sigma) \circ (B_h \Sigma)) \mathbf{e}_j \\ = (\mathbf{e}'_i (R_h \Sigma) \mathbf{e}_j)^2 + 2 (\mathbf{e}'_i R_h \Sigma \mathbf{e}_j) (\mathbf{e}'_i B_h \Sigma \mathbf{e}_j) + (\mathbf{e}'_i B_h \Sigma \mathbf{e}_j)^2 \quad \square$$

Appendix C. Gibbs sampling details

We introduce for $k \geq 1$ the set of indexes of the coefficients allocated to the k th mixture component, $\mathcal{D}_{ik} = \{j \in 1, \dots, n_i : d_{ij} = k, \xi_{ij} = 1\}$ and the set of the non-empty mixture components, $\mathcal{D}^* = \{k | \cup_i \mathcal{D}_{ik} \neq \emptyset\}$. The number of stick-breaking components is denoted by $D^* = \max_i \{\max_{j \in \{1, \dots, n_i\}} d_{ij}\}$. As noted by Kalli et al. (2011), the sampling of infinitely many elements of Θ and V is not necessarily, since only the elements in the full conditional distributions of (D, Ξ) are needed and the maximum number is $N^* = \max_i \{N_i^*\}$, where N_i^* is the smallest integer such that $\sum_{k=1}^{N_i^*} w_{ik} > 1 - u_i^*$, where $u_i^* = \min_{1 \leq j \leq n_i} \{u_{ij}\}$.

Update the stick-breaking and slice variables V and U

We treat V as three blocks: $V^* = \{V_k : k \in \mathcal{D}^*\}$, $V^{**} = (v_{kD^*+1}, \dots, v_{kN^*})$ and $V^{***} = \{V_k : k > N^*\}$. In order to sample (U, V) , a further blocking is used:

(i) Sampling from the full conditional of V^* , by using

$$f(v_{ij} | \dots) \propto \text{Be} \left(1 + \sum_{j=1}^{n_i} \mathbb{I}(d_{ij} = d, \xi_{ij} = 1), \alpha + \sum_{j=1}^{n_i} \mathbb{I}(d_{ij} > d, \xi_{ij} = 1) \right).$$

for $k \leq D^*$. The elements of (V^{**}, V^{***}) are sampled from the prior $\text{Be}(1, \alpha)$.

(ii) Sampling from full conditional posterior distribution of U

$$f(u_{ij} | \dots) \propto \begin{cases} \mathbb{I}(u_{ij} < w_{1d_{ij}})^{\xi_{ij}} & \text{if } \xi_{ij} = 1, \\ \mathbb{I}(u_{ij} < 1)^{1-\xi_{ij}} & \text{if } \xi_{ij} = 0, \end{cases}$$

Update the mixing parameters λ

Regarding the mixing parameters λ_{ij} , the full conditional posterior distribution is

$$f(\lambda_{ij} | \dots) \propto \lambda_{ij}^{C_{ij}-1} \exp \left\{ -\frac{1}{2} \left[A_{ij} \lambda_{ij} + \frac{B_{ij}}{\lambda_{ij}} \right] \right\} \propto \mathcal{G}i\mathcal{G}(A_{ij}, B_{ij}, C_{ij}),$$

where $\mathcal{G}i\mathcal{G}$ denotes a Generalized Inverse Gaussian with $A_{ij} > 0$, $B_{ij} > 0$ and $C_{ij} \in \mathbb{R}$

$$A_{ij} = [(1 - \xi_{ij})\tau_0 + \xi_{ij}\tau_{id_{ij}}], \quad B_{ij} = [(1 - \xi_{ij})\beta_{ij}^2 + \xi_{ij}(\beta_{ij} - \mu_{id_{ij}})^2], \\ C_{ij} = \left[(1 - \xi_{ij})\gamma_0 + \gamma_{id_{ij}}\xi_{ij} - \frac{1}{2} \right].$$

The matrix $\Lambda_i = \text{diag}\{\lambda_i\}$ has the elements of $\lambda_i = (\lambda_{i1}, \dots, \lambda_{in_i})'$ on the diagonal.

Update the atoms Θ

We consider the sparse and the non-sparse case. In the sparse case, the full conditional distribution of μ_0 is $f(\mu_0 | \dots) = \delta_{\{0\}}(\mu_0)$ and the full conditional distribution of (γ_0, τ_0) is:

$$f((\gamma_0, \tau_0) | \dots) \propto \mathcal{G}\mathcal{S} \left(\gamma_0, \tau_0 | v_0 + \sum_{i=1}^M n_{i,0}, p_0 \prod_{i=1}^M \prod_{j|\xi_{ij}=0} \lambda_{ij}, s_0 + \frac{1}{2} \sum_{i=1}^M \sum_{j|\xi_{ij}=0} \lambda_{ij}, n_0 + \sum_{i=1}^M n_{i,0} \right),$$

where we assume $n_{i,0} = \sum_{j=1}^{n_i} (1 - \xi_{ij}) = n_i - n_{i,1}$ and $n_{i,1} = \sum_{j=1}^{n_i} \xi_{ij}$. In the non-sparse case, we generate samples $(\mu_{ik}, \gamma_{ik}, \tau_{ik})$, $k = 1, \dots, N^*$, by applying a single move Gibbs sampler. The full conditional for μ_{ik} is

$$f(\mu_{ik} | \dots) \propto \mathcal{N}(\mu_{ik} | C, d) \prod_{j|\xi_{ij}=1, d_{ij}=k} \mathcal{N}(\beta_{ij} | \mu_{ik}, \lambda_{ij}) \propto \mathcal{N}(\tilde{E}_k, \tilde{V}_k)$$

with $\tilde{E}_k = \tilde{V}_k \left(\frac{c}{d} + \sum_{j|\xi_{ij}=1, d_{ij}=k} \frac{\beta_{ij}}{\lambda_{ij}} \right)$ and $\tilde{V}_k = \left(\frac{1}{d} + \sum_{j|\xi_{ij}=1, d_{ij}=k} \frac{1}{\lambda_{ij}} \right)^{-1}$, the mean and variance, respectively. The joint conditional posterior of (γ_{ik}, τ_{ik}) is:

$$f((\gamma_{ik}, \tau_{ik}) | \dots) \propto \mathcal{GS} \left(\gamma_{ik}, \tau_{ik} | v_1 + n_{i,1k}, p_1 \prod_{j|\xi_{ij}=1, d_{ij}=k} \lambda_{ij}, s_1 + \frac{1}{2} \sum_{j|\xi_{ij}=1, d_{ij}=k} \lambda_{ij}, n_1 + n_{i,1k} \right),$$

for $k \in \mathcal{D}^*$, where $n_{i,1k} = \sum_{j=1}^{n_i} \xi_{ij} \mathbb{I}(d_{ij} = k)$ and from the prior H for $k \notin \mathcal{D}^*$. In order to draw samples from $\mathcal{GS}$ in both cases, we apply a collapsed Gibbs sampler. Samples from $f(\gamma)$ are obtained by a Metropolis–Hastings (MH) algorithm and samples from $f(\tau | \gamma)$ are obtained from a Gamma distribution.

Update the coefficients β

The full conditional posterior distribution of β is:

$$f(\beta_i | \dots) \propto \mathcal{N}_{n_i}(\tilde{\mathbf{v}}_i, M_i),$$

with mean $\tilde{\mathbf{v}}_i = M_i \left(\sum_t X_t' \Sigma^{-1} \mathbf{y}_t + \Lambda_i^{-1} (\mu_i^* \odot \xi_i) \right)$ and variance $M_i = \left(\sum_t X_t' \Sigma^{-1} X_t + \Lambda_i^{-1} \right)^{-1}$; and $\mu_i^* = (\mu_{id_{i1}}, \dots, \mu_{id_{in_i}})'$, $\xi_i = (\xi_{i1}, \dots, \xi_{in_i})'$.

Update the covariance matrix Σ

By using the sets $\mathcal{S}$ and $\mathcal{P}$ as described in [Appendix A](#) and a decomposable graph, the likelihood of the graphical Gaussian model can be approximated as the ratio between the likelihood in the prime and in the separator components. Thus, the posterior for Σ factorizes as follows:

$$p(\Sigma | \dots) \propto \mathcal{HITW}_G \left(b + T, L + \sum_{t=1}^T (\mathbf{y}_t - X_t' \beta)' (\mathbf{y}_t - X_t' \beta) \right).$$

Update the graph G

We apply a MCMC for multivariate graphical models for G (see [Jones et al. \(2005\)](#)) and due to prior independence assumption,

$$p(\mathbf{y} | G) = \int \prod_{t=1}^T (2\pi)^{-n/2} |\Sigma|^{-n/2} \exp \left(-\frac{1}{2} (\mathbf{y}_t - X_t' \beta)' \Sigma^{-1} (\mathbf{y}_t - X_t' \beta) \right) p(\beta) p(\Sigma | G) d\beta d\Sigma.$$

Following [Jones et al. \(2005\)](#) we apply a local-move MH based on the conditional posterior $p(G | \dots)$ with an add/delete edge move proposal.

Update the allocation variables D and $\mathcal{E}$

Sampling from the full conditional of D is obtained from

$$P(d_{ij} = d, \xi_{ij} = 1 | \dots) \propto \frac{(1 - \pi_i) \mathcal{N}(\beta_{ij} | \mu_{id}, \lambda_{ij}) \mathcal{G}a(\lambda_{ij} | \gamma_{id}, \tau_{id}/2)}{\sum_{k \in A_{w_i}(u_{ij})} \mathcal{N}(\beta_{ij} | \mu_{ik}, \lambda_{ij}) \mathcal{G}a(\lambda_{ij} | \gamma_{ik}, \tau_{ik}/2)}$$

in the non-sparse case for every $d \in A_{w_i}(u_{ij})$. While in the sparse case, from

$$P(d_{ij} = d, \xi_{ij} = 0 | \dots) \propto \pi_i \mathbb{I}(u_{ij} < \tilde{w}_{id}), \quad d \in A_{\tilde{w}}(u_{ij})$$

where $A_{\tilde{w}}(u_{ij}) = \{k : u_{ij} < \tilde{w}_k\} = \{0\}$, because $\tilde{w}_k = 0, \forall k > 0$,

$$P(d_{ij} = d, \xi_{ij} = 0 | \dots) \propto \begin{cases} \pi_i \mathbb{I}(u_{ij} < 1) \mathcal{N}(\beta_{ij} | 0, \lambda_{ij}) \mathcal{G}a(\lambda_{ij} | \gamma_0, \tau_0/2) & \text{if } d = 0, \\ 0 & \text{if } d > 0. \end{cases}$$

Update the prior restriction probabilities π

The full conditional for π_i is,

$$f(\pi_i | \dots) \propto \mathcal{Be} \left(n_i + 1 - \sum_{i=1}^{n_i} \mathbb{I}(\xi_{ii} = 1), \alpha_i + \sum_{i=1}^{n_i} \mathbb{I}(\xi_{ii} = 1) \right).$$

Appendix D. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jeconom.2019.04.022>.

References

- Acemoglu, D., Carvalho, V.M., Ozdaglar, A., Tahbaz-Salehi, A., 2012. The network origins of aggregate fluctuations. *Econometrica* 80 (5), 1977–2016.
- Acemoglu, D., Ozdaglar, A., Tahbaz-Salehi, A., 2015. Systemic risk and stability in financial networks. *Am. Econom. Rev.* 105 (2), 564–608.
- Ahelgebey, D.F., Billio, M., Casarin, R., 2016a. Bayesian graphical models for structural vector autoregressive processes. *J. Appl. Econometrics* 31 (2), 357–386.
- Ahelgebey, D.F., Billio, M., Casarin, R., 2016b. Sparse graphical vector autoregression: A bayesian approach. *Ann. Econom. Statist.* 123, 333–361.
- Banbura, M., Giannone, D., Reichlin, L., 2010. Large Bayesian vector autoregressions. *J. Appl. Econometrics* 25 (1), 71–92.
- Barabási, A.-L., Albert, R., 1999. Emergence of scaling in random networks. *Science* 286 (5439), 509–512.
- Bargigli, L., Gallegati, M., 2013. Finding communities in credit networks. *Economics* 7 (17), 1.
- Barigozzi, M., Brownlees, C., 2018. NETS: Network Estimation for Time Series. Working Paper.
- Bassetti, F., Casarin, R., Leisen, F., 2014. Beta-product dependent Pitman-Yor processes for Bayesian inference. *J. Econometrics* 180 (1), 49–72.
- Battiston, S., Puliga, M., Kaushik, R., Tasca, P., Caldarelli, G., 2012. Debtrank: Too central to fail? Financial networks, the fed and systemic risk. *Sci. Rep.* 2.
- Bianchi, D., Billio, M., Casarin, R., Guidolin, M., 2018. Modeling systemic risk with Markov switching graphical SUR models. *J. Econometrics* 210 (1), 58–74.
- Billio, M., Getmansky, M., Lo, A.W., Pelizzon, L., 2012. Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *J. Financ. Econ.* 104 (3), 535–559.
- Brownlees, C., Engle, R., 2016. SRISK: A conditional capital shortfall measure of systemic risk. *Rev. Financ. Stud.* 30 (1), 48–79.
- Canova, F., Ciccarelli, M., 2004. Forecasting and turning point prediction in a Bayesian panel VAR model. *J. Econometrics* 120 (2), 327–359.
- Caron, F., Fox, E.B., 2017. Sparse graphs using exchangeable random measures. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 79 (5), 1295–1366.
- Carriero, A., Clark, T., Marcellino, M., 2015. Bayesian VARs: Specification choices and forecast accuracy. *J. Appl. Econometrics* 30 (1), 46–73.
- Carvalho, C.M., Massam, H., West, M., 2007. Simulation of hyper-inverse wishart distributions in graphical models. *Biometrika* 94 (3), 647–659.
- De Souza, S.R.S., Silva, T.C., Tabak, B.M., Guerra, S.M., 2016. Evaluating systemic risk using bank default probabilities in financial networks. *J. Econom. Dynam. Control* 66, 54–75.
- Demirer, M., Diebold, F.X., Liu, L., Yilmaz, K., 2018. Estimating global bank network connectedness. *J. Appl. Econometrics* 33 (1), 1–15.
- Di Lucca, M., Guglielmi, A., Muller, P., Quintana, F., 2013. A simple class of Bayesian nonparametric autoregression models. *Bayesian Anal.* 8 (1), 63–88.
- Diebold, F.X., Yilmaz, K., 2014. On the network topology of variance decompositions: Measuring the connectedness of financial firms. *J. Econometrics* 182 (1), 119–134.
- Diebold, F.X., Yilmaz, K., 2015. Measuring the Dynamics of Global Business Cycle Connectedness. Oxford University Press, pp. 45–89.
- Diebold, F.X., Yilmaz, K., 2016. Trans-Atlantic equity volatility connectedness: U.S. and European financial institutions, 2004–2014. *J. Financ. Econ.* 141 (1), 81–127.
- Doan, T., Litterman, R., Sims, C.A., 1984. Forecasting and conditional projection using realistic prior distributions. *Econometric Rev.* 3 (1), 1–100.
- Eichler, M., 2013. CaUsal inference with multiple time series: principles and problems. *Phil. Trans. R. Soc. A* 371 (1997).
- Ferguson, T.S., 1973. A Bayesian analysis of some nonparametric problems. *Ann. Statist.* 1 (2), 209–230.
- Francis, N., Owyang, M., Savascin, O., 2017. An endogenously clustered factor approach to international business cycles. *J. Appl. Econometrics* 32 (7), 1261–1276.
- Gefang, D., 2014. Bayesian doubly adaptive elastic-net Lasso for VAR shrinkage. *Int. J. Forecast.* 30 (1), 1–11.
- George, E.I., Sun, D., Ni, S., 2008. Bayesian stochastic search for VAR model restrictions. *J. Econometrics* 142 (1), 553–580.
- Griffin, J., Kalli, M., 2018. Bayesian nonparametric vector autoregressive models. *J. Econometrics* 203 (2), 267–282.
- Griffin, J., Steel, M.F.J., 2006. Inference with non-gaussian ornstein–uhlenbeck processes for stochastic volatility. *J. Econometrics* 134 (2), 605–644.
- Griffin, J.E., Steel, M.F.J., 2011. Stick-breaking autoregressive processes. *J. Econometrics* 162 (2), 383–396.
- Hatjispyros, S.J., Nicolieris, T.N., Walker, S.G., 2011. Dependent mixtures of Dirichlet processes. *Comput. Statist. Data Anal.* 55 (6), 2011–2025.
- Hautsch, N., Schaumburg, J., Schienle, M., 2015. Financial network systemic risk contributions. *Rev. Financ.* 19 (2), 685–738.
- Hirano, K., 2002. Semiparametric Bayesian inference in autoregressive panel data models. *Econometrica* 70 (2), 781–799.
- Hjort, N.L., Homes, C., Müller, P., Walker, S.G., 2010. Bayesian Nonparametrics. Cambridge University Press.
- Huber, F., Feldkircher, M., 2017. Adaptive shrinkage in Bayesian vector autoregressive models. *J. Bus. Econom. Statist.* 1–13.
- Jackson, M., 2008. Social and Economic Networks. Princeton University Press.
- Jensen, J.M., Maheu, M.J., 2010. Bayesian semiparametric stochastic volatility modeling. *J. Econometrics* 157 (2), 306–316.
- Jones, B., Carvalho, C., Dobra, A., Hans, C., Carter, C., West, M., 2005. Experiments in stochastic computation for high-dimensional graphical models. *Statist. Sci.* 20 (4), 388–400.
- Kalli, M., Griffin, J.E., Walker, S.G., 2011. Slice sampling mixture models. *Stat. Comput.* 21 (1), 93–105.
- Karlsson, S., 2013. Forecasting with Bayesian Vector Autoregression. In: Elliott, G., Granger, C., Timmermann, A. (Eds.), *HandBook of Economic Forecasting*, Vol. 2, Chapter 15. Elsevier, pp. 791–897.
- Kaufmann, S., 2010. Dating and forecasting turning points by Bayesian clustering with dynamic structure: a suggestion with an application to Austrian data. *J. Appl. Econometrics* 25 (2), 309–344.
- Kaufmann, S., Schumacher, C., 2017. Identifying relevant and irrelevant variables in sparse factor models. *J. Appl. Econometrics* 32 (6), 1123–1144.
- Kivela, M., Arenas, A., Barthélemy, M., Gleeson, J.P., Moreno, Y., Porter, M.A., 2014. Multilayer networks. *J. Complex Netw.* 2 (3), 203.
- Koop, G., 2013. Forecasting with medium and large Bayesian VARs. *J. Appl. Econometrics* 28 (2), 177–203.
- Koop, G., Korobilis, D., 2016. Model uncertainty in panel vector autoregressions. *Eur. Econom. Rev.* 81, 115–131.
- Korobilis, D., 2013. VAR forecasting using Bayesian variable selection. *J. Appl. Econometrics* 28 (2), 204–230.
- Korobilis, D., 2016. Prior selection for panel vector autoregressions. *Comput. Statist. Data Anal.* 101, 110–120.
- Korobilis, D., Pettenuzzo, D., 2019. Adaptive hierarchical priors for high-dimensional vector autoregressions. *J. Econometrics*.
- Kose, M.A., Otrok, C., Whiteman, C.H., 2003. International business cycles: World, region and country specific factors. *Amer. Econ. Rev.* 93 (4), 1216–1239.
- Leicht, E.A., Newman, M.E., 2008. Community structure in directed networks. *Phys. Rev. Lett.* 100 (11), 118703.
- Litterman, R., 1986. Forecasting with bayesian vector autoregressions-five years of experience. *J. Bus. Econom. Statist.* 4 (1), 25–38.
- Lo, A.Y., 1984. On a class of bayesian nonparametric estimates: I. density estimates. *Ann. Statist.* 12 (1), 351–357.
- MacEachern, S.N., 1999. Dependent nonparametric processes. In: In in ASA Proceedings of the Section on Bayesian Statistical Science, Alexandria, VA. American Statistical Association, pp. 50–55.
- MacEachern, S.N., 2001. Decision theoretic aspects of dependent nonparametric processes. In: George, E. (Ed.), *Bayesian Methods with Applications To Science, Policy and Official Statistics*. Creta: ISBA, pp. 551–560.
- MacLehose, R., Dunson, D., 2010. Bayesian semiparametric multiple shrinkage. *Biometrics* 66 (2), 455–462.
- Magnus, J.R., Neudecker, H., 1999. Matrix Differential Calculus with Applications in Statistics and Econometrics, second ed. John Wiley.

- McCracken, M.W., Ng, S., 2016. Fred-md: A monthly database for macroeconomic research. *J. Bus. Econom. Statist.* 34 (4), 574–589.
- Miller, R., 1980. Bayesian analysis of the two-parameter gamma distribution. *Technometrics* 22 (1), 65–69.
- Müller, P., Quintana, F., Rosner, G., 2004. A method for combining inference across related nonparametric Bayesian models. *J. R. Statist. Soc. B* 66, 735–749.
- Newman, M.E., 2006. Modularity and community structure in networks. *Proc. Natl. Acad. Sci.* 103 (23), 8577–8582.
- Newman, M., 2010. *Networks: An Introduction*. Oxford University Press.
- Nieto-Barajas, L.E., Quintana, F.A., 2016. A Bayesian non-parametric dynamic AR model for multiple time series analysis. *J. Time Series Anal.* 37 (5), 675–689.
- Opsahl, T., Agneessens, F., Skvoretz, J., 2010. Node centrality in weighted networks: Generalizing degree and shortest paths. *Social Networks* 32 (3), 245–251.
- Park, T., Casella, G., 2008. The Bayesian Lasso. *J. Amer. Statist. Assoc.* 103 (482), 681–686.
- Pennell, M.L., Dunson, D.B., 2006. Bayesian semiparametric dynamic frailty models for multiple event time data. *Biometrics* 62, 1044–1052.
- Pitman, J., Yor, M., 1997. The two parameter Poisson-Dirichlet distribution derived from a stable subordinator. *Ann. Prob.* 25, 855–900.
- Puliga, M., Flori, A., Pappalardo, G., Chessa, A., Pammolli, F., 2016. The accounting network: How financial institutions react to systemic crisis. *PLoS One* 11 (10), e0162855.
- Richardson, S., Green, P.J., 1997. On Bayesian analysis of mixtures with an unknown number of components. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 59, 731–792.
- Scott, H., 2016. *Connectedness and Contagion: Protecting the Financial System from Panics*. MIT Press.
- Scott, S.L., Varian, H.R., 2014. Predicting the present with bayesian structural time series. *Int. J. Math. Modell. Numer. Opt.* 5 (1–2), 4–23.
- Sethuraman, J., 1994. A constructive definition of the Dirichlet process prior. *Statist. Sinica* 4, 639–650.
- Sims, C.A., Zha, T., 1998. Bayesian methods for dynamic multivariate models. *Internat. Econom. Rev.* 39 (4), 949–968.
- Stock, J.H., Watson, M.W., 2012. Generalized shrinkage methods for forecasting using many predictors. *J. Bus. Econom. Statist.* 30 (4), 481–493.
- Taddy, M.A., Kottas, A., 2009. Markov switching Dirichlet process mixture regression. *Bayesian Anal.* 4 (4), 793–816.
- Walker, S.G., 2007. Sampling the Dirichlet mixture model with slices. *Comm. Statist. Simulation Comput.* 36 (1), 45–54.
- Wang, H., 2010. Sparse seemingly unrelated regression modelling: Applications in finance and econometrics. *Comput. Statist. Data Anal.* 54 (11), 2866–2877.
- Zou, H., Hastie, T., 2005. Regularization and variable selection via the elastic net. *J. R. Statist. Soc. B* 67 (2), 301–320.