462354,0, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/icre.12577 by HS Education for Scotland NES, Edinburgh Central Office, Wiley Online Library on [171/0202], See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable, Centive Commons License

APPROXIMATE BAYESIAN INFERENCE AND FORECASTING IN HUGE-DIMENSIONAL MULTICOUNTRY VARS*

BY MARTIN FELDKIRCHER, FLORIAN HUBER, GARY KOOP, AND MICHAEL PFARRHOFER

Vienna School of International Studies, Austria; University of Salzburg, Austria; University of Strathclyde, U.K.; University of Salzburg, Austria

Panel vector autoregressions (PVARs) are a popular tool for analyzing multicountry data sets. However, the number of estimated parameters can be enormous, leading to computational and statistical issues. In this article, we develop fast Bayesian methods for estimating PVARs using integrated rotated Gaussian approximations. We exploit the fact that domestic information is often more important than international information and group the coefficients accordingly. Fast approximations are used to estimate the latter whereas the former are estimated with precision using Markov chain Monte Carlo techniques. We illustrate, using a huge model of the world economy, that it produces competitive forecasts quickly.

1. INTRODUCTION

There is much evidence that working with multicountry time-series models improves macroeconomic forecasting and structural analysis (see, among many others, Pesaran et al., 2004, 2009; Canova and Ciccarelli, 2016). This is to be expected in the modern globalized economy where countries are linked together through trade and financial flows and events in one country can spill over into others. However, the relevant data sets can be enormous. In this article, we work with a 38-country data set that contains 487 variables. If all of these are treated as endogenous variables in an unrestricted multicountry vector autoregression (VAR), the number of equations in the VAR will be huge, as will the number of right-hand side variables in each equation. The resulting model will be overparameterized.

Much of the existing literature deals with this problem by imposing restrictions on the model or compressing the data (e.g., by using factor methods). For instance, the popular class of global VARs (GVARs, see, e.g., Pesaran et al., 2004, 2009; Cuaresma et al., 2016; Dovern et al., 2016; Huber, 2016) assumes that information from all other countries impacts a country solely through a single weighted average of other country information. The weights in the average typically are based on bilateral trade flows. By contrast, the literature on panel VARs (PVARs) mainly deals with overparameterization issues through constraints on the parameters describing the dynamic and static relations across countries (see, e.g., Canova and Ciccarelli, 2009).

We propose an unrestricted PVAR specification where any variable can affect any or all other variables either contemporaneously or with a lag. This implies that the influence of

*Manuscript received March 2021; revised February 2022.

We thank the editor Jesús Fernández-Villaverde and two anonymous referees for their valuable suggestions and constructive comments. We also thank Maximilian Böck, Todd Clark, Niko Hauzenberger, Ed Knotek, James Mitchell, Anna Stelzer, Saeed Zaman, and participants of the IAAE 2021 Annual Conference in Rotterdam. Florian Huber and Michael Pfarrhofer gratefully acknowledge financial support from the Austrian Science Fund (FWF, Grant no. ZK 35) and the Jubiläumsfonds of the Oesterreichische Nationalbank (Grant no. 18304). Please address correspondence to: Florian Huber, Department of Economics, University of Salzburg, Mönchsberg 2a, 5020 Salzburg, Austria. E-mail: *florian.huber@plus.ac.at.*

^{© 2022} The Authors. *International Economic Review* published by Wiley Periodicals LLC on behalf of the Economics Department of the University of Pennsylvania and the Osaka University Institute of Social and Economic Research Association

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

foreign variables is neither a priori restricted (e.g., using trade weights) nor based on the output of a dimension-reduction procedure such as principal components. In other words, we want to let the data decide the exact nature and extent of linkages between countries. This feature of our approach turns out to be a substantial improvement over multicountry models commonly used in the literature.

Bayesian shrinkage methods are increasingly used for overcoming the overparameterization problems which arise when working with unrestricted PVARs (or large VARs in general). Influential early contributions such as Canova and Ciccarelli (2009) used simple methods for choosing the prior (e.g., subjective elicitation or training sample methods). More recently, Bayesians have been working with global–local shrinkage or variable selection priors (see, e.g., Koop and Korobilis, 2016; Korobilis, 2016; Bai et al., 2022). These priors are commonly used in Big Data problems where models involve a large number of parameters. They automatically sort through all the parameters and decide which ones to shrink to zero and which ones to estimate freely. Our goal in this article is to use a global–local shrinkage prior in a large unrestricted PVAR and let it decide which cross-country linkages are important and which can be ignored. In our empirical work, we mainly use the Horseshoe prior of Carvalho et al. (2010) although the econometric methods developed in this article will work with any hierarchical shrinkage prior that takes a conditionally Gaussian form including the least absolute shrinkage and selection operator (LASSO) and Dirichlet–Laplace priors of Park and Casella (2008) and Bhattacharya et al. (2015), respectively.

Bayesian estimation and forecasting in VARs using global-local shrinkage priors typically requires the use of computationally demanding Markov chain Monte Carlo (MCMC) methods. These algorithms are impractical in the very large PVARs that arise when working with many countries and many variables. For this reason, the existing Bayesian literature which uses unrestricted PVARs has focused on relatively small models. For instance, Koop and Korobilis (2016) use a PVAR involving four variables for each of seven European countries which is much smaller than the one considered in this article.

To overcome the computational hurdle, we develop an integrated rotated Gaussian approximation (IRGA) for the PVAR. IRGAs were recently developed by Van den Boom et al. (2021) as a machine learning tool to speed up computation in high-dimensional models. These methods build on the intuition that some parameters are more important than other parameters. The other parameters, which in our case are mainly coefficients associated with other countries' lagged endogenous variables and covariance terms, are estimated using efficient approximations. The more important parameters are then estimated conditional on these approximations using precise MCMC techniques. We adapt these methods for use with PVARs. The resulting IRGA-based algorithm leads to vast improvements in speed of computation and has appealing approximation properties which we illustrate through simulations.

In our empirical work, we estimate a huge model of the world economy that contains 487 endogenous variables. In a forecasting exercise, we compare the performance of our unrestricted PVAR with a Horseshoe prior, estimated using IRGA methods (PVAR-IRGA), to a range of alternatives including a GVAR, factor-augmented VARs (FAVARs) with the factors constructed using other country variables, and single-country Bayesian VARs (BVARs). In terms of computation, our key finding is that the computational improvements are large enough to enable Bayesian forecasting and structural analysis to be done even in huge PVARs. In terms of empirical results, we find our large approximate model forecasts well and often outperforms competing models. The forecast improvements are particularly strong for short-run density forecasts of stock market returns and longer-run inflation and output predictions. In these cases, we find PVAR-IRGA, with few exceptions, to forecast substantially better than all the alternatives.

We then proceed by analyzing the properties of the forecasts of the PVAR-IRGA via recent techniques used for analyzing social networks (Holland et al., 1983; Karrer and Newman, 2011; Žiberna, 2014; Pati and Bhattacharya, 2015). These stochastic block models build on the correlation matrix and sort correlations into clusters which facilitate interpretation. This analysis provides novel insights on the properties of the forecasts which are consistent with actual developments during the global financial crisis and the euro area sovereign debt crisis.

Our forecasting exercise is complemented by additional empirical results on the degree of cross-country spillovers. Considering different variants of the Diebold and Yilmaz (2009) spillover index shows that our model detects sizable international relations across output, prices, long-term interest rates (LT-IR), and stock markets which sharply increase throughout the holdout period. These increases are particularly pronounced during times of economic turmoil. Our findings hence provide novel insights on truly global connectivity since comparable large-scale analyses are not feasible with standard econometric techniques.

The remainder of the article is structured as follows: In Section 2, we define the PVAR likelihood and the prior we use to carry out Bayesian inference and prediction. Section 3 discusses computation and develops our IRGA methods which allow for fast computation. It also includes a theoretical discussion on the approximation properties of IRGA. In Section 4, we carry out a simulation exercise to complement the theoretical discussion whereas in Section 5 we present results for our forecasting exercise which compares our PVAR-IRGA to a range of other approaches. This section also includes information on the extent of crosscountry spillovers. The final section summarizes and concludes the article. The Appendix provides additional technical details and further empirical results such as robustness checks.

2. THE PVAR

This section develops the basic PVAR and briefly discusses the main specification issues commonly faced by researchers. We then discuss how Bayesian techniques can be used to deal with overparameterization concerns and thus, in an automatic fashion, solve several of these specification issues.

2.1. The Likelihood Function. Our goal is to model dynamic and static relations in an international panel of macroeconomic and financial time series which are stored in an *n*-dimensional vector $\mathbf{y}_t = (\mathbf{y}'_{1t}, \dots, \mathbf{y}'_{Nt})'$ for $t = 1, \dots, T$. This vector is composed of N country-specific subvectors \mathbf{y}_{it} which are $M \times 1$ -dimensional.¹ Assuming that each \mathbf{y}_{it} depends on the lagged values of \mathbf{y}_t , we obtain a PVAR given by

(1)
$$\mathbf{y}_{it} = \mathbf{\Gamma}_{i1}\mathbf{y}_{it-1} + \dots + \mathbf{\Gamma}_{ip}\mathbf{y}_{it-p} + \mathbf{\Xi}_{i}\mathbf{z}_{it} + \boldsymbol{\epsilon}_{it},$$

where Γ_{ij} are $M \times M$ coefficient matrices associated with the lagged endogenous variables of country *i*. Lags of variables from countries other than *i* are denoted by $z_{it} = (\mathbf{y}'_{-i,t-1}, \ldots, \mathbf{y}'_{-i,t-p})'$ with $\mathbf{y}_{-i,t} = (\mathbf{y}'_{1t}, \ldots, \mathbf{y}'_{i+1,t}, \ldots, \mathbf{y}'_{Nt})'$. The coefficient matrix on other country lags, Ξ_i , is an $M \times K_{other}$ matrix where $K_{other} = (N-1)Mp$. Note that Ξ_i will contain an enormous number of parameters unless N and/or M are small. The matrix Ξ_i encodes the dynamic relations across countries (which are commonly referred to as dynamic interdependencies in the literature) whereas the $M \times k(=Mp)$ matrix $\Gamma_i = (\Gamma_{i1}, \ldots, \Gamma_{ip})$ captures domestic dynamics.

The usual VAR representation in terms of y_t is obtained by stacking all country-specific models and reshuffling the columns of Γ_i and Ξ_i appropriately:

(2)
$$\mathbf{y}_t = \tilde{\mathbf{\Gamma}}_1 \mathbf{y}_{t-1} + \dots + \tilde{\mathbf{\Gamma}}_p \mathbf{y}_{t-p} + \boldsymbol{\epsilon}_t.$$

The coefficient matrices $\tilde{\Gamma}_j$ are of dimension $n \times n$ and the errors $\epsilon_t = (\epsilon'_{1t}, \dots, \epsilon'_{Nt})'$ are i.i.d. Gaussian with Σ being an $n \times n$ -dimensional variance–covariance matrix. The off-diagonal elements of this matrix determine both contemporaneous relations across variables *within* a

¹ Note that M may differ across countries but is used here to simplify notation. Our approach naturally allows for different covariates across equations and countries. In the empirical work, variable coverage differs across countries.

country and instantaneous dependencies *across* countries. The latter relations are typically referred to as static interdependencies in the literature. Notice that unrestricted estimation of all contemporaneous relations across countries implies estimating n(n-1)/2 covariances. For large *n*, this adds to the already huge number of parameters in the Γ_i 's and Ξ_i 's.

In the literature on PVARs, estimation is often facilitated by introducing restrictions (see, e.g., Canova and Ciccarelli, 2009, 2016) on the coefficients in (1) and (2). For instance, the so-called cross-sectional homogeneity restriction arises if $\Gamma_i = \Gamma_s$ for $i \neq s$. This implies that domestic dynamics across countries are identical—a rather restrictive assumption if the panel of countries includes, for example, developed and developing economies. Another restriction often introduced is $\Xi_i = \mathbf{0}$ for some (or even all) *i*. This rules out dynamic relations across some countries but substantially reduces the number of free parameters. GVARs are also not restriction free. The assumption here is that cross-country linkages can be approximated by cross-country weighted averages and hence restrict Ξ_i .² Another restriction sometimes considered assumes that shocks across countries are uncorrelated, that is, $\operatorname{cov}(\epsilon_{it}, \epsilon_{st}) = \mathbf{0}$. This implies introducing zero restrictions on the relevant elements in Σ . All these restrictions, however, have serious implications for forecasting and structural inference and potentially introduce misspecification if chosen wrongly. These considerations inspire us to use Bayesian variable selection methods via a global–local shrinkage prior so as to choose the appropriate restrictions in a data-based manner.

If left unrestricted, estimation of the PVAR using traditional Bayesian MCMC methods is computationally cumbersome. For large data sets such as the one used in this article, the computational burden becomes impractical. Hence, a goal of this article is to speed up computation. In a first step, we greatly simplify computation by transforming the PVAR to allow for equation-by-equation estimation. This can be achieved by taking a Cholesky-type decomposition of $\Sigma = UHU'$. Here, we let U denote a lower triangular matrix with unit diagonal and H is a diagonal matrix with main diagonal $\sigma^2 = (\sigma_{1}^{2'}, \ldots, \sigma_{N}^{2'})'$. The *M*-dimensional vector $\sigma_i^2 = (\sigma_{\varepsilon,i1}^2, \ldots, \sigma_{\varepsilon,iM}^2)'$ stores the idiosyncratic variances $\sigma_{\varepsilon,ij}^2$ associated with the shock in country *i* and equation *j*. We can use this decomposition to recover the structural form of (2):

$$\mathbf{y}_t = \mathbf{A}_1 \mathbf{y}_{t-1} + \cdots + \mathbf{A}_p \mathbf{y}_{t-p} + \mathbf{W} \mathbf{y}_t + \mathbf{\varepsilon}_t, \quad \mathbf{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{H}),$$

where $W = (I - U^{-1})$ encodes the contemporaneous relations across the shocks in the system and the matrices A_j (j = 1, ..., p) denote structural coefficients. Within a given country, we can easily obtain a representation similar to (1) by reshuffling the explanatory variables:

$$\mathbf{y}_{it} = \mathbf{A}_{i1}\mathbf{y}_{it-1} + \cdots + \mathbf{A}_{ip}\mathbf{y}_{it-p} + \mathbf{B}_i\mathbf{z}_{it} + \mathbf{W}_i\mathbf{y}_t + \boldsymbol{\varepsilon}_{it},$$

with W_i denoting the *M* rows of *W* associated with the *i*th country. This allows us to rewrite the *j*th (>1) equation in the country-specific model *i*(> 1) as follows:³

(3)
$$y_{ij,t} = \mathbf{A}'_{ij,\bullet} \mathbf{x}_{it} + \mathbf{B}'_{ij,\bullet} \mathbf{z}_{it} + \sum_{s=1}^{j-1} w_{is} y_{is,t} + \sum_{v=1}^{i-1} \mathbf{u}'_{iv} \mathbf{y}_{vt} + \varepsilon_{ij,t},$$

where $y_{ij,t}$ denotes the *j*th element of $y_{it}, x_{it} = (y'_{it-1}, \dots, y'_{it-p})'$ whereas $A_{ij,\bullet}$ and $B_{ij,\bullet}$ denote the *j*th rows of A_i and B_i , respectively. $u_{ij,\bullet} = (w_{i1}, \dots, w_{i,j-1}, u'_{i1}, \dots, u'_{ii-1})'$ are the covariance parameters associated with the relevant row in W_i .

Note that the errors are now independent across equations (i.e., both within and across countries) with error variance given by $\sigma_{\varepsilon,ii}^2$. This independence allows for estimating one

² These weights have to be specified exogenously and are often based on measures of economic connectivity such as bilateral trade flows. For an overview, see Feldkircher and Huber (2016).

³ For j = 1 and i = 1, the equation simplifies and only x_{it} and z_{it} appear as regressors.

equation at a time which greatly speeds up computation. Cholesky-based formulations such as this have been used in many recent papers, including Carriero et al. (2019), Koop et al. (2019), Huber et al. (2021), and Carriero et al. (2021). As opposed to Carriero et al. (2021), our approach includes the contemporaneous values of the endogenous variables and is thus not order-invariant. In the Appendix, we show that the results are robust to different orderings of the countries in y_t , implying only negligible empirical differences.

Equation (3) is a simple regression model which regresses $y_{ij,t}$ on the lags of y_{it} , the lags of the other countries' endogenous variables in z_{it} , the contemporaneous values of the preceding j-1 variables domestic variables in country i as well as the contemporaneous values of the preceding i-1 countries. Our approach builds on the notion that x_{it} is more important in explaining $y_{ij,t}$ than all other quantities in (3).

To simplify notation, let $\tilde{z}_{ij,t} = (z'_{it}, y_{i1,t}, \dots, y_{ij-1,t}, y'_{1t}, \dots, y'_{i-1t})'$ denote a $K_{ij}(=K_{other} + j - 1 + (i - 1)M)$ vector which stores the international quantities (both lagged and contemporaneously) as well as the time t values of the endogenous variables up to equation i. Moreover, we stack the corresponding regression coefficients in a K_{ij} vector $\tilde{B}_{ij,\bullet} = (B'_{ij,\bullet}, u'_{ij,\bullet})'$. Notice that K_{ij} is much larger than k which implies that $\tilde{B}_{ij,\bullet}$ is difficult to estimate for large M, N, and p. We can rewrite (3) in full-data form by stacking the T observations into vectors to obtain the PVAR equations which define the likelihood function of our model:

(4)
$$\mathbf{y}_{ij} = \mathbf{x}_i \mathbf{A}_{ij,\bullet} + \tilde{\mathbf{z}}_{ij} \tilde{\mathbf{B}}_{ij,\bullet} + \boldsymbol{\varepsilon}_{ij} \Leftrightarrow \mathbf{y}_{ij} \sim \mathcal{N} \Big(\mathbf{x}_i \mathbf{A}_{ij,\bullet} + \tilde{\mathbf{z}}_{ij} \tilde{\mathbf{B}}_{ij,\bullet}, \sigma_{\varepsilon,ij}^2 \mathbf{I}_T \Big).$$

with x_i being a $T \times k$ matrix where k = Mp with tth row given by x'_{it} . \tilde{z}_{ij} is a $T \times K_{ij}$ matrix with tth row \tilde{z}'_{ijt} . This equation is a regression model which discriminates between a highdimensional set of predictors related to covariances and international quantities in \tilde{z}_{ij} and a low-dimensional set of domestic quantities in x_i .

2.2. *The Prior.* The methods developed in this article apply for any prior which has a hierarchical Gaussian form and thus leads to a full conditional posterior which is Gaussian. This is due to the fact that IRGA methods exploit the property that Gaussian distributions are invariant to rotations. A popular class of priors which has this form is the class of Gaussian global–local shrinkage priors. These can be represented as scale mixtures of Gaussians.⁴

At a general level, consider the *j*th coefficient in a model, ϕ_j . A global-local shrinkage prior can be written as:

$$\phi_j \sim \mathcal{N} \Big(0, \psi_j^2 \lambda^2 \Big), \quad \psi_j \sim f, \quad \lambda \sim g,$$

where f and g are mixing densities and many different choices for them have been proposed. In a global–local shrinkage prior, λ controls global shrinkage (common to all parameters). Having global shrinkage has often been found useful in BVARs (e.g., the Minnesota prior has a global shrinkage parameter) to reduce overfitting concerns.⁵ ψ_j does local shrinkage (specific to the *j*th parameter). That is, if ψ_j is estimated to be close to zero then ϕ_j is shrunk to be close to zero.

Suitably chosen mixing densities f and g result in a wide range of popular shrinkage priors such as the LASSO (Park and Casella, 2008), Normal-Gamma (Griffin and Brown, 2010), Dirichlet–Laplace (Bhattacharya et al., 2015), or the Horseshoe (Carvalho et al., 2010). Due to its empirical success and ease of implementation, we use the Horseshoe prior. It takes the form:

$$\lambda \sim \mathcal{C}^+(0,1), \quad \psi_i \sim \mathcal{C}^+(0,1),$$

⁴ Cadonna et al. (2020) provides a taxonomy of a range of priors in this class and discusses their properties.

⁵ In Appendix A.2, we show how to implement a hierarchical Minnesota-type prior within this framework.

whereby C^+ denotes the half-Cauchy distribution. Makalic and Schmidt (2016) show that the Horseshoe can be equivalently stated in terms of inverse Gamma distributions using suitable auxiliary variables. Specifically,

$$\psi_i^2 | v_i \sim \mathcal{G}^{-1}(1/2, 1/v_i), \quad \lambda^2 | \xi \sim \mathcal{G}^{-1}(1/2, 1/\xi), \quad v_i, \xi \sim \mathcal{G}^{-1}(1/2, 1),$$

with \mathcal{G}^{-1} denoting the inverse Gamma distribution and ν_j , ξ being auxiliary parameters. These auxiliary parameters are merely used to simplify posterior inference.

Up to this point, we have assumed that all parameters of the model are forced to zero through a single global shrinkage parameter λ . However, in the PVAR model we have different sets of parameters, countries, and variable types (equations within a given country). Specifying a single shrinkage parameter would imply that global shrinkage is symmetric across these different dimensions, a rather restrictive assumption. In the PVAR, we will assume that the λ does not shrink *all* parameters toward zero but is specified to differ across countries, equations, and types of parameters. This implies that for each coefficient vector $A_{ij,\bullet}$ and $\tilde{B}_{ij,\bullet}$ we will replace λ with $\lambda_{A,ij}$ and $\lambda_{B,ij}$, respectively. In other words, each equation will have its own global shrinkage parameters and there will be two of them: one for own country coefficients and one for other country and contemporaneous coefficients.

Since $\tilde{B}_{ij,\bullet}$ includes the covariance parameters as well, our discussion suggests that we use a Horseshoe also on the off-diagonal elements of U. This implies that our prior allows for detecting whether static interdependencies across countries and variables are present and, if not, introduces shrinkage. We choose independent weakly informative inverse Gamma prior on the variances collected on the main diagonal of H. In particular, $\sigma_{\epsilon,ij}^2 \sim \mathcal{G}^{-1}(a_{\sigma}, b_{\sigma})$ with $a_{\sigma} = b_{\sigma} = 0.01$ for all countries and equations.

3. APPROXIMATE BAYESIAN INFERENCE IN PVARS

In this section, we provide a framework that is capable of estimating huge PVAR models at reasonable computational cost. The next subsection introduces IRGA to the PVAR case and provides some information on how posterior simulation can be carried out. Since this approach relies on approximating certain regions of the parameter space we then discuss the theoretical properties of the approximation.

3.1. Posterior Computation Using IRGAs. As written in (4), the PVAR simply involves MN regression models. Bayesian MCMC methods for posterior and predictive inference in the regression model using a Horseshoe prior are well-established (see Makalic and Schmidt, 2016). In theory, we could simply use such methods with our PVAR. However, the problem is that MCMC methods are simply too slow for dealing with the high-dimensional parameter spaces that arise with PVARs. The main source of this high-dimensionality is that $\tilde{B}_{ij,\bullet}$ (the coefficients on contemporaneous and other country variables in the equation for the *j*th variable for country *i*) potentially contains tens of thousands of coefficients. The matrix of own country coefficients, $A_{ij,\bullet}$, is much smaller.

One empirical regularity often found in multicountry data sets is that own country effects are usually more important than other country effects. Hence, the literature sometimes sets B_i equal to a zero matrix to rule out dynamic interdependencies. This, however, could translate into a misspecified model. In this article, we allow the data to speak about the degree of sparsity in $\tilde{B}_{ij,\bullet}$. A second empirical regularity is that the time series in \tilde{z}_{ij} often display substantial components (Kose et al., 2003). A potential solution would be to extract a low number of principal components from \tilde{z}_{ij} . Another solution, which we adopt here, builds on the notion that if elements in \tilde{z}_{ij} are very similar to each other, it might pay off to not include all of them and effectively control for collinearity. This is also consistent with $\tilde{B}_{ij,\bullet}$ being very sparse. By contrast, $A_{ij,\bullet}$ is likely nonsparse. This consideration motivates the way we implement IRGA with the PVAR. The general idea of IRGA is to use MCMC methods on important (low-dimensional) parameters but compute a fast approximation of the posterior for other (high-dimensional) parameters of less importance. In the PVAR, we consider $A_{ij,\bullet}$ as the important parameters and $\tilde{B}_{ij,\bullet}$ the less important ones.⁶

We now provide details of how we implement IRGA in the PVAR. Let Q_i be the $T \times T$ rotation matrix obtained from the QR-decomposition of x_i and partition it as $Q_i = (Q_{i1}, Q_{i2})$ with Q_{i1} being $T \times k$ and Q_{i2} being $T \times (T - k)$.

Multiplying the equation for the *j*th variable in country *i* by Q_i and exploiting the rotation invariance of the Gaussian distribution yields an equivalent representation of (4):

(5)
$$\boldsymbol{Q}_{i1}^{\prime}\boldsymbol{y}_{ij} \sim \mathcal{N}(\boldsymbol{Q}_{i1}^{\prime}\boldsymbol{x}_{i}\boldsymbol{A}_{ij,\bullet} + \boldsymbol{Q}_{i1}^{\prime}\tilde{\boldsymbol{z}}_{ij}\tilde{\boldsymbol{B}}_{ij,\bullet}, \sigma_{\varepsilon,ij}^{2}\boldsymbol{I}_{k}),$$

(6)
$$\boldsymbol{Q}'_{i2}\boldsymbol{y}_{ij} \sim \mathcal{N}(\boldsymbol{Q}'_{i2}\tilde{\boldsymbol{z}}_{ij}\tilde{\boldsymbol{B}}_{ij,\bullet}, \sigma^2_{\varepsilon,ij}\boldsymbol{I}_{T-k}).$$

The second equation follows since $Q'_{i2}x_i = 0$. Note that $A_{ij,\bullet}$ does not appear in it. This gives rise to a computational strategy which estimates $\tilde{B}_{ij,\bullet}$ independently of $A_{ij,\bullet}$. IRGA involves calculating the posteriors based on the two likelihood functions defined by (5) and (6). An approximate posterior for the (high-dimensional) $\tilde{B}_{ij,\bullet}$ and $\sigma_{\varepsilon,ij}^2$ is obtained using (6). Conditional on this approximate posterior $\hat{p}(\tilde{B}_{ij,\bullet}|Q'_{i2}y_{ij})$, the posterior of $A_{ij,\bullet}$ is obtained using MCMC based on (5).

Any Gaussian approximation can be used for $\hat{p}(\tilde{B}_{ij,\bullet}|Q'_{i2}y_{ij})$. We use vector approximate message passing (VAMP). Our choice of VAMP is driven by its scalability in huge dimensions and the fact that recent papers in machine learning and econometrics (see, e.g., Korobilis, 2021) have shown that it works extremely well for forecasting purposes. The specific implementation of the VAMP algorithm is the one proposed in Rangan et al. (2019) and details are given in Appendix A.⁷ The result is a Gaussian approximation: $\mathcal{N}(\bar{B}_{ij,\bullet}, \bar{V}_{ij,\bullet})$.

Rewriting (5) and plugging in the approximate moments of $\hat{p}(\tilde{B}_{ij,\bullet}|Q'_{i2}y)$ yields:

$$\boldsymbol{Q}_{i1}^{\prime}(\boldsymbol{y}_{ij}-\tilde{\boldsymbol{z}}_{ij}\overline{\boldsymbol{B}}_{ij,\bullet})\sim\mathcal{N}\Big(\boldsymbol{Q}_{i1}^{\prime}\boldsymbol{x}_{i}\boldsymbol{A}_{ij,\bullet},\boldsymbol{Q}_{i1}^{\prime}\tilde{\boldsymbol{z}}_{ij}\overline{\boldsymbol{V}}_{ij,\bullet}\tilde{\boldsymbol{z}}_{ij}^{\prime}\boldsymbol{Q}_{i1}+\sigma_{\varepsilon,ij}^{2}\boldsymbol{I}_{k}\Big).$$

This gives us a Gaussian likelihood for a regression model which can be combined with any (conditionally) Gaussian prior on $A_{ij,\bullet}$ leading to a textbook form for the posterior of $A_{ij,\bullet}$ which can be estimated using MCMC methods. Additional details on the full conditional posterior distributions and how we approximate the error variances as well as the hyperparameters of the prior are given in Appendix A.2.

3.2. Approximation Accuracy. In this subsection, we briefly discuss why a Gaussian approximation is reasonable and how this approximation impacts the posterior distribution of $A_{ij,\bullet}$. Intuitively, the accuracy of the approximate posterior of the domestic coefficients $\hat{p}(A_{ij,\bullet}|\mathbf{y}_{ij})$ depends on the goodness of the approximation to $p(\tilde{B}_{ij,\bullet}|\mathbf{Q}_{j2}\mathbf{y}_{ij})$.

To investigate this relationship more formally, let $KL(p(a)||\hat{p}(a))$ denote the Kullback– Leibler (KL) divergence between an exact and an approximating distribution. Van den Boom et al. (2021) show that the expected (with respect to the conditional distribution of y_{ij} given

⁶ It is worth noting that other choices of important and less important coefficients are possible. For instance, if it is felt likely that the United States is a dominant unit, then coefficients on U.S. variables could always be treated as important even for countries other than the United States. The key consideration is that the number of parameters in $A_{ij,\bullet}$ should not be too large so as to prevent practical use of MCMC methods on it.

⁷ We have also experimented with variational Bayes to approximate $\hat{\rho}(\hat{B}_{ij,e}|Q_{i2}y_{ij})$. However, computation of the expected lower bound on evidence in these dimensions becomes computationally and numerically cumbersome. In addition, in our experiments the VAMP-based algorithm led to more precise predictions of the model.

 $Q'_{i2}y_{ij}$) KL divergence between $p(A_{ij,\bullet}|y_{ij})$ and $\hat{p}(A_{ij,\bullet}|y_{ij})$ is bounded from above by the approximation error to $p(\tilde{B}_{ij,\bullet}|Q'_{i2}y_{ij})$.

Formally, Van den Boom et al. (2021, Theorem 1) establish a link between the approximation quality of $\hat{p}(\tilde{B}_{ij,\bullet}|Q'_{i2}y_{ij})$ and how this impacts the approximate full conditional posterior $\hat{p}(A_{ij,\bullet}|y_{ij})$:

$$\mathbb{E}[\mathrm{KL}(p(\mathbf{A}_{ij,\bullet}|\mathbf{y}_{ij}) || \hat{p}(\mathbf{A}_{ij,\bullet}|\mathbf{y}_{ij}))|\mathbf{Q}'_{i2}\mathbf{y}_{ij}] \leq \\ \mathrm{KL}(p(\mathbf{Q}'_{i1}\tilde{\mathbf{z}}_{ij}\tilde{\mathbf{B}}_{ij,\bullet}|\mathbf{Q}'_{i2}\mathbf{y}_{ij}) \times \mathcal{N}_{\sigma_{ij}} || \hat{p}(\mathbf{Q}'_{i1}\tilde{\mathbf{z}}_{ij}\tilde{\mathbf{B}}_{ij,\bullet}|\mathbf{Q}'_{i2}\mathbf{y}_{ij}) \times \mathcal{N}_{\sigma_{ij}}),$$

with $\mathcal{N}_{\sigma_{ij}} = \mathcal{N}(0, \sigma_{\varepsilon,ij}^2 \mathbf{I}_k)$. This equation has three main implications. First, if the Gaussian approximation $\hat{p}(\tilde{B}_{ij,\bullet}|Q_{i2}\mathbf{y}_{ij})$ is close to the exact full conditional posterior $p(\tilde{B}_{ij,\bullet}|Q_{i2}\mathbf{y}_{ij})$ the corresponding conditional distribution $\hat{p}(A_{ij,\bullet}|\mathbf{y}_{ij})$ will be close to $p(A_{ij,\bullet}|\mathbf{y}_{ij})$. Second, the prior on $A_{ij,\bullet}$ does not impact the error bound. Third, it does not depend on concentration properties around the true value of $\tilde{B}_{ij,\bullet}$ which makes the result relevant for settings with T being small.

In the next step, we justify our Gaussian approximation. Under certain mild assumptions on $p(\tilde{B}_{ij,\bullet}|Q'_{i2}y_{ij})$ and \tilde{z}_{ij} and if $k \ll K_{ij}$, a multivariate central limit theorem implies that $Q'_{i1}\tilde{z}_{ij}\tilde{B}_{ij,\bullet}$ is close to a Gaussian distribution (see Diaconis and Freedman, 1984) even if elements in $\tilde{B}_{ij,\bullet}$ are non-Gaussian. This motivates a Gaussian approximation to $p(\tilde{B}_{ij,\bullet}|Q'_{i2}y_{ij})$. Van den Boom et al. (2021), in Theorem 2, show that the expected KL divergence (with respect to $Q'_{i2}\tilde{z}_{ij}$) between the actual full conditional and the Gaussian approximation is bounded by two constants ϖ_1 and ϖ_2 :⁸

$$(8\mathbb{E}[\mathrm{KL}(p(\boldsymbol{Q}'_{i1}\tilde{\boldsymbol{z}}_{ij}\tilde{\boldsymbol{B}}_{ij,\bullet}|\boldsymbol{Q}'_{i2}\boldsymbol{y}_{ij},\boldsymbol{Q}'_{i2}\tilde{\boldsymbol{z}}_{ij})\times\mathcal{N}_{\sigma_{ij}}\mid|\hat{p}(\boldsymbol{Q}'_{i1}\tilde{\boldsymbol{z}}_{ij}\tilde{\boldsymbol{B}}_{ij,\bullet}|\boldsymbol{Q}'_{i2}\boldsymbol{y}_{ij},\boldsymbol{Q}'_{i2}\tilde{\boldsymbol{z}}_{ij})\times\mathcal{N}_{\sigma_{ij}})] \leq \varpi_1 + \varpi_2,$$

 ϖ_1 depends on the concentration properties of $p(\tilde{B}_{ij,\bullet}|Q'_{i2}y_{ij})$ around its mean and ϖ_2 on the average correlation between the elements in $\tilde{B}_{ij,\bullet}$.

If the posterior covariance of $\tilde{B}_{ij,\bullet}$ is small relative to $\sigma_{\varepsilon,ij}^2$ and when $k \ll K_{ij}, \varpi_1$ approaches zero. The second constant measures approximation errors between the first two moments of $p(\tilde{B}_{ij,\bullet}|Q'_{i2}y_{ij})$ and the approximating density. This quantity depends on the posterior covariance of the approximating density to the posterior of $\tilde{B}_{ij,\bullet}$ and thus the error bound can be small even if the approximate variance–covariance matrix differs sharply from the true covariance of the posterior of $\tilde{B}_{ij,\bullet}$. This finding also has important implications for our estimates of $A_{ij,\bullet}$ since (7) and (8) can be combined to arrive at an upper bound for the approximation error to the posterior distribution of $A_{ij,\bullet}$.

4. ILLUSTRATION USING SYNTHETIC DATA

The theoretical discussion in the previous section builds on certain assumptions about the rows of \tilde{z}_{ij} and the correlation properties of the posterior of $p(\tilde{B}_{ij,\bullet}|Q'_{i2}y_{ij})$. In this section, we will use synthetic data to illustrate how our approach performs under a realistic data-generating process (DGP).

Our DGP assumes that N = 10, K = 2, n = 20, and T = 500 and features a single lag. It is given by

$$\mathbf{y}_t = \hat{\mathbf{\Gamma}}_1 \mathbf{y}_{t-1} + \boldsymbol{\epsilon}_t, \quad \boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}_{20}, \boldsymbol{UHU}'),$$

where $\boldsymbol{H} = \sigma_{\varepsilon}^2 \times \boldsymbol{I}_{20}$ and $u_{ij} \sim \mathcal{N}(0, \sigma_{\varepsilon}^2/10)$ for i = 2, ..., n; j = 1, ..., i - 1. The matrix $\hat{\boldsymbol{\Gamma}}_1$ is obtained as follows: The blocks referring to the domestic coefficients are centered around the

⁸ The constants ϖ_1 and ϖ_2 take a complicated form. To avoid introducing additional notation, we summarize their main properties here. Precise definitions can be found in Van den Boom et al. (2021).

APPROXIMATE INFERENCE IN PANEL VARS

σ_{ε}	σ_{β}				
	0.01	0.025	0.05		
0.01	1.99	1.33	1.06		
	(1.96)	(1.29)	(1.00)		
0.025	1.19	1.08	1.04		
	(1.19)	(1.07)	(0.97)		
0.05	1.11	1.09	1.07		
	(1.14)	(1.10)	(1.03)		

 $T_{ABLE \ 1}$ simulation evidence: irga-pvar versus pvar estimated through MCMC

Notes: The table shows the relative mean absolute error (MAE) ratios between our IRGA-PVAR and a PVAR estimated using MCMC techniques for different values of σ_{ε} and σ_{β} . The numbers are averages across 100 draws from the DGP. The rows with numbers in parentheses refer to the MSE ratios for the full coefficient vector (i.e., including both international and domestic coefficients) whereas the other rows include MSE ratios for the domestic coefficients only.

same mean vector and we add Gaussian shocks:

$$\operatorname{vec}(\boldsymbol{\Gamma}_{i1}) = \operatorname{vec}\left[\begin{pmatrix} 0.8 & 0.2 \\ 0.3 & 0.6 \end{pmatrix}\right] + \sigma_{\beta}\boldsymbol{\eta}_{t}, \quad \boldsymbol{\eta}_{t} \sim \mathcal{N}(\boldsymbol{0}_{4}, \boldsymbol{I}_{4})$$

for all countries *i*. We simulate the elements in Ξ_i from $\mathcal{N}(0, \sigma_{\beta}^2)$. To obtain sparsity in Ξ_i , we randomly zero out elements such that we have around 60% zeroes in Ξ_i . The initial value of y_t is sampled from a multivariate zero-mean Gaussian with variance 0.01. To analyze how the goodness of our approximation changes with different measurement errors and cross-country heterogeneity, we consider $\sigma_{\varepsilon}, \sigma_{\beta} \in \{0.01, 0.025, 0.05\}$.

The main question is whether our IRGA-based approach yields estimates of the domestic coefficients which are close to the ones obtained from exact methods. Since exact methods become prohibitively slow in truly large data sets, we make this comparison operational by simulating from a moderately large DGP and estimate a PVAR model with a Horseshoe prior and without using IRGA (i.e., all coefficients are estimated using MCMC). For both models, we include p = 2 lags of y_t .

To avoid mixing up approximation errors arising from using IRGA to any errors in estimates that come from our equation-by-equation estimation approach, both the exact and approximate models are estimated based on the same (correct) ordering and in what follows we compare the coefficients A_{i1} and B_i to their true (implied) values.

Table 1 shows averages of relative mean absolute error (MAE) ratios between the IRGA-PVAR and the PVAR estimated through MCMC across 100 replications of the DGP. The white rows refer to MAE ratios for the domestic coefficients whereas the gray-shaded rows denote MAE ratios for all regression coefficients. For both models, we use the posterior median as our point estimator.

When we consider the results for the domestic coefficients (i.e., the white rows), we find a great deal of relative MAEs close to 1 (except for cases in which both σ_{β} and σ_{ε} are very small). In principle, when we fix a given row and consider increasing values of σ_{β} we find that the relative MAEs approach unity. This is consistent with the theoretical predictions in the previous section and shows that, at least when the posterior mean is considered, approximation accuracy of our IRGA-based approach increases with the ratio $\sigma_{\beta}/\sigma_{\varepsilon}$. When we fix a given column, as long as σ_{ε} is not too small, we find no discernible differences across values of σ_{β} .

Turning to the relative MAEs across all coefficients reveals that our approximate method also yields estimates of B_i which are competitive to the ones obtained using MCMC. In fact, for $\sigma_{\beta} = 0.05$ we find that the mean estimates essentially equal to the ones obtained from the MCMC-based method.



NOTES: $A_{ij,s}$ denotes the *s*th element of $A_{ij,\bullet}$ for s = 1, ..., k. The solid black line refers to the approximate posterior $\hat{p}(A_{ij,s}|\mathbf{y}_{ij})$ whereas the dashed black line denotes the posterior distribution $p(A_{ij,s}|\mathbf{y}_{ij})$ obtained through estimating the model using MCMC.

MARGINAL POSTERIOR DISTRIBUTION OF THE DOMESTIC COEFFICIENTS IN $A_{ii,\bullet}$ For three countries

To illustrate our approach using a single draw from the DGP for $\sigma_{\beta} = \sigma_{\varepsilon} = 0.05$, Figure 1 shows the marginal posterior distributions of the domestic coefficients for our IRGAbased PVAR (in solid black) and the MCMC-based estimates (in dashed black). This figure shows that in most cases, posterior distributions are similar. Especially when we focus on the mean/median, we observe only small differences across coefficients (with some few cases suggesting a larger disagreement between MCMC and approximate estimates). When we focus on the higher moments of the marginal distributions, we find similar variances, tail behavior, and skewness properties. This small discussion has shown that, at least when synthetic data are considered, our approach yields reasonable estimates.

5. A HUGE MODEL OF THE WORLD ECONOMY

In this section, we develop a huge-dimensional model of the world economy. The model is used to forecast output (measured by industrial production), inflation, LT-IR, and stock prices for a large panel of countries. We moreover analyze the properties of the forecasts using novel stochastic block models as well as Diebold–Yilmaz (DY) spillover indices.

5.1. *Data.* We have collected macroeconomic and financial data from the Organisation for Economic Co-operation and Development's (OECD's) short-term indicator data base. The data are monthly and span the period from 2001m2 to 2019m12. In principle, we have a panel of 18 series for 38 OECD countries but not all variables are available with the same country coverage (see Table 2). In total, we have 487 variables in our PVAR.

The series fall into three categories: macroeconomic, financial, and leading indicators. For macroeconomic variables, we consider measures of economic activity (industrial production growth, the output gap and the unemployment rate), export and import growth, consumer and producer price inflation as well as changes in earnings in the manufacturing sector.

Financial data cover short- and long-term interest rates (overnight, three-month money market rates and long-term government bond yields), changes in stock prices and broad money growth. For euro area countries, we include the three-month euribor as a measure of short-term interest rates.

APPROXIMATE INFERENCE IN PANEL VARS

OECD Code	Description	Trans.	Coverage (percent)	
Macroeconomic data				
PRINTO01	Industrial production, s.a.	1	81.6	
LORSGPRT	GDP, ratio to trend (output gap)	0	92.1	
LRHUTTTT	Harmonized unemployment rate: all persons, s.a.	0	76.3	
XTEXVA01	Exports in goods, s.a.	1	100.0	
XTIMVA01	Imports in goods, s.a.	1	100.0	
CPALTT01	Consumer prices: all items	1	92.1	
PIEAMP01	Producer prices—Manufacturing	1	60.5	
LCEAMN01	Hourly earnings: manufacturing, s.a.	2	39.5	
Financial data				
IRSTCI01	Overnight interbank rate	0	96.3	
IR3TIB01	Three-month interbank rate	0	92.6	
IRLTLT01	Long-term interest rate	0	78.9	
SPASTT01	Share prices	2	100.0	
MABMM301	Broad money, s.a.	2	28.9	
Leading indicators				
LOLITOAA	Leading indicator, amplitude adjusted	0	92.1	
BSCICP02	Manufacturing confidence indicator, s.a.	0	78.9	
CSCICP02	Consumer confidence indicator, s.a.	0	65.8	
SLRTCR03	Passenger car registrations, s.a.	2	28.9	
ODCNPI03	Permits issued for dwellings, s.a.	2	31.6	

TABLE 2 DATA DESCRIPTION

Notes: All data are from the OECD's short-term indicator database, accessed via dbnomics (https://db.nomics. world/). Transformations refer to no-transformation (0), year-on-year growth rates (1), and month-on-month growth rates (2). Country coverage of each variable in percent. The sample consists of Austria, Australia, Belgium, Brazil, Canada, Switzerland, China, Czechia, Germany, Denmark, Estonia, Spain, Finland, France, Great Britain, Greece, Hungary, Indonesia, Ireland, Israel, Italy, Japan, Korea, Lithuania, Latvia, Mexico, the Netherlands, Norway, New Zealand, Poland, Portugal, Russia, Sweden, Slovenia, Slovakia, Turkey, the United States, and South Africa.

In addition to this rather standard macrofinancial data set, we gathered data on confidence/sentiment and short-term leading indicators. These comprise the OECD's leading indicator, which is a constructed measure to provide early signals of turning points in business cycles,⁹ measures of manufacturers' and consumers' confidence, as well as changes in passenger car registration and newly permitted dwellings. Confidence measures potentially contain additional information to predict economic activity (Batchelor and Dua, 1998; Ludvigson, 2004) as do car registrations and new dwellings for consumer expenditures. Changes or growth rates refer to either year-on-year or month-on-month growth rates. For a detailed overview, see Table 2.

5.2. Design of the Forecasting Exercise and Competing Models. We carry out a forecasting exercise comparing an unrestricted PVAR with a Horseshoe prior to a range of alternatives. In what follows, we focus on predictions for four *target variables*: consumer price inflation (labeled Infl. in the tables and figures), industrial production (Ind. prod.), stock returns (Equities), and LT-IR.

Our forecasting design is recursive, implying that we use the data from 2001m2 to 2006m12 as an initial estimation period. We use data through 2006m12 to produce one-month up to 12-month-ahead forecast distributions. This initial estimation period is then extended by data for 2007m1 and forecast densities for 2007m2 (up to 2008m1) are constructed. We repeat this procedure until we reach the end of the holdout period.

⁹ This measure shows fluctuations of an economy's activity around its long-term trend (i.e., potential output). The main difference to the output gap measure we employ is that the leading indicator is amplitude adjusted. For more information, see http://www.oecd.org/sdd/leading-indicators/oecdcompositeleadingindicators referenceturningpointsandcomponentseries.htm.

FELDKIRCHER ET AL.

TABLE 3						
ESTIMATION TIME IN MINUTES						

BVAR	FAVAR-10	GVAR	PVAR-IRGA
235.4	663.7	566.4	113.4

NOTES: BVAR and FAVAR-10 are estimated one country at a time. The indicated time marks the time required for estimating all forecasts for the full system. GVAR and PVAR-IRGA estimation times are based on joint estimation of the multicountry system.

To compare point forecasts across models, we use MAEs. Since this disregards higher-order moments of the predictive distribution we also use log predictive likelihoods (LPLs) to compare the density forecast performance of these alternatives.

The set of competing models is chosen not only to reflect a variety of popular approaches, but also to be computationally practical. In particular, alongside the proposed PVAR-IRGA approach, we consider the following models. First, to assess the role of allowing for cross-country spillovers, we estimate single-country BVARs that rule out both static and dynamic interdependencies. These are estimated country-by-country. Second, to compare our approach to other models incorporating international information, we consider two types of specifications. The first is a factor-augmented VAR (FAVAR-10) model, which augments the single-country BVARs with 10 factors extracted from the nondomestic country variables. This procedure serves to obtain a lower-dimensional representation of the international information set. As a second option to include international information, we use a Bayesian GVAR.

To ensure that differences in forecast performance are not driven by the respective priors on the VAR coefficients, we estimate all models with Horseshoe priors and set the number of lags equal to two. As a robustness check, we also repeat the forecasting exercise replacing the Horseshoe prior with a conventional Minnesota prior for each of the PVAR-IRGA, BVAR, FAVAR-10, and the GVAR. Further specification details and results are given in Appendices A.2, A.3, and A.4.

Other than PVAR-IRGA, all of the models are estimated using exact MCMC methods. We stress that approaches which involve using MCMC methods for an unrestricted 487 dimensional VAR would simply be computationally impractical. Computation times (average estimation time per model over all periods in the holdout) for doing the pseudo real-time forecasting exercise are provided in Table 3. It can be seen that PVAR-IRGA is substantially faster than any of the competing approaches, even though all of the latter are much more parsimonious models. Notice that estimating the PVAR-IRGA is even faster than estimating a set of *N* country-specific VARs of the same size. This is because MCMC sampling of the domestic coefficients based on the IRGA posterior is faster since it only relies on a part of the likelihood function to form the conditional posterior distributions.

5.3. Summary of Forecasting Results. The results of our forecasting exercise are summarized in Table 4. It presents absolute values of MAEs and LPLs for our PVAR-IRGA approach for two forecast horizons. Results for the other approaches are benchmarked relative to these. To be precise, for MAEs we take ratios relative to PVAR-IRGA (with numbers exceeding unity implying that the PVAR improves upon the competitors), for LPLs we take differences relative to PVAR-IRGA (with values below zero suggesting that the PVAR is outperforming the respective model). The numbers in the table are GDP-weighted averages of country-specific LPLs computed by using GDP in 2015 U.S. dollars averaged over the period 2002–19. To provide a rough gauge of model performance across countries, the numbers in parentheses represent the percentage of countries in which a given model performs best in absolute terms.

The most important finding is that PVAR-IRGA works—it produces sensible forecasts quickly. Bayesian estimation of huge-dimensional PVARs has been made possible through the use of IRGA methods. The other main finding is that (with some exceptions) PVAR-IRGA

APPROXIMATE INFERENCE IN PANEL VARS

MAE LPS Ind. Prod. LT-IR Infl. LT-IR Model Equities Equities Ind. Prod. Infl. h=1**BVAR** 1.008 0.988 0.927 1.000 -4.6862.538 13.855 -0.503(31.6)(9.7)(36.7)(22.9)(13.2)(12.9)(13.3)(14.3)FAVAR-10 1.004 0.981 0.923 0.994 -4.5114.022 14.664 0.713 (23.7)(51.6)(53.3)(60.0)(2.6)(51.6)(66.7)(45.7)PVAR-IRGA 0.274 -177.799-39.815 -13.3870.825 0.410 0.386 36.181 (44.7)(25.8)(10.0)(17.1)(68.4)(25.8)(20.0)(40.0)-42.623 GVAR 1.051 1.092 1.090 1.199 -12.541-21.086-44.480(0.0)(12.9)(0.0)(0.0)(15.8)(9.7)(0.0)(0.0)h = 120.984 0.912 9.251 -10.936**BVAR** 1.081 1.020 16.181 -4.867(39.5)(32.3)(40.0)(40.0)(42.1)(25.8)(50.0)(34.3)FAVAR-10 5.545 -15.8780.987 1.089 0.896 1.020 14.098 -0.372(28.9)(3.2)(16.7)(11.4)(5.3)(6.5)(20.0)(22.9)PVAR-IRGA 0.875 0.664 0.570 0.759 -204.541-149.824-104.962-174.713(28.9)(48.4)(36.7)(45.7)(47.4)(61.3)(30.0)(40.0)

 Table 4

 summary of forecast exercise, all models with horseshoe prior

Notes: GDP-weighted average over countries, win percentage across countries in parentheses. Root mean squared error (RMSE) and log predictive score (LPS) relative to the benchmark. The benchmark PVAR-IRGA shows actual values, all other models are in ratios to the benchmark for RMSEs and in differences for LPSs.

1.135

(2.9)

-1.333

(5.3)

-13.215

(6.5)

-15.904

(0.0)

-30.407

(2.9)

0.981

(6.7)

GVAR

1.048

(2.6)

1.072

(16.1)

works well and is highly competitive with competing approaches. These improvements are limited for point forecasts but sometimes very pronounced for LPLs which measure density forecast performance. For short-term forecasts of equity returns in particular, PVAR-IRGA is producing strong improvements in LPLs whereas its performance is slightly weaker for the remaining variables under consideration.

When we focus on higher-order forecasts, the relative performance of PVAR-IRGA improves. Whereas we observe that predictive accuracy is deteriorating for equities at the 12-month-ahead horizon, forecasts of industrial production and inflation improve considerably. For the latter two variables, the PVAR-IRGA is the single best performing model.

BVAR is the only alternative that does not allow for any cross-country spillovers. On average, it forecasts fairly well for both forecast horizons. Improvements, however, are more pronounced at the one-month-ahead horizon. For 12-month-ahead, we find that taking cross-country linkages into account helps forecast accuracy. This suggests that for short-run forecasts, cross-country spillovers are not that strong (or at least do not significantly help in predicting output, inflation, and long-term rates).

The fact that the single-country BVARs improve upon the GVAR indicates that having a smaller-sized sparse model seems to be more important than taking into account crosscountry linkages for improving forecasts. The good performance of our PVAR-IRGA suggests that the Horseshoe prior successfully strikes a balance between exploiting cross-country information and sparsity. Since the model is unrestricted it also allows the data to speak about the precise form of cross-country linkages.

The statements in the preceding paragraphs are based on an examination of LPLs. Analyzing MAEs reveals similar patterns, but to a weaker extent. This indicates that the benefits of unrestricted modeling of the high-dimensional PVAR offers some benefits in terms of point forecast performance, but the benefits are larger for density forecasts.

The discussion has focused on averages across countries (and time periods). However, it could be that a model (such as our PVAR-IRGA) yields lower-average LPLs but still provides the best performance for individual countries in our sample. Considering the percentage



cumulative LPS for pvar-irga relative to bvar for h=1 forecasts across countries and over time

of wins for each model corroborates the findings based on LPLs and MAEs. But it is worth emphasizing that even though the PVAR-IRGA is sometimes outperformed by simpler competitors such as BVAR, for most variables we still find a sizable fraction of wins across countries for our proposed model. In the case of one-month-ahead inflation density forecasts, this share is about 34% whereas it is around 47% for longer-run forecasts of equity returns (in which BVAR outperforms the PVAR-IRGA if we consider LPLs). These sizable shares suggest that it might be worthwhile to carefully analyze country-specific results.

5.4. Forecast Comparison across Countries and over Time. In the preceding subsection, we compared the average (over countries and time) forecast performance of PVAR-IRGA to various alternatives. In this subsection, we look behind the average performance to investigate forecast performance at the country level and see how it changes over time. We do so through heatmaps of cumulative LPLs for PVAR-IRGA for the individual countries. To assess when it pays off to allow for cross-country linkages (and because it is the strongest competitor to our approach) we benchmark our results to the LPLs of the BVAR.

Figures 2 and 3 contain these heatmaps for the two forecast horizons. The figures are grouped into four categories: Advanced European, Emerging European, Advanced Other, and Emerging Other and individual countries are labeled using ISO country codes. Intensifying shades of blue (red) indicate stronger support for PVAR-IRGA (BVAR).



cumulative LPS for pvar-irga relative to bvar for h=12 forecasts across countries and over time

We first focus on one-month-ahead forecasts. Starting with inflation predictions, we observe that the stronger performance of the BVAR is mainly driven by a weak performance of the PVAR-IRGA in emerging economies (most notably Brazil and Turkey, with some exceptions, such as India) whereas it performs best for developed economies such as the United States, France, and Ireland. This brief discussion shows why it is important to also consider results at the country level. If the researcher is interested in short-term forecasting of U.S. inflation and has to decide on one of the models we consider, focusing on overall LPLs masks the particularly strong U.S.-specific forecasting performance.

When we focus on industrial production we find a somewhat different pattern, with PVAR-IRGA being outperformed by the BVAR for several countries in Advanced Europe (except for Austria, Finland, Norway, and the Netherlands) and some gains in several countries located in Emerging Europe (Slovakia, Russia, and the Baltics). For some countries (such as Sweden and Italy), the PVAR-IRGA performs well prior to the global financial crisis. The rapid decline in output in the final half of 2008, however, led to a deterioration in forecast performance. This is because PVAR-IRGA yields predictive distributions which are sometimes too tight and thus capturing outliers becomes increasingly difficult.

Turning to the results for equity returns reveals a great number of blue-colored cells. In principle, our model works well for most economies (with a slightly weaker performance for, e.g., Portugal, Norway, Slovenia, and Japan). Interestingly, we also find some heterogeneity

with respect to model performance over time. In the case of the United States, for instance, our model only improves upon the BVAR from 2012 onward. We conjecture that the slightly weaker performance prior and during the financial crisis is, again, driven by too tight predictive intervals. But these tight intervals then help in predicting returns after the financial crisis, a period characterized by steady increases in U.S. stock markets.

The PVAR-IRGA displays the weakest performance for LT-IR. For some few countries (e.g., Portugal, Greece, Latvia, and Israel) the PVAR performs well. In general, the weak performance for long-term rates is driven by the fact that these display a downward trend during the holdout period for most countries. PVAR-IRGA captures this downward trend rather well but the predictive variance is considerably smaller than the one of the single-country BVAR. Hence, under the predictive distribution of the PVAR-IRGA, even relatively small changes in long rates have strong effects on LPLs. The countries which depart from this general pattern (such as Portugal and Greece) feature large spikes in LT-IR. The PVAR captures this well and quickly adjusts the predictive variance. Since it takes slightly longer for the BVAR to adjust we conjecture that this quick increase is mostly driven by the large information set.

In the previous subsection, we have shown that on average, the PVAR-IRGA produces the most precise density forecasts for inflation and industrial production when the forecast horizon is increased whereas the performance for equity returns deteriorates. When we focus on 12-month-ahead predictions (see Figure 3) we find that the strong overall performance for inflation is mostly driven by excellent forecasts in major developed economies such as the United States or Japan. A similar pattern is found for industrial production. Again, we find that the PVAR-IRGA produces precise density forecasts for most developed economies located in Europe as well as the United States and Japan. However, it is also worth stressing that the PVAR also produces accurate forecasts for developing economies such as Turkey as well as several countries located in Central Eastern Europe. This strong performance is driven by more precise point forecasts but also by the fact that the predictive distributions for multistep-ahead forecasts seem to be heavy tailed and thus make observing outliers more probable.

For equities and LT-IR, we find that the PVAR performs slightly weaker than the BVAR. Especially for equities, this is driven by a particularly bad performance in the United States. For LT-IR, we again find that the PVAR is competitive when used to forecast long-term rates in Portugal and Israel but it appears to be outperformed in countries such as Denmark.

In summary, PVAR-IRGA yields precise equity return predictions for short-term forecasts and shows good performance when used to forecast inflation in major economies such as the United States. For higher-order forecasts, the results somewhat reverse and the PVAR-IRGA works particularly well when it is used to predict industrial production and inflation. In general, and this is consistent with the findings based on average LPLs, we find that more predictive evidence in favor of cross-country spillovers increases with the forecast horizons. We will provide additional evidence on the increasing importance of cross-country spillovers in the following subsections.

5.4.1. A quantitative analysis of the point forecasts. In the previous section, we have shown that our PVAR yields highly competitive forecasts and often improves upon other single- or multicountry models. In this section, our aim is to quantitatively analyze the country-specific forecasts to investigate the role of cross-country correlations in the point predictions. For brevity, we focus on one-month-ahead forecasts for industrial production and LT-IR.

Since the corresponding correlation matrices are high-dimensional and thus difficult to interpret, we use techniques from network analysis (see, e.g., Žiberna, 2014) to search for clusters in the correlation matrices of the posterior median and the posterior standard deviation of the forecasts.¹⁰ Intuitively speaking, we reorder the rows and columns such that correlations between countries are grouped into R = 8 distinct blocks.¹¹ The relations (i.e., positive correlations) within a given block are maximized whereas the relations of countries within a block to economies outside of a block are much less important (or even negative).

Since this algorithm needs a correlation matrix of forecasts, we compute the initial correlation matrix based on the first 12 observations and then expand this window until we reach the end of the holdout period which is used to compute forecast distributions. This yields a sequence of correlation matrices which we then analyze using a stochastic block model.

In what follows, we focus on one-month-ahead forecasts for three distinct periods in our holdout sample. First, we examine correlation structures among our forecasts for 2009m1, the onset of the global financial crisis. The second period we consider is 2012m6, the month prior to Mario Draghi's famous "whatever it takes" speech, which marks the height of the euro area sovereign debt crisis. The final period is 2019m12, the end of our sample. Using all available information is a natural choice to investigate how forecasts are related. Figure 4 shows the correlation matrices (multiplied by 10) for industrial production forecasts sorted using a stochastic block model.

The figure shows the R = 8 clusters on the main diagonal of the matrix. The first two blocks are each defined by a single country (Norway and Portugal, respectively). Forecasts between the blocks are (modestly) negatively correlated, as indicated by the red shading in the offdiagonal elements between the blocks. Considering the remaining countries, we see that forecasts for Norway are not only negatively correlated with those for Portugal, but also with the rest of the sample. As an oil-based economy, Norway was considerably less affected by the global financial crisis than the rest of the sample. Looking at the remaining economies reveals two large clusters. One of them contains forecasts for Greece, Italy, Spain, and the Baltics countries that showed massive contractions in output during the global financial crisis. The Irish economy, which also significantly contracted during the global financial crisis, appears in another cluster. Another large cluster contains countries that were comparably less affected by the crisis. Importantly though and with the exception of Norway, all clusters are positively correlated indicated by the gray to black color shading in Figure 4. This reflects the global nature of the financial crisis.

Considering the period of the euro area sovereign debt crisis, depicted in the upper right panel of Figure 4 reveals a very similar picture: Norway stands out, and forecasts of the remaining countries are positively correlated. One big cluster emerges which covers mostly European economies and Russia—the latter which shares strong trade ties with the European Union. Other European countries that were not affected by the sovereign debt crisis, such as the Baltics, are allocated into different blocks and Ireland again appears in a separate cluster. Turning now to one-month-ahead forecasts for the end of the holdout sample, we still find isolated countries such as Norway and Ireland, two European clusters and one large, international cluster. The latter one shows only a modest within-correlation, indicating that at the end of our sample, with the pandemic yet to fully unfold, forecasts are less strongly correlated as in periods of severe downturns.

Figure 5 shows the same analysis for LT-IR forecasts. During the period of the global financial crisis, long-term rates spiked in countries like Greece, but have been downward-trending in countries considered as safe havens: the United States, Canada, and Japan. PVAR-IRGA forecasts are consistent with these historical observations putting the aforementioned country groups into separate blocks. Two further clusters emerge: one solely consisting of European economies, whereas the other one contains both Advanced European and non-European economies. At the height of the sovereign debt crisis, LT-IR shot up for most economies but to a different extent. This is mirrored in the first four correlation blocks depicted on the main diagonal of the correlation matrix. Forecasts for crisis-stricken economies, such as Ireland, Italy, and Spain are strongly correlated and appear in a separate block. Interestingly, Greece



CORRELATION MATRICES OF ONE-MONTH-AHEAD FORECASTS FOR INDUSTRIAL PRODUCTION SORTED USING A STOCHASTIC BLOCK MODEL

and Portugal appear in one block, but forecasts between these clusters are positively correlated. Both blocks of crisis-stricken countries are clearly separated from the rest of the sample. Last and looking at the end of sample period, we see a very different picture. With the exception of South Africa and to some extent Greece, global long-term interest forecasts are very homogeneous and positively correlated.

Summing up, examining correlation structures of one-month-ahead forecasts revealed insights as to which extent the model is capable of mirroring correlation structures present in the data. Correlations are strongest during periods of simultaneous contractions such as witnessed during the global financial crisis. Looking at the episode of the sovereign debt crisis, which unfolded on a more regional basis, also reveals clusters but they differ from those identified during the global financial crisis. This highlights the overall flexibility of the model. For some countries, forecasts are always separated from the rest of the sample (e.g., Norway as an oil-exporting economy). This implies that the PVAR-IRGA can take cross-country links into





CORRELATION MATRICES OF ONE-MONTH-AHEAD FORECASTS FOR LONG-TERM INTEREST RATES SORTED USING A STOCHASTIC BLOCK MODEL

account when they are important but at the same time does not enforce them on the whole set of countries—a flexibility which is of ample importance when dealing with large, heterogeneous cross-sections.

5.4.2. Other evidence of cross-country spillovers. Another way of looking at cross-country correlations is to use the DY spillover index (see Diebold and Yilmaz, 2009). This index distinguishes to which extent forecast error variance can be explained by its own history as opposed to effects through all other variables in the system. The latter effects are dubbed "spillovers" and serve as a measure of overall connectivity. For instance, in this article, we calculate the index based on a generalized forecast error variance decomposition (GFEVD) which avoids order dependence with respect to the elements in y_t . We report the total share of effects from nondomestic variables (spillovers from one variable to another) in the GFEVD. The calculation of the underlying GFEVD is recursive (in the same manner as was done for the



NOTES: This index indicates the share of spillovers across countries excluding spillovers to variables within a given country (summed across variable types). Estimated based on an expanding window. The solid line is the posterior median alongside the 68% posterior credible set.

DIEBOLD–YILMAZ SPILLOVER INDEX FOR ALL COUNTRIES BASED ON DECOMPOSITION OF FORECAST ERROR VARIANCE OF THE FULL SYSTEM

forecasting exercise), and we show results for a forecast horizon of 12 and 24 months. Our focus on higher-order spillovers is motivated by evidence reported in the literature on GVARs (see, e.g., Feldkircher and Huber, 2016) which shows that cross-country spillovers (measured through FEVDs) become sizable only after several quarters.

Figure 6 shows the DY index over the holdout period for both horizons. Most importantly, our results point at a sizable degree of cross-country connectedness. At the end of our sample period, the DY indices amount to about 45 at the 12-month-ahead forecast horizon and to 68% at the 24-month-ahead horizon. Investigating the DY index over time, reveals a steady increase of the index until 2016 after which the indicator levels out.¹² A particularly pronounced increase of the index can be observed during the period of the euro area sovereign debt crisis (between 2010 and 2012). In general, economic variables tend to comove more strongly during turbulent times (see, e.g., Pham and Sala, 2022)—a pattern that we also observe with our data. It is worth stressing that our estimates are surrounded by considerable posterior uncertainty, which tends to attribute more posterior mass above than below the posterior median. This behavior is even more pronounced during the period of the sovereign debt crisis.

The discussion up to this point focused on an overall measure of cross-country connectivity considering all focus variables jointly. To investigate whether connectivity plays a larger role for certain variables, we display the DY index for each of the focus variables separately in Figure 7.

The figure reveals some interesting variable-specific differences. For example, the evolution of the DY index for industrial production, inflation, and LT-IR is very similar to the behavior of the overall index displayed in Figure 6. The degree of connectedness, is however, comparably smaller (35%). It is worth stressing that the posterior distribution is again strongly tilted toward higher levels of connectivity. A high degree of cross-country dependence between inflation rates is in line with Borio and Filardo (2007), Ciccarelli and Mojon (2010), and Kabukçuoğlu and Martínez-García (2018) who stress the importance of a global component in determining domestic inflation rates.

¹² Note that consistent with the forecasting exercise, we use an expanding window to calculate the DY index which automatically introduces a certain degree of persistence. Appendix A.4 provides additional empirical results for estimates using a rolling window of observations.



Notes: This index indicates the share of spillovers between countries one variable at a time (summed across countries). Estimated based on an expanding window. The solid line is the posterior median alongside the 68% posterior credible set.

DIEBOLD–YILMAZ SPILLOVER INDEX BY VARIABLE BASED ON DECOMPOSITION OF FORECAST ERROR VARIANCE OF THE FULL SYSTEM

Spillovers between equity prices show a distinct dynamic. The degree of connectivity in financial markets is generally higher compared to that of the remaining variables. The associated DY index is about 25% at the beginning of the holdout period and rises sharply during the global financial crisis. This finding is in line with Demirer et al. (2018), who demonstrate a strong increase in equity connectedness between banks during periods of financial stress. At the end of the sample period, the DY index amounts to about 75%. Notably, there is also little difference between the two forecast horizons.

In Figure 8, we repeat the calculation and measure the role of cross-country effects for each country. In general, most of the countries considered display sustained increases in their respective DY indices peaking at about 75%. In some emerging economies, such as China and Indonesia, the index increases sharply right after 2008. Some notable exceptions are Turkey and Mexico which experienced a decrease in connectivity after the global financial crisis surrounding the Fed's tapering statement in June 2013. In the case of Greece, we observe a gradual decline in connectivity during the euro area sovereign debt crisis. This finding reflects a strong domestic component which complies with the fact that Greece was the epicenter of the crisis at the time.

Summing up and jointly considered with the forecasting results, the cross-variable DY indices paint a similar and consistent picture. Benefits from estimating a large multicountry system tend to increase with the forecast horizon. This finding is mirrored by increasing levels of connectivity if higher-order forecasts are considered. Our results indicate that the PVAR is capable of, in light of heavy shrinkage introduced through the Horseshoe prior, capturing cross-country relations flexibly. As opposed to other models (such as the FAVAR or the GVAR), the PVAR-IRGA introduces no particular restrictions on the coefficients. This feature is crucial to appropriately control for cross-country heterogeneity in light of a diverse set of countries such as the one we have in our data set.

Appendix A.4 provides additional results for the spillover indices computed using rolling windows of varying length. Using a rolling instead of an expanding window implies that past observations do not impact the estimates at some point. Thus, parameters are quicker to



NOTES: This index indicates the share of foreign spillovers to the indicated country (summed across variables). Estimated based on an expanding window. The solid line is the posterior median alongside the 68% posterior credible set.

DIEBOLD–YILMAZ SPILLOVER INDEX BY COUNTRY BASED ON DECOMPOSITION OF FORECAST ERROR VARIANCE OF THE FULL SYSTEM

adjust to new information (at the cost of discarding past information). Results for both the overall and the variable-specific DY indices tell a similar story to the ones based on an expanding window. The main difference is that the estimates feature more movements in periods of economic turmoil such as the global financial crisis and the sovereign debt crisis in Europe.

6. CONCLUDING REMARKS

Multicountry VARs have the potential to be enormous and simply working with unrestricted versions of them leads to overparameterization and computational problems. The existing literature typically deals with these problems by imposing restrictions or reducing the dimension of the data. But the former strategy risks misspecification and the latter risks losing information. Accordingly, in this article, we have developed Bayesian methods for working with unrestricted VARs and rely on the Horseshoe prior to gain in parsimony by imposing shrinkage in a data-based fashion. Existing Bayesian work with VARs with such a prior has been done using MCMC methods. These are too computationally demanding to be used with PVARs with hundreds of dependent variables. In this article, we have used IRGA methods to overcome this computational hurdle. We show that these allow for practical inference even in PVARs of huge dimension. Our macroeconomic empirical application demonstrates the benefits of being able to work with such large PVARs.

APPENDIX A

A.1 Vector Approximate Message Passing (VAMP). In this section, we provide a brief introduction on VAMP. Before we provide details on the VAMP algorithm for the panel vector auto regression (PVAR), a (very) brief introduction to message passing and the sum-product algorithm is in order.¹³ In general, message passing involves factorizing a joint density in an efficient way. Let $p(x_1, \ldots, x_K)$ denote such a joint density over K discrete random variables $\mathbf{x} = (x_1, \ldots, x_K)'$. Moreover, let us assume that we are interested in the marginal distribution of some variable x_j with \mathbf{x}_{-j} denoting all except the *j*th variable. The marginal distribution of x_j is then simply given by

$$p(x_j) = \sum_{x_{-j}} p(\boldsymbol{x})$$

which is a summation over K - 1 random variables. For large K, calculating this marginal distribution using this formula becomes computationally infeasible. To get around this problem, we write the marginal density in a different way based on the concept of a factor graph. The factors are functions of the random variables. Under certain conditions (which are almost always fulfilled in econometric models), we can factorize the joint density $p(\mathbf{x})$ as follows (see, e.g., (8.59) in Bishop, 2006):

$$p(\mathbf{x}) = \prod_{s} f_{s}(\mathbf{x}_{s})$$

with x_s denoting a subset of x and f_s being a factor. These factors encode the relationship between variables. We make use of this idea to represent the linear Bayesian regression model in terms of a factor graph and derive the relevant message passing algorithm which we then approximate using VAMP.

To set the stage, let us assume that $\mathbf{x} = (x_1, x_2)'$ and these two random variables can be factorized as follows:

$$p(x_1, x_2) = f_a(x_1) f_b(x_1, x_2) f_c(x_2).$$

This implies that, for instance, the factor f_b establishes a relationship between x_1 and x_2 . The precise functional form of f depends on the application. Notice that the factors in this decomposition depend on subsets of variables with $\mathbf{x}_a = \{x_1\}, \mathbf{x}_b = \{x_1, x_2\}$, and $\mathbf{x}_c = \{x_2\}$. The corresponding factor graph is depicted in Figure A.1.

In this toy example, the marginal distribution of, say, x_1 is proportional to:

$$p(x_1) \propto \sum_{x_2} p(x_1, x_2)$$

= $\sum_{x_2} f_a(x_1) f_b(x_1, x_2) f_c(x_2)$
= $f_a(x_1) \sum_{x_2} f_b(x_1, x_2) f_c(x_2).$

¹³ For a textbook introduction, see Bishop (2006).



Using the decomposition in the final line is computationally more efficient than computing the marginal distribution naively through the joint distribution.

In what follows, we let $m_{f_i \to x_j}$ denote the message (or information) from factor node f_i ($i \in \{a, b, c\}$) to variable x_j (j = 1, 2) whereas $m_{x_j \to f_i}$ refers to a message from a factor to a variable node. Since $f_c(x_3)$ depends exclusively on x_3 the message sent is equal to the factor node itself, $m_{f_c \to x_3} = f_c(x_3)$. These messages convey all relevant information from a factor to a node and vice versa. In the next step, we derive all relevant messages and show that the marginal distribution of x_1 is proportional to the product of all incoming messages.

Starting on the right part of Figure A.1(a), the message from factor node f_c to x_2 is

$$m_{f_c \to x_2} = f_c(x_2).$$

The message from x_2 to factor node f_b equals

$$m_{x_2 \to f_b} = f_c(x_2).$$

Finally, the messages from node f_b to x_1 are

$$m_{f_b \to x_1} = \sum_{x_2} f_b(x_1, x_2) m_{x_2 \to f_b}.$$

Intuitively speaking, $m_{f_b \to x_1}$ captures the information flow from the nodes to the right of the variable node x_1 . If we would use this information exclusively we would miss all information that arises from the node to the left of x_1 . Because f_a is an exterior node, the message from f_a to x_1 is simply $m_{f_a \to x_1} = f_a(x_1)$.

The product over all incoming messages is proportional to the marginal distribution $p(x_1)$:

(A.2)
$$p(x_1) \propto m_{f_b \to x_1} m_{f_a \to x_1} = m_{f_b \to x_1} f_a(x_1).$$

Deriving the marginal distribution of x_2 analogously yields:

(A.3)
$$p(x_2) \propto m_{f_b \to x_2} \mu_{f_c \to x_2} = \mu_{f_b \to x_2} f_c(x_2).$$

Bishop (2006) highlights that there exists a close relationship between the belief about a variable x_j (which is defined as the product of all incoming messages to the variable node), $b(x_j)$, and the corresponding marginal distribution. If the graphical model is a tree, the beliefs converge to the marginal distribution after one iteration of a message passing algorithm.

Algorithms exploiting Equations (A.2) and (A.3) are labeled sum-product algorithms (see, e.g., Korobilis, 2021). Notice that to arrive at the marginals of x_1 (x_2), we need to marginalize over x_2 (x_1). This summation is often difficult to compute. As a solution, researchers often rely

on approximations to these sums (or integrals more generally) and arrive at so-called approximate message passing algorithms (Donoho et al., 2009).

Before we discuss such approximations in more detail, we show how to design a message passing algorithm for the general regression model. Let y denote a T-dimensional response vector and X is a $T \times K$ matrix of regressors. The corresponding regression coefficients are denoted by β and the error variance is given by σ^2 . Furthermore, we let $\Pi(\beta)$ denote the prior. We assume throughout that any hyperparameters associated with the prior and σ^2 are known.¹⁴

The posterior of $\boldsymbol{\beta}$ (conditional on σ^2) is given by

(A.4)
$$p(\boldsymbol{\beta}|\sigma^2, \boldsymbol{y}, \boldsymbol{X}) \propto \mathcal{N}(\boldsymbol{y}|\boldsymbol{X}\boldsymbol{\beta}, \sigma^2 \boldsymbol{I}_T) \Pi(\boldsymbol{\beta}).$$

To derive the graphical representation of the regression model, it proves convenient to introduce a copy of β , labeled β^* . This allows us to rewrite the posterior in (A.4) as follows:

(A.5)
$$p(\boldsymbol{\beta}|\sigma^2, \boldsymbol{y}, \boldsymbol{X}) \propto \mathcal{N}(\boldsymbol{y}|\boldsymbol{X}\boldsymbol{\beta}^*, \sigma^2 \boldsymbol{I}_T)\delta(\boldsymbol{\beta} - \boldsymbol{\beta}^*)\Pi(\boldsymbol{\beta}),$$

with δ denoting the Dirac Delta function which equals one if $\beta = \beta^*$.

The factorization in (A.5) closely resembles the one used in our toy example. If $\mathbf{x}_1 = \boldsymbol{\beta}$, $f_a(\mathbf{x}_1) = \Pi(\boldsymbol{\beta})$, $\mathbf{x}_2 = \boldsymbol{\beta}^*$, $f_b(\mathbf{x}_1, \mathbf{x}_2) = \delta(\boldsymbol{\beta} - \boldsymbol{\beta}^*)$, and $f_c(\mathbf{x}_2) = \mathcal{N}(\mathbf{y}|\mathbf{X}\boldsymbol{\beta}^*, \sigma^2 \mathbf{I}_T)$, we can easily derive the belief of $\boldsymbol{\beta}$ and $\boldsymbol{\beta}^*$ as follows:

(A.6)
$$b(\boldsymbol{\beta}) \propto m_{\Pi \to \boldsymbol{\beta}} m_{\delta \to \boldsymbol{\beta}} = m_{\delta \to \boldsymbol{\beta}} \Pi(\boldsymbol{\beta}),$$

(A.7)
$$b(\boldsymbol{\beta}^*) \propto m_{\delta \to \boldsymbol{\beta}^*} m_{\mathcal{N} \to \boldsymbol{\beta}^*} = \mathcal{N}(\boldsymbol{y} | \boldsymbol{x} \boldsymbol{\beta}^*, \sigma^2 \boldsymbol{I}_T) m_{\delta \to \boldsymbol{\beta}^*}$$

Van den Boom et al. (2021) show how (A.6) and (A.7) can be used to derive a message passing algorithm which exploits certain properties of the structure of the graphical model outlined above. This message passing algorithm cycles between the following updating steps:

(A.8)
$$b(\boldsymbol{\beta}) \propto m_{\delta \to \boldsymbol{\beta}} \Pi(\boldsymbol{\beta})$$

(A.9)
$$m_{\beta \to \delta} \propto \frac{b(\boldsymbol{\beta})}{m_{\delta \to \boldsymbol{\beta}}},$$

(A.10)
$$b(\boldsymbol{\beta}^*) \propto m_{\boldsymbol{\beta} \to \delta} \mathcal{N}(\boldsymbol{y} | \boldsymbol{X} \boldsymbol{\beta}^*, \sigma^2 \boldsymbol{I}_T)$$

(A.11)
$$m_{\delta \to \beta^*} = m_{\beta^* \to \delta} \propto \frac{b(\beta^*)}{m_{\beta \to \delta}}$$

The corresponding updates are then obtained as follows: Initialize $m_{\delta \to \beta} = \mathcal{N}(\boldsymbol{\beta}|\boldsymbol{\xi}^{(0)}, \boldsymbol{\zeta}^{(0)}\boldsymbol{I}_K)$. For iteration *j*, the belief about $\boldsymbol{\beta}$ is approximated through a Gaussian distribution

(A.12)
$$b(\boldsymbol{\beta}) \approx \mathcal{N}\left(\boldsymbol{\beta} | \overline{\boldsymbol{\beta}}^{(j)}, s^{(j)} \boldsymbol{I}_{K}\right)$$

¹⁴ In practice, we will update them using expectation maximization steps using the quantities presented below.

with $\overline{\boldsymbol{\beta}}^{(j)} = \mathbb{E}_{b(\boldsymbol{\beta})}(\boldsymbol{\beta})$ and $s^{(j)} = \text{Tr}(\boldsymbol{\Omega}^k)/K$ with $\boldsymbol{\Omega}^{(j)} = \text{cov}_{b(\boldsymbol{\beta})}(\boldsymbol{\beta})$ with Tr denoting the trace operator. Using (A.8) yields the first updating step of the algorithm:

(A.13)
$$\overline{\beta}_i^j = \mathbb{E}\left(\beta_i | \xi_i^{(j)}, \zeta^{(j)}\right) \quad \text{with } s^{(j)} = \sum_{i=1}^K \operatorname{var}\left(\beta_i | \xi_i^{(j)}, s^j\right).$$

Given that we approximate $m_{\delta \to \beta}$ using a Gaussian distribution and, under a Gaussian prior, it directly follows that $b(\beta)$ is Gaussian as well, the message $m_{\beta \to \delta}$ is also Gaussian with:

(A.14)
$$m_{\boldsymbol{\beta}\to\boldsymbol{\delta}}\sim \mathcal{N}(\boldsymbol{\beta}^*|\boldsymbol{\xi}^{*j},\boldsymbol{\zeta}^{*j}\boldsymbol{I}_K),$$

where

In the next step, we notice that because $m_{\beta \to \delta}$ is Gaussian and the likelihood based on the copy of β is Gaussian too, the belief $b(\beta^*)$ is also approximated using a Gaussian distribution with a covariance matrix that is proportional to the identity matrix (Rangan et al., 2019; Van den Boom et al., 2021). The corresponding approximating density is

(A.15)
$$b(\boldsymbol{\beta}^*) \approx \mathcal{N}\left(\boldsymbol{\beta}^* | \overline{\boldsymbol{\beta}}^{*(j)}, s^{*(j)} \boldsymbol{I}_K\right).$$

Simply computing the moments of this distribution requires inverting a $K \times K$ matrix. However, the formula for these moments can be written in a different way so as to avoid inverting such a matrix by decomposing X using the singular value decomposition (SVD) of X = UDV'. The corresponding mean and variance are then given by

$$\overline{\boldsymbol{\beta}}^{*(j)} = \xi^{*(j)} + \boldsymbol{V}(\sigma^2 \boldsymbol{D}^{-1} / \zeta^{*(j)} + \boldsymbol{D})^{-1} (\boldsymbol{U}' \boldsymbol{y} - \boldsymbol{D} \boldsymbol{V}' \xi^{*(j)}),$$

$$s^{*(j)} = \zeta^{*(j)} (1 - \operatorname{Tr}(\boldsymbol{V}(\sigma^2 \boldsymbol{D}^{-1} / \zeta^{*(j)} + \boldsymbol{D})^{-1}) / K).$$

These expressions require only the inversion of diagonal matrices and are thus computationally cheap to implement.

Finally, we set

(A.16)
$$1/\zeta^{(j+1)} = 1/s^{*(j)} - 1/\zeta^{*(j)},$$

(A.17)
$$\xi^{(j+1)} = \frac{\zeta^{*(j)} \overline{\beta}^{*(j)} - s^{*(j)} \xi}{\zeta^{*(j)} - s^{*(j)}}.$$

Repeating steps (A.13)–(A.17) *S* times yields the approximate posterior which is $\mathcal{N}(\overline{\beta}^{(S)}, s^{(S)}I_K)$. In practice, we stop iterating the algorithm if $||\overline{\beta}^{(j)} - \overline{\beta}^{(j-1)}||^2$ is smaller than a certain threshold close to zero. This algorithm shows good convergence properties, is computationally fast and has been shown to work well empirically.

A.1.1 Updating the error variance. Conditional on the data and the current state of the parameter vector $\boldsymbol{\beta}$, let $\boldsymbol{\epsilon} = \boldsymbol{y} - \boldsymbol{X}\boldsymbol{\beta}$. Under the prior $\sigma^2 \sim \mathcal{G}^{-1}(a_{\sigma}, b_{\sigma})$, the error variance update is given by

$$\sigma^2 = 2b_\sigma + \frac{\epsilon'\epsilon}{2a_\sigma + T}$$

A.1.2 Updating the Horseshoe prior. Based on the general Horseshoe prior described in Subsection 2.2, the updating quantities for the hierarchical prior structure in the VAMP algorithm for a parameter vector $\boldsymbol{\phi} = (\phi_1, \dots, \phi_k)'$ are given by

$$\psi_i^{-2} = rac{1}{
u_i^{-1} + \phi_i^2 \lambda^{-2}/2}, \quad \lambda^{-2} = rac{k+1}{2\xi^{-1} + \psi_i^{-2} \sum_{i=1}^k \phi_i^2}$$
 $u_i = rac{1}{1 + \psi_i^{-2}}, \quad \xi = rac{1}{1 + \lambda^{-2}}.$

Note that in the empirical application, we have country- and equation-specific global and local shrinkage parameters.

A.2 Inference for panel VAR integrated rotated Gaussian approximation (PVAR-IRGA).

A.2.1 Markov chain Monte Carlo (MCMC) sampling of the VAR coefficients. Let $\tilde{\mathbf{y}}_{ij} = \mathbf{Q}'_{i1}(\mathbf{y}_{ij} - \tilde{\mathbf{z}}_{ij}\overline{\mathbf{B}}_{ij,\bullet}), \tilde{\mathbf{x}}_{ij} = (\tilde{\mathbf{Q}}'_{i1}\mathbf{x}_{i})$, and $\boldsymbol{\Sigma}_{ij} = \mathbf{Q}'_{i1}\tilde{\mathbf{z}}_{ij}\overline{\mathbf{V}}_{ij,\bullet}\tilde{\mathbf{z}}'_{ij}\mathbf{Q}_{i1} + \sigma_{\varepsilon,ij}^{2}\mathbf{I}_{k}$, and recall that the likelihood of the approximate model is

$$\boldsymbol{Q}_{i1}^{\prime}(\boldsymbol{y}_{ij}-\tilde{\boldsymbol{z}}_{ij}\overline{\boldsymbol{B}}_{ij,\bullet})\sim\mathcal{N}\Big(\boldsymbol{Q}_{i1}^{\prime}\boldsymbol{x}_{i}\boldsymbol{A}_{ij,\bullet},\boldsymbol{Q}_{i1}^{\prime}\tilde{\boldsymbol{z}}_{ij}\overline{\boldsymbol{V}}_{ij,\bullet}\tilde{\boldsymbol{z}}_{ij}^{\prime}\boldsymbol{Q}_{i1}+\sigma_{\varepsilon,ij}^{2}\boldsymbol{I}_{k}\Big).$$

The conditional posterior distribution of the VAR coefficients collected in $A_{ij,\bullet}$ is

$$\begin{split} \boldsymbol{A}_{ij,\bullet} &\sim \mathcal{N}(\bar{\boldsymbol{\mu}}_{ij}, \bar{\boldsymbol{S}}_{ij}), \\ \bar{\boldsymbol{S}}_{ij} &= \left(\boldsymbol{V}_{A,ij}^{-1} + \tilde{\boldsymbol{x}}_{ij}' \boldsymbol{\Sigma}_{ij}^{-1} \tilde{\boldsymbol{x}}_{ij}\right)^{-1}, \\ \bar{\boldsymbol{\mu}}_{ij} &= \bar{\boldsymbol{S}}_{ij} \left(\boldsymbol{V}_{A,ij}^{-1} \boldsymbol{\mu}_{A,ij} + \tilde{\boldsymbol{x}}_{ij}' \boldsymbol{\Sigma}_{ij}^{-1} \tilde{\boldsymbol{y}}_{ij}\right), \end{split}$$

where $\mu_{A,ij}$ is the prior mean and $V_{A,ij}$ the prior covariance matrix. For our empirical application, the prior mean is set to a zero vector, $\mu_{A,ij} = \mathbf{0}_k$ for all equations and countries, and the prior variance is defined by the Horseshoe prior: $V_{A,ij} = (\lambda_{A,ij})^2 \times \text{diag}(\psi_{ij,1}^2, \dots, \psi_{ij,k}^2)$.

A.2.2 Minnesota prior for PVAR-IRGA. The Minnesota prior for PVAR-IRGA, in the spirit of Chan (2021), is constructed as follows: The global scaling parameter for equation $i, \lambda_i \sim C^+(0, 1)$, follows a half-Cauchy distribution as in the case of the Horseshoe. This parameter governs the overall tightness of the Minnesota prior on an equation-by-equation basis. Let $\hat{\sigma}_i^2$ denote the residual variance of a univariate AR(p) model for the *i*th variable. The local scales ψ_{ij} are set deterministically in the Minnesota tradition, by introducing a quadratic lag-penalty scaled by the ratio $\hat{\sigma}_i^2/\hat{\sigma}_j^2$ for the shrinkage parameter related to variable *j*.

A.2.3 *Posteriors related to the Horseshoe prior.* We present the conditional posterior distributions of the Horseshoe prior for the general example in Subsection 2.2 to economize on notation. Note that for the PVAR, we have country- and equation-specific global and local shrinkage parameters.

Our implementation is based on the auxiliary representation of the Horseshoe prior discussed in Makalic and Schmidt (2016). For a k-dimensional parameter vector $\boldsymbol{\phi} = (\phi_1, \dots, \phi_k)'$, with ϕ_j indexing the $j = 1, \dots, k$ coefficient, we obtain the following inverse Gamma distributed posteriors:

$$\psi_j^2 |\phi_j, \lambda, \nu_j \sim \mathcal{G}^{-1}\left(1, \frac{1}{\nu_j} + \frac{\phi_j^2}{2\lambda^2}
ight), \quad \lambda^2 |\phi_j, \psi_j, \xi \sim \mathcal{G}^{-1}\left(\frac{k+1}{2}, \frac{1}{\xi} + \sum_{j=1}^k \frac{\phi_j^2}{2\psi_j^2}
ight),$$

 $u_j |\psi_j \sim \mathcal{G}^{-1}\left(1, 1 + \frac{1}{\psi_j^2}
ight), \quad \xi |\lambda \sim \mathcal{G}^{-1}\left(1, 1 + \frac{1}{\lambda^2}
ight).$

For the case of the Minnesota prior, only the auxiliary variables related to the global parameter need to be drawn.

A.3 Other Models.

A.3.1 Single-country VARs. The Bayesian vector autoregressions (BVARs) for single countries are given by

$$\mathbf{y}_{it} = \mathbf{\Gamma}_{i1}\mathbf{y}_{it-1} + \cdots + \mathbf{\Gamma}_{ip}\mathbf{y}_{it-p} + \boldsymbol{\epsilon}_{it}, \quad \boldsymbol{\epsilon}_{it} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{i}).$$

Note that this specification corresponds to Equation (1) with $\Xi_i = \mathbf{0}_{M \times K_{other}}$, ruling out both dynamic and static interdependencies across countries by assuming $\operatorname{cov}(\boldsymbol{\epsilon}_{it}, \boldsymbol{\epsilon}_{st}) = \mathbf{0}$. The models are estimated one country at a time.

We use a Horseshoe prior on the own and cross-variable VAR coefficients. Alternatively, we use a hierarchical Minnesota-type prior similar to the one discussed in Subsection A.2.2. The VARs are estimated equation-by-equation. The priors on the covariances are independent Gaussian distributions with mean zero and variance 10. The variances of the structural errors are assigned independent weakly informative inverse Gamma priors.

A.3.2 Factor-augmented vector autoregressions (FAVAR). Define a vector $\tilde{y}_{it} = (y'_{it}, f'_{it})'$, with y_{it} including the domestic variables of country *i* and f_{it} denoting a set of principal components extracted from the matrix $F_i = (y_{-i,1}, \ldots, y_{-i,T})'$. The domestic variables and the "foreign" factors f_{it} are modeled jointly in a VAR:

$$\tilde{\mathbf{y}}_{it} = \mathbf{\Gamma}_{i1}\tilde{\mathbf{y}}_{it-1} + \cdots + \mathbf{\Gamma}_{ip}\tilde{\mathbf{y}}_{it-p} + \tilde{\boldsymbol{\epsilon}}_{it}, \quad \tilde{\boldsymbol{\epsilon}}_{it} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{i}).$$

We use the same hierarchical Minnesota and Horseshoe priors and estimation algorithm as for the BVAR and estimate the models one country at a time.

A.3.3 Global vector autoregression (GVAR). The GVAR model incorporates crosscountry dependencies via including weighted averages of the foreign variables in the domestic equations. In particular, the GVAR is given by

$$\mathbf{y}_{it} = \mathbf{\Gamma}_{i1}\mathbf{y}_{it-1} + \dots + \mathbf{\Gamma}_{ip}\mathbf{y}_{it-p} + \mathbf{\Xi}_{i0}\mathbf{y}_{it}^* + \mathbf{\Xi}_{i1}\mathbf{y}_{it-1}^* + \dots + \mathbf{\Xi}_{iq}\mathbf{y}_{it-q}^* + \boldsymbol{\epsilon}_{it}, \quad \boldsymbol{\epsilon}_{it} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_i),$$

where the y_{it}^* are defined as

$$\mathbf{y}_{it}^* = \sum_{j=1}^N w_{ij} \mathbf{y}_{jt}.$$

APPROXIMATE INFERENCE IN PANEL VARS

		MAE				LPS			
	Model	Equities	Ind. Prod.	LT-IR	Infl.	Equities	Ind. Prod.	LT-IR	Infl.
h = 1									
	BVAR	1.003	0.992	0.955	0.965	-5.877	0.058	6.796	3.241
		(42.1)	(16.1)	(40.0)	(45.7)	(7.9)	(29.0)	(40.0)	(31.4)
	FAVAR-10	1.003	1.007	0.948	0.971	-4.905	-0.789	8.324	3.417
		(31.6)	(12.9)	(56.7)	(45.7)	(15.8)	(22.6)	(60.0)	(40.0)
	PVAR-IRGA	0.829	0.414	0.265	0.397	-178.797	-42.377	43.308	-18.394
		(23.7)	(54.8)	(3.3)	(8.6)	(63.2)	(38.7)	(0.0)	(28.6)
	GVAR	1.045	1.083	1.127	1.166	-11.543	-18.524	-49.750	-39.472
		(2.6)	(16.1)	(0.0)	(0.0)	(13.2)	(9.7)	(0.0)	(0.0)
h = 12									
	BVAR	1.000	1.033	0.940	1.010	-1.092	-11.992	5.480	8.696
		(42.1)	(6.5)	(40.0)	(28.6)	(15.8)	(16.1)	(26.7)	(45.7)
	FAVAR-10	0.997	1.034	0.923	1.008	-0.220	-16.291	9.441	5.539
		(31.6)	(9.7)	(43.3)	(20.0)	(10.5)	(12.9)	(56.7)	(17.1)
	PVAR-IRGA	0.856	0.700	0.548	0.772	-192.860	-158.364	-95.594	-184.650
		(26.3)	(45.2)	(16.7)	(42.9)	(68.4)	(48.4)	(16.7)	(31.4)
	GVAR	1.071	1.018	1.021	1.115	-13.015	-4.675	-25.273	-20.469
		(0.0)	(38.7)	(0.0)	(8.6)	(5.3)	(22.6)	(0.0)	(5.7)

TABLE A.1 SUMMARY OF FORECAST EXERCISE, ALL MODELS WITH MINNESOTA PRIOR

Notes: GDP-weighted average over countries, win percentage across countries in parentheses. Root mean squared error (RMSE) and log predictive score (LPS) relative to the benchmark. The benchmark PVAR-IRGA shows actual values, all other models are in ratios to the benchmark for RMSEs and in differences for LPSs.

TABLE A.2

SUMMARY OF FORECAST EXERCISE, COMPARISON OF PVAR-IRGA ESTIMATED WITH HORSESHOE (HS), AND MINNESOTA (MN) PRIOR

	MAE				LPS				
	Model	Equities	Ind. Prod.	LT-IR	Infl.	Equities	Ind. Prod.	LT-IR	Infl.
h = 1									
	PVAR-IRGA-HS	0.825	0.410	0.274	0.386	-177.799	-39.815	36.181	-13.387
		(63.2)	(51.6)	(46.7)	(77.1)	(39.5)	(61.3)	(46.7)	(77.1)
	PVAR-IRGA-MN	1.005	1.009	0.967	1.029	-0.998	-2.562	7.127	-5.007
		(36.8)	(48.4)	(53.3)	(22.9)	(60.5)	(38.7)	(53.3)	(22.9)
h = 12									
	PVAR-IRGA-HS	0.875	0.664	0.570	0.759	-204.541	-149.824	-104.962	-174.713
		(47.4)	(61.3)	(63.3)	(65.7)	(36.8)	(77.4)	(43.3)	(74.3)
	PVAR-IRGA-MN	0.979	1.054	0.962	1.017	11.681	-8.540	9.368	-9.938
		(52.6)	(38.7)	(36.7)	(34.3)	(63.2)	(22.6)	(56.7)	(25.7)

Notes: GDP-weighted average over countries, win percentage across countries in parentheses. Root mean squared error (RMSE) and log predictive score (LPS) relative to the benchmark. The benchmark PVAR-IRGA-HS shows actual values, PVAR-IRGA-MN is in ratios to the benchmark for RMSEs and in differences for LPSs.

The prespecified set of weights satisfy the restrictions $\sum_{j=1}^{N} w_{ij} = 1$ and $w_{ii} = 0$. We estimate the GVAR using the R-package by Böck et al. (2020), using again a hierarchical Minnesota and a Horseshoe prior similar to the BVAR and the FAVAR.

A.4 Additional Empirical Results.

A.4.1 *Minnesota prior.* To provide empirical evidence on the role of the Horseshoe prior that we choose to impose shrinkage in the baseline specification of PVAR-IRGA, we re-do the forecast exercise with all models estimated using a Minnesota-type prior. Table A.1 shows the forecast results for this set of models relative to PVAR-IRGA estimated with a Minnesota prior. Table A.2 provides a direct comparison between the two competing priors for PVAR-IRGA. On average, the performance of PVAR-IRGA-HS is stronger than PVAR-IRGA-MN.



FIGURE A.2

CORRELATION OF SPILLOVER INDEX ESTIMATES ACROSS PERMUTATIONS OF COUNTRY BLOCKS



Maximum horizon — 12 — 24

NOTES: This index indicates the share of spillovers across countries excluding spillovers to variables within a given country (summed across variable types). Estimated based on rolling windows of varying length. The thicker line is the posterior median alongside the 68% posterior credible set (thin lines).

FIGURE A.3

DIEBOLD–YILMAZ SPILLOVER INDEX FOR ALL COUNTRIES BASED ON DECOMPOSITION OF FORECAST ERROR VARIANCE OF THE FULL SYSTEM USING A ROLLING WINDOW OF OBSERVATIONS



Notes: This index indicates the share of spillovers between countries one variable at a time (summed across countries). Estimated based on rolling windows of varying length. The thicker line is the posterior median alongside the 68% posterior credible set (thin lines).

FIGURE A.4

DIEBOLD–YILMAZ SPILLOVER INDEX BY VARIABLE BASED ON DECOMPOSITION OF FORECAST ERROR VARIANCE OF THE FULL SYSTEM USING A ROLLING WINDOW OF OBSERVATIONS

However, it is worth mentioning that PVAR-IRGA-MN produces modest improvements in our aggregate measures of predictive accuracy for equity prices and long-term interest rates (LT-IR) with respect to both point and density forecasts.

A.4.2 Ordering of variables. Our implementation of the equation-by-equation estimation algorithm of the PVAR implies that the posterior estimates are not invariant with respect to different orderings of the variables.

To shed light on the robustness of our approach in this context, we assess the sensitivity of the posterior median estimates of the spillover index to different orderings of country blocks (defined by the four categories Advanced European, Emerging European, Advanced Other, and Emerging Other) in the system. We choose this approach for two reasons. First, it is not possible to compute all possible permutations of individual variables. Consequently, we compute the same metric, the spillover index, for all possible permutations of the four country blocks. Second, it is infeasible to show robustness with respect to estimates of all parameters of the model. Thus, we rely on the spillover index as a relevant summary statistic for this purpose.

Figure A.2 shows a pairwise correlation matrix of the posterior estimates of the spillover index across all possible permutations of country blocks. The correlations range from 0.97 to



Window — 60 months ···· 84 months Maximum horizon — 12 — 24

NOTES: This index indicates the share of foreign spillovers to the indicated country (summed across variables). Estimated based on rolling windows of varying length. The lines are the posterior median estimates.

FIGURE A.5

DIEBOLD–YILMAZ SPILLOVER INDEX BY COUNTRY BASED ON DECOMPOSITION OF FORECAST ERROR VARIANCE OF THE FULL SYSTEM USING A ROLLING WINDOW OF OBSERVATIONS

close to one, suggesting almost identical estimates of the spillover indices across orderings. Hence, we conclude that the ordering of the variables in our model only has negligible effects on the results.

A.4.3 Spillover index with rolling window. Figures A.3 to A.5 show the DY indices estimated using rolling windows of varying length instead of an expanding window. In Figure A.3, we show the results of the overall DY index for different rolling window sizes (60 and 84 months) and for two horizons (12 and 24 months). Two facts that are in line with the literature on using the DY emerge: First, we see a significant increase in spillovers during crisis periods. More specifically, the DY increases strongly during the period from 2009 and 2011. This period was characterized by the emergence of the global financial crisis and—for Europe—the sovereign debt crisis. A second peak emerges in 2013/2014, the period of the so-called taper tantrum, that is, the fear that the U.S. Federal Reserve Bank could wind down its asset purchase program, causing large swings in exchange rates and capital outflows from emerging economies. The third peak occurred around 2016/2017, the period when the Brexit referendum caused severe strains in the European Union and Donald Trump was elected president of the United States. Second, we see that, in general, spillover indices are larger for higher forecast horizons, which is a finding that is consistent with the analysis provided in Section 5.

In Figure A.4, we show variable-specific spillover indices. We see that for example, spillovers are much larger for equity prices compared to spillovers of inflation, output, and LT-IR. This finding is also in line with results provided in Section 5. Spillovers in equity prices increased strongly in the aftermath of the global financial crisis as well as from 2015 onward. In both periods, major central banks launched large-scale asset purchase programs which affected financial markets (see, e.g., D'Amico and King, 2013; Rogers et al., 2014). Consistent with our analysis from Section 5, we also find a peak in spillovers between long-term rates in 2012, which might be explained by the events related to the sovereign debt crisis in Europe. For completeness, we also show country-specific spillover indices in Figure A.5.

REFERENCES

- BAI, Y., A. CARRIERO, T. E. CLARK, AND M. MARCELLINO, "Macroeconomic Forecasting in a Multi-country Context," *Journal of Applied Econometrics*, forthcoming.
- BATCHELOR, R., AND P. DUA, "Improving Macro-Economic Forecasts: The Role of Consumer Confidence," International Journal of Forecasting 14 (1998), 71-81.
- BHATTACHARYA, A., D. PATI, N. S. PILLAI, AND D. B. DUNSON, "Dirichlet–Laplace Priors for Optimal Shrinkage," *Journal of the American Statistical Association* 110 (2015), 1479–90.
- BISHOP, C. M., Pattern Recognition and Machine Learning (Springer, New York 2006).
- BÖCK, M., M. FELDKIRCHER, AND F. HUBER, "BGVAR: Bayesian Global Vector Autoregressions with Shrinkage Priors in R," Globalization and Monetary Policy Institute Working Paper 395, 2020.
- BORIO, C. E. V., AND A. FILARDO, "Globalisation and Inflation: New Cross-Country Evidence on the Global Determinants of Domestic Inflation," BIS Working Papers 227, May 2007.
 CADONNA, A., S. FRÜHWIRTH-SCHNATTER, AND P. KNAUS, "Triple the Gamma—A Unifying Shrinkage
- CADONNA, A., S. FRÜHWIRTH-SCHNATTER, AND P. KNAUS, "Triple the Gamma—A Unifying Shrinkage Prior for Variance and Variable Selection in Sparse State Space and TVP Models," *Econometrics* 8 (2020), 1–36.
 - , and _____, "Estimating Multicountry VAR Models," *International Economic Review* 50 (2009), 929–59.

, "Panel Vector Autoregressive Models: A survey," Advances in Econometrics 32 (2016), 205–46.

- CARRIERO, A., J. CHAN, T. E. CLARK, AND M. MARCELLINO, "Corrigendum to: Large Bayesian Vector Autoregressions with Stochastic Volatility and Non-Conjugate Priors," Manuscript, 2021.
 T. E. CLARK, AND M. MARCELLINO, "Large Bayesian Vector Autoregressions with Stochastic
- Volatility and Non-Conjugate Priors," *Journal of Econometrics* 212 (2019), 137–54.
 CARVALHO, C., N. POLSON, AND J. SCOTT, "The Horseshoe Estimator for Sparse Signals," *Biometrika* 97
- (2010), 465–80.
- CHAN, J. C., "Minnesota-Type Adaptive Hierarchical Priors for Large Bayesian VARs," *International Journal of Forecasting* 37 (2021), 1212–26.
- CICCARELLI, M., AND B. MOJON, "Global Inflation," *Review of Economics and Statistics* 92 (August 2010), 524–35.
- CUARESMA, J. C., M. FELDKIRCHER, AND F. HUBER, "Forecasting with Global Vector Autoregressive Models: A Bayesian Approach," *Journal of Applied Econometrics* 31 (2016), 1371–91.
- D'AMICO, S., AND T. B. KING, "Flow and Stock Effects of Large-Scale Treasury Purchases: Evidence on the Importance of Local Supply," *Journal of Financial Economics* 108 (2013), 425–48.
- DEMIRER, M., F. X. DIEBOLD, L. LIU, AND K. YILMAZ, "Estimating Global Bank Network Connectedness," *Journal of Applied Econometrics* 33 (2018), 1–15.
- DIACONIS, P., AND D. FREEDMAN, "Asymptotics of Graphical Projection Pursuit," Annals of Statistics (1984), 793-815.
- DIEBOLD, F. X., AND K. YILMAZ, "Measuring Financial Asset Return and Volatility Spillovers, with Application to Global Equity Markets," *Economic Journal* 119 (2009), 158–71.
- DONOHO, D. L., A. MALEKI, AND A. MONTANARI, "Message-Passing Algorithms for Compressed Sensing," Proceedings of the National Academy of Sciences 106 (2009), 18914–919.
- DOVERN, J., M. FELDKIRCHER, AND F. HUBER, "Does Joint Modelling of the World Economy Pay Off? Evaluating Global Forecasts from a Bayesian GVAR," *Journal of Economic Dynamics and Control* 70 (2016), 86–100.
- FELDKIRCHER, M., AND F. HUBER, "The International Transmission of US Shocks—Evidence from Bayesian Global Vector Autoregressions," *European Economic Review* 81 (2016), 167–88.
- GRIFFIN, J. E., AND P. J. BROWN, "Inference with Normal-Gamma Prior Distributions in Regression Problems," Bayesian Analysis 5 (2010), 171–88.
- HOLLAND, P. W., K. B. LASKEY, AND S. LEINHARDT, "Stochastic Blockmodels: First Steps," Social Networks 5 (1983), 109–37.

HUBER, F., "Density Forecasting Using Bayesian Global Vector Autoregressions with Stochastic Volatility," *International Journal of Forecasting* 32 (2016), 818–37.

——, G. KOOP, AND L. ONORANTE, "Inducing Sparsity and Shrinkage in Time-Varying Parameter Models," *Journal of Business & Economic Statistics* 39 (2021), 669–83.

- KABUKÇUOĞLU, A., AND E. MARTÍNEZ-GARCÍA, "Inflation as a Global Phenomenon—Some Implications for Inflation Modeling and Forecasting," *Journal of Economic Dynamics and Control* 87 (2018), 46–73.
- KARRER, B., AND M. E. NEWMAN, "Stochastic Blockmodels and Community Structure in Networks," *Physical Review E* 83 (2011), 016107.
- KOOP, G., AND D. KOROBILIS, "Model Uncertainty in Panel Vector Autoregressive Model," *European Economic Review* 81 (2016), 115–31.

—, D. KOROBILIS, AND D. PETTENUZZO, "Bayesian Compressed Vector Autoregressions," *Journal of Econometrics* 210 (2019), 135–54.

KOROBILIS, D., "Prior Selection for Panel Vector Autoregressions," *Computational Statistics & Data Analysis* 101 (2016), 110–20.

——, High-Dimensional Macroeconomic Forecasting using Message Passing Algorithms," Journal of Business & Economic Statistics 39 (2021), 493–504.

- Kose, M. A., C. OTROK, AND C. H. WHITEMAN, "International Business Cycles: World, Region, and Country-Specific Factors," *American Economic Review* 93 (2003), 1216–39.
- LUDVIGSON, S., "Consumer Confidence and Consumer Spending," *Journal of Economic Perspectives* 18 (2004), 29–50.
- MAKALIC, E., AND D. SCHMIDT, "A Simple Sampler for the Horseshoe Estimator," IEEE Signal Processing Letters 23 (2016), 179–82.
- PARK, T., AND G. CASELLA, "The Bayesian Lasso," Journal of the American Statistical Association 103 (2008), 681–86.
- PATI, D., AND A. BHATTACHARYA, "Optimal Bayesian Estimation in Stochastic Block Models," *arXiv* preprint arXiv:1505.06794 (2015).

—, —— and V. SMITH, "Forecasting Economic and Financial Variables with Global VARs," International Journal of Forecasting 25 (2009), 642–75.

- ——, and S. WEINER, "Modeling Regional Interdependencies using a Global Error-Correcting Macroeconometric Model," *Journal of Business & Economic Statistics* 22 (2004), 129–62.
- PHAM, B. T., AND H. SALA, "Cross-Country Connectedness in Inflation and Unemployment: Measurement and Macroeconomic Consequences," *Empirical Economics* 62 (2022), 1123–46.
- RANGAN, S., P. SCHNITER, AND A. K. FLETCHER, "Vector Approximate Message Passing," *IEEE Transactions on Information Theory* 65 (2019), 6664–84.
- ROGERS, J. H., C. SCOTTI, AND J. H. WRIGHT, "Evaluating Asset-Market Effects of Unconventional Monetary Policy: A Cross-Country Comparison," International Finance Discussion Papers 1101, March 2014.
- VAN DEN, BOOM, W., G. REEVES, AND D. DUNSON, "Approximating Posteriors with High-Dimensional Nuisance Parameters via Integrated Rotated Gaussian Approximation," *Biometrika* 108 (2021), 269–82.
- ŽIBERNA, A., "Blockmodeling of Multilevel Networks," Social Networks 39 (2014), 46–61.