A Structural Matrix Autoregression Framework for International Spillovers

Ignacio Moreira Lara
University of Duisburg-Essen, 45141 Essen, Germany, and
TU Dortmund University, 44221 Dortmund, Germany
ignacio.moreira-lara@vwl.uni-due.de

Jan Prüser
EBS University of Business and Law, 65375 Oestrich-Winkel, Germany
Jan.Prueser@ebs.edu

and

Christoph Hanck
University of Duisburg-Essen, 45141 Essen, Germany
christoph.hanck@vwl.uni-due.de

Abstract

Understanding how macroeconomic shocks propagate across countries requires structural models that can jointly identify country-specific shocks and their international transmission. Yet, extending structural vector autoregressions (SVARs) to large multi-country systems is challenging in view of rapidly increasing dimensionality, computational costs, and the proliferation of identifying restrictions. This paper develops a Bayesian Structural Matrix Autoregression (BSMAR) framework that exploits the natural matrix structure of international macroeconomic data. By separating dependence across economic variables from dependence across countries, the framework provides a parsimonious representation that substantially reduces the dimensionality of large structural systems. We develop a Bayesian sampling algorithm for posterior inference that accommodates zero, sign, and ranking (magnitude) restrictions, allowing established SVAR identification schemes to be combined with a novel approach to identifying contemporaneous international spillovers. Applying the model to quarterly supply and demand data for 17 economies, we find evidence of spillovers between countries in both shocks, with the United States as an important transmitter.

Keywords: Structural vector autoregression; Identification; International spillover; High dimension; Bayesian inference

1 Introduction

Structural vector autoregressions (SVARs) have become a central tool for identifying macroeconomic shocks and studying their dynamic effects on the economy, see, e.g., Sims (1980), Blanchard & Quah (1989), and Kilian & Lütkepohl (2017). Recent applications of sign-restricted SVARs have, in particular, revisited the role of supply and demand disturbances in explaining the post-pandemic inflation episode, see, e.g., Giannone & Primiceri (2025), Bergholt et al. (2026), and Chang et al. (2023). Yet, despite the inherently international nature of many macroeconomic questions, most structural VAR applications continue to study countries in isolation. This is increasingly difficult to reconcile with the structure of modern economies, where deep trade integration, globally interconnected financial markets and complex production networks allow shocks originating in one country to propagate rapidly across borders. Understanding these spillovers requires structural models that can simultaneously identify country-specific shocks and characterize their transmission across economies.

Extending structural VARs to a large panel of countries, however, presents a fundamental challenge. As the number of countries and variables increases, the dimension of the system and the number of unknown parameters grow rapidly, resulting in severe overparameterization, high computational costs, and a proliferation of identifying restrictions. Consequently, unrestricted multi-country SVARs quickly become infeasible, even for a moderate number of economies given the size of typical balanced macroeconomic datasets.

To address these challenges this paper develops a Bayesian Structural Matrix Autoregression (BSMAR) approach, together with an efficient sampling algorithm for drawing from its joint posterior distribution. The key insight is that international macroeconomic data naturally have a matrix structure, with one dimension representing economic variables and the other representing countries. Rather than vectorizing the data into a single high-

dimensional system, we exploit this two-way structure directly. The resulting framework separates dependence across economic variables from dependence across countries, yielding a parsimonious representation that preserves the multivariate structure of the data while substantially reducing the dimensionality of the model. See the seminal work of Chen et al. (2021) for the introduction of this class of models.

The matrix formulation resolves the three principal obstacles that have prevented the estimation of large structural international systems. First, it drastically reduces the number of unknown parameters relative to an unrestricted SVAR, mitigating overparameterization and improving statistical efficiency. Second, the resulting parsimonious representation makes Bayesian estimation computationally feasible even in systems containing many countries. Third, and most importantly, it fundamentally simplifies the structural identification problem. Under our framework, the contemporaneous impact matrix admits a Kronecker product decomposition into two variable and country specific components. Rather than identifying one prohibitively large structural impact matrix, identification is carried out separately for economic relationships among variables and for contemporaneous spillovers across countries. This decomposition substantially reduces the number of structural parameters and identifying restrictions while preserving a transparent economic interpretation.

Our Bayesian framework further allows researchers to incorporate a broad class of identifying restrictions, including zero, sign, ranking, and more general nonlinear restrictions. This flexibility makes it possible to combine established SVAR identification schemes for the variable-specific component[1] with a novel identification framework for the country-specific component. In particular, we use economically motivated restrictions on the country-specific impact matrix to identify structural spillovers across countries. More specifically, a central assumption is that a country-specific shock has its largest contemporaneous impact in the

[1]Some examples are sign restriction identification of supply and demand shocks.

country where it originates. We also allow for soft restrictions that incorporate prior beliefs about the direction and relative magnitude of spillovers without ruling out deviations from these beliefs.

We demonstrate the empirical relevance of the framework using quarterly data for 17 economies. Motivated by the recent literature on the role of supply and demand shocks in explaining inflation dynamics (e.g., Giannone & Primiceri (2025), Bergholt et al. (2026), Chang et al. (2023)) we simultaneously identify 34 country-specific structural shocks, two supply and demand shocks for each economy, and trace the spillovers into other countries. To the best of our knowledge, this is the first study to jointly identify such a large number of country-specific structural shocks. We find substantial evidence of spillovers between countries, with the US acting as an important transmitter. Furthermore, we find that demand shocks emerge as the main driver of the increase in inflation following the COVID-19 pandemic. Lastly, we provide a comparison of the matrix framework with the state-of-the-art Bayesian SVAR with Minnesota shrinkage designed to handle a large number of restrictions (Arias et al. 2025).

Our paper contributes to two strands of literature. First, we contribute to the rapidly growing literature on matrix-valued time-series models. Matrix autoregressive (MAR) models exploit the two-way structure of matrix-valued data to achieve substantial dimension reduction relative to vectorized VARs. Recent developments include Bayesian estimation, variable selection, nonlinear models, and factor structures, among others, see Celani & Pagnottoni (2024), Celani et al. (2024), Bucci (2025), Chan & Zhang (2024), Kong et al. (2026) and Ye et al. (2026). Our contribution in this regard is to develop a structural matrix autoregression that exploits this two-way structure to make high-dimensional structural analysis tractable while allowing researchers to identify contemporaneous structural relationships within countries as well as structural spillovers across countries.

Second, we contribute to the literature on international macroeconomic spillovers. Existing approaches address the curse of dimensionality by imposing structure on cross-country dependence. For example, dynamic factor models summarize large datasets using a small number of latent factors and are highly effective for studying common business-cycle fluctuations, see, e.g., Forni & Lippi (2001), Charnavoki & Dolado (2014), Mumtaz & Surico (2009), and Stock & Watson (2005). However, they are not designed to identify a large number of country-specific structural shocks and typically focus on a relatively small number of shocks, see, e.g., Charnavoki & Dolado (2014). Global vector autoregressive (GVAR) models provide another influential framework for studying international macroeconomic linkages, see, e.g., Pesaran et al. (2004), Dees et al. (2007), Crespo Cuaresma et al. (2019), and Feldkircher & Huber (2016). By connecting country-specific models through predetermined foreign aggregates, GVARs remain tractable for large panels, but at the cost of imposing substantial structure on contemporaneous cross-country interactions. More generally, the existing literature faces a fundamental trade-off: approaches that remain tractable for large panels necessarily impose restrictions on the cross-country dependence structure (Camehl & von Schweinitz 2026, Pfarrhofer & Stelzer 2025), limiting the identification of country-specific structural spillovers, whereas unrestricted structural models become infeasible as the number of countries and variables increases. Our approach addresses this trade-off by combining the dimensionality reduction of matrix autoregressions with structural identification.

The remainder of the paper is organized as follows. Section 2 introduces the structural matrix framework, discusses general identification restrictions, and develops the Bayesian sampling algorithm. Section 3 describes the data, presents our empirical identification strategy, and discusses the main empirical findings. The Online Appendix, Section E, compares our model with a benchmark from the literature. Section 4 concludes.

2 A Large Multi-Country Framework

In this section, we first present the reduced-form Bayesian Matrix Autoregressive (MAR) model. We then formulate its structural representation and discuss the identification of the structural parameters. Finally, we develop a sampling algorithm for drawing from the joint posterior distribution of the elements in the model.

2.1 The Reduced Form MAR

Let $\mathbf{Y}_t$ denote an $n \times k$ matrix of endogenous time series, where $i = 1, \ldots, n$ indexes the macroeconomic variables for each country and $j = 1, \ldots, k$ indexes the countries, for $t = 1, \ldots, T$. That is, the rows of $\mathbf{Y}_t$ correspond to economic variables, while the columns correspond to countries. Hence, the matrix structure reflects the two dimensions that characterize international macroeconomic data: interactions among economic variables and interactions across countries. In total, there are $r = nk$ time series.

The Matrix Autoregressive model, as presented in Chan & Qi (2026), Samadi & Billard (2025), Chen et al. (2021), exploits this structure and is given by

$$\mathbf{Y}_t = \mathbf{A}_1 \mathbf{Y}_{t-1} \mathbf{B}_1' + \ldots + \mathbf{A}_p \mathbf{Y}_{t-p} \mathbf{B}_p' + \mathbf{U}_t, \tag{1}$$

where $\mathbf{A}_1, \ldots, \mathbf{A}_p$ are $n \times n$ coefficient matrices and $\mathbf{B}_1, \ldots, \mathbf{B}_p$ are $k \times k$ coefficient matrices. The matrices $\mathbf{A}_\ell$ capture dynamic interactions among economic variables, whereas $\mathbf{B}_\ell$ capture dynamic interactions across countries. The model therefore allows for a rich but parsimonious representation of the dynamics of the data, making it suitable for high-dimensional settings where the number of variables and countries is jointly large relative to the number of observations.

The error term is assumed to follow a matrix normal distribution,

$$\mathbf{U}_t \sim \mathcal{MN}_{n \times k}(\mathbf{0}, \mathbf{\Sigma}_c, \mathbf{\Sigma}_r), \tag{2}$$

where $\mathbf{\Sigma}_c$ is a $k \times k$ positive definite column covariance matrix and $\mathbf{\Sigma}_r$ is an $n \times n$ positive definite row covariance matrix.[2] For a more compact representation, define $\mathbf{X}_t = \text{blkdiag}(\mathbf{Y}_{t-1}, \ldots, \mathbf{Y}_{t-p})$, $\mathbf{A} = (\mathbf{A}_1, \ldots, \mathbf{A}_p)'$ and $\mathbf{B} = (\mathbf{B}_1, \ldots, \mathbf{B}_p)'$ denote the stacked autoregressive coefficient matrices. The MAR model can then be written as

$$\mathbf{Y}_t = \mathbf{A}'\mathbf{X}_t\mathbf{B} + \mathbf{U}_t.$$

To illustrate the MAR form, consider the special case with $p = 1$ in which the variable-side matrix is the identity, $\mathbf{A}_1 = \mathbf{I}_n$. In a panel of GDP and CPI series for the United States, Canada and Japan, the block of US equations is given by

$$\overset{\text{USA}}{\begin{pmatrix} \text{GDP} \\ \text{CPI} \end{pmatrix}_t} = b_1^{11} \overset{\text{USA}}{\begin{pmatrix} \text{GDP} \\ \text{CPI} \end{pmatrix}_{t-1}} + b_1^{12} \overset{\text{CAN}}{\begin{pmatrix} \text{GDP} \\ \text{CPI} \end{pmatrix}_{t-1}} + b_1^{13} \overset{\text{JPN}}{\begin{pmatrix} \text{GDP} \\ \text{CPI} \end{pmatrix}_{t-1}} + \overset{\text{USA}}{\begin{pmatrix} u^{\text{GDP}} \\ u^{\text{CPI}} \end{pmatrix}_t},$$

where $b_1^{j_1 j_2}$ denotes the (j_1, j_2) element of $\mathbf{B}_1$, and u^{GDP} and u^{CPI} are the reduced-form innovations of the US block. Each country block is thus a linear combination of the corresponding blocks of all countries, its own included, with b_1^{1j} measuring the dynamic propagation from country j to the United States, acting identically on GDP and CPI. When $\mathbf{A}_\ell \neq \mathbf{I}_n$ we have the interactions across variables within a country, since each lagged block is then premultiplied by $\mathbf{A}_\ell$ and GDP and CPI no longer respond proportionally.

A notable characteristic of the MAR model is the equivalent vector representation. Vectorizing yields

$$\mathbf{y}_t = \mathbf{\Phi}_1\mathbf{y}_{t-1} + \ldots + \mathbf{\Phi}_p\mathbf{y}_{t-p} + \mathbf{u}_t, \tag{3}$$

where $\mathbf{y}_t = \text{vec}(\mathbf{Y}_t)$, $\mathbf{u}_t = \text{vec}(\mathbf{U}_t)$, $\mathbf{u}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_u)$, $\mathbf{\Phi}_\ell = \mathbf{B}_\ell \otimes \mathbf{A}_\ell$, and $\mathbf{\Sigma}_u = \mathbf{\Sigma}_c \otimes \mathbf{\Sigma}_r$. Hence, the MAR model can be viewed as a VAR in which the large $r \times r$ coefficient matrices and

[2] Note that under the matrix normal distribution, $\text{Cov}(\text{vec}(\mathbf{U}_t)) = \mathbf{\Sigma}_c \otimes \mathbf{\Sigma}_r$. Thus, the error structure mirrors the two-dimensional organization of the data by separately capturing dependence across countries and across economic variables.

covariance matrix are restricted to have a Kronecker product structure. Besides providing a familiar framework, the VAR representation of the MAR allows us to summarize our findings with the standard SVAR tools (Kilian & Lütkepohl 2017): structural impulse response functions (SIRFs), forecast error variance decompositions (FEVDs) and historical decompositions (HDs). The causal interpretation of these measures depends on the stability condition of the MAR process[3]. Other interpretations of the MAR model as a GVAR and factor model are presented in Chan & Qi (2026).

To illustrate the efficiency gains in our model, consider that an unrestricted VAR with nk variables and p lags requires $(nk)^2p$ autoregressive parameters. In contrast, the MAR representation requires only $(n^2 + k^2)p$ parameters. For example, with $n = 2$ economic variables and $k = 17$ countries, an unrestricted VAR requires $1156p$ autoregressive coefficients, compared with only $293p$ in the MAR model. Of course, the reduction becomes increasingly important as the number of countries or variables grows. The MAR model therefore provides a parsimonious yet economically plausible representation that is particularly well suited to high-dimensional settings where estimating an unrestricted VAR would be computationally demanding and statistically inefficient (Chen et al. 2021).

2.2 Structural Matrix Autoregression

To recover economically meaningful shocks, the reduced-form innovations must be mapped into the vector of structural shocks $\mathbf{e}_t$:

$$\mathbf{u}_t = \mathbf{B}_0\mathbf{e}_t, \qquad \mathbf{e}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_r), \tag{4}$$

where $\mathbf{B}_0$ is the structural impact matrix. However, the covariance matrix of the reduced-form innovations only identifies $\mathbf{B}_0$ up to an orthogonal rotation. Specifically, $\mathbf{\Sigma}_u = \mathbf{B}_0\mathbf{B}_0'$, and any orthogonal matrix $\mathbf{Q} \in \mathcal{O}(r)$, for the orthogonal group $\mathcal{O}(r) = \{\mathbf{Q} \in \mathbb{R}^{r\times r} : \mathbf{Q}\mathbf{Q}' =$

[3] Samadi & Billard (2025) define stability as the lag polynomial having all roots outside of the unit circle.

$\mathbf{I}_r\}$, yields an observationally equivalent model

$$(\mathbf{B}_0\mathbf{Q}')(\mathbf{B}_0\mathbf{Q}')' = \mathbf{B}_0\mathbf{Q}'\mathbf{Q}\mathbf{B}_0' = \mathbf{B}_0\mathbf{B}_0'.$$

Therefore, an unrestricted SVAR requires sufficiently many economic restrictions to select a unique rotation among infinitely many observationally equivalent structural representations. This identification problem becomes particularly challenging in large systems. An unrestricted SVAR has r^2 structural impact parameters, the number of elements in $\mathbf{B}_0$. Point identification requires at least $r(r-1)/2$ identifying assumptions on $\mathbf{B}_0$. When n or k increases, the number of restrictions required quickly becomes impractical.

To overcome this limitation, we exploit the matrix structure of the VAR and propose a parsimonious representation of the structural impact matrix. Namely we propose the structural bridge:

$$\mathbf{U}_t = \mathbf{B}_r\,\mathbf{E}_t\,\mathbf{B}_c', \qquad \mathbf{E}_t \sim \mathcal{MN}_{n\times k}(\mathbf{0}, \mathbf{I}_k, \mathbf{I}_n), \tag{5}$$

where $\mathbf{E}_t$ collects the structural shocks in matrix form, so that $\mathbf{e}_t = \text{vec}(\mathbf{E}_t)$. As a by-product, we obtain the structural impact matrix as

$$\mathbf{B}_0 = \mathbf{B}_c \otimes \mathbf{B}_r.$$

Hence, this restriction provides two important advantages. First, it substantially reduces the number of unknown structural parameters. Instead of identifying $r^2 = n^2k^2$ elements of $\mathbf{B}_0$, the researcher only needs to identify the $n^2 + k^2 - 1$ free elements contained in $\mathbf{B}_r$ and $\mathbf{B}_c$, one element being fixed by the scale normalization discussed below. Second, the Kronecker decomposition provides a more transparent economic interpretation of the identifying restrictions. The matrix $\mathbf{B}_r$ captures contemporaneous relationships among variables within a unit and therefore allows researchers to exploit existing identification schemes from the single-country SVAR literature. This is essentially the same separability as for the dynamics, where a single set of $\mathbf{A}_\ell$ matrices governs the variable dynamics of every country and a

single set of $\mathbf{B}_\ell$ matrices the cross-country dynamics of every variable. In contrast, $\mathbf{B}_c$ captures cross-unit transmission mechanisms. For $j_1, j_2 \in \{1, \ldots, k\}$, we write $b_c^{j_1 j_2}$ for the (j_1, j_2) element of $\mathbf{B}_c$. This represents the relative strength of the contemporaneous spillover from country j_2 to country j_1, and b_c^{jj} is the own effect of country j. Consequently, restrictions on $\mathbf{B}_c$ can be directly interpreted as assumptions about economic linkages, such as international spillovers or network effects. Hence, the identification problem is transformed from selecting a rotation of a large unrestricted $r \times r$ matrix $\mathbf{B}_0$ to identifying two lower-dimensional matrices with economically meaningful interpretations. An additional implication of this decomposition is that the covariance structure can be separated into row and column components as $\mathbf{\Sigma}_r = \mathbf{B}_r \mathbf{B}_r'$ and $\mathbf{\Sigma}_c = \mathbf{B}_c \mathbf{B}_c'$. The price to pay for the gains in efficiency is the rigidity of the model when estimating own country dynamics and spillovers. All interactions are restricted to this bilinear form and the flexibility is constrained. Whether the Kronecker restrictions are warranted is an empirical question, which we address by formal Bayesian model comparison, see the Online Appendix, Section E.2.

An interesting direction for future research is the development of matrix models that relax some separability restrictions while preserving the computational and statistical advantages of the matrix representation. One promising approach is to approximate the single $(\mathbf{B}_\ell \otimes \mathbf{A}_\ell)$ with a sum of Kronecker products, thereby allowing for multiple channels of interaction between countries and indicators while still achieving some dimension reduction. Another possibility is to combine a Kronecker-structured component with an unrestricted component in the structural impact matrix $\mathbf{B}_0$. Such specifications would provide a flexible continuum between the fully separable Matrix SVAR model and an unrestricted SVAR. We believe these extensions could be particularly useful in applications involving a relatively small number of countries, where the additional flexibility can be estimated with reasonable precision. However, as the cross-sectional dimension grows, the number of parameters

increases rapidly, leading to a substantial rise in estimation uncertainty and diminishing the practical advantages of these more flexible specifications. We therefore do not apply these weaker restrictions in our application but present both extensions in the Online Appendix, Section F.

2.3 General Structural Identification Restrictions

Because the structural-form representation is not identified, additional identifying restrictions are required to narrow the identified set and recover the structural shocks. In this section, we present a general framework for imposing such restrictions within our sampling algorithm presented in Section 2.4. The particular restrictions adopted in the empirical application are described in Section 3.2.

2.3.1 Parameter Restrictions

Define $\boldsymbol{\theta} = (\text{vec}(\mathbf{B}_r)', \text{vec}(\mathbf{B}_c)')' \in \mathbb{R}^{d\times 1}$, with $d = n^2 + k^2$, which stacks all structural parameters of the model. Following the notation commonly used in the literature (Hou 2024), we impose linear restrictions on $\boldsymbol{\theta}$ by organizing them into three classes: *zero* restrictions, *sign* restrictions and *relative magnitude* restrictions. Each is a set of linear conditions on $\boldsymbol{\theta}$, obtained by stacking rows using an appropriate selection matrix in $\mathbb{R}^d$, whose i-th vector we write $\mathbf{q}_i$.

Zero Restrictions. The r_0 zero restrictions fix selected elements of $\boldsymbol{\theta}$ at zero. If $\theta_{i_1} = \cdots = \theta_{i_{r_0}} = 0$, stack the corresponding basis vectors into

$$\mathbf{R}_0 = \begin{bmatrix} \mathbf{q}_{i_1} & \dots & \mathbf{q}_{i_{r_0}} \end{bmatrix}', \qquad \mathbf{R}_0\,\boldsymbol{\theta} = \mathbf{0}.$$

This is the special case $\mathbf{d}_i = \mathbf{0}$ of the column-specific equality restriction $\mathbf{R}_i^E \mathbf{b}_i = \mathbf{d}_i$ considered in Hou (2024). Such exclusion restrictions rule out specific contemporaneous

relationships, for example by setting the contemporaneous effect of one unit on another to zero.

Sign Restrictions. Sign restrictions fix the direction of selected elements of $\boldsymbol{\theta}$. Let $\psi_1, \ldots, \psi_{r_s}$, each in $\{-1, 1\}$, be the signs required of $\theta_{i_1}, \ldots, \theta_{i_{r_s}}$, with $+1$ denoting a non-negative element and -1 a non-positive one. Stacking the corresponding signed basis vectors for the r_s sign restrictions gives

$$\mathbf{R}_s = \begin{bmatrix} \psi_1 \mathbf{q}_{i_1} & \ldots & \psi_{r_s} \mathbf{q}_{i_{r_s}} \end{bmatrix}', \qquad \mathbf{R}_s \, \boldsymbol{\theta} \geq \mathbf{0} \text{ elementwise.}$$

Typical applications include restricting the direction of a contemporaneous response to a structural shock and/or imposing a sign restriction on cross-unit spillover effects.

Relative Magnitude Restrictions. Similarly the r_Δ relative magnitude, or ranking, restrictions order two elements of $\boldsymbol{\theta}$ rather than bound either one. Consider the restrictions $\theta_{i_\delta} \geq \theta_{j_\delta}$ for $i_\delta, j_\delta \in \{1, \ldots, r_\Delta\}$[4], and stack the differences of the corresponding basis vectors into

$$\mathbf{R}_\Delta = \begin{bmatrix} \mathbf{q}_{i_1} - \mathbf{q}_{j_1} & \ldots & \mathbf{q}_{i_{r_\Delta}} - \mathbf{q}_{j_{r_\Delta}} \end{bmatrix}', \qquad \mathbf{R}_\Delta \, \boldsymbol{\theta} \geq \mathbf{0}.$$

This class carries the dominance restriction that we will discuss in Section 3.2, which requires a structural coefficient to be at least as large as the competing coefficients in absolute value,

$$b_c^{j_2 j_2} \geq |b_c^{j_1 j_2}|, \qquad \forall j_1 \neq j_2.$$

The condition is stated on the own effect including its sign and therefore also implies $b_c^{j_2 j_2} \geq 0$. Replacing the difference $\mathbf{q}_i - \mathbf{q}_j$ by the pair $\mathbf{q}_i$ and $-\mathbf{q}_i$, and the zero right-hand side by $\underline{\eta}_i$ and $-\overline{\eta}_i$, gives $\underline{\eta}_i \leq \theta_i \leq \overline{\eta}_i$.

Collecting the three classes of restrictions into

$$\mathcal{S} = \{\mathbf{R}_0, \mathbf{R}_s, \mathbf{R}_\Delta\},$$

[4]The same construction yields the two-sided linear inequality restrictions $\mathbf{l}_i < \mathbf{R}_i^I \mathbf{b}_i < \mathbf{u}_i$ of Hou (2024)

the identified set is given by (cf. Arias et al. 2018, Read & Zhu 2025)

$$\mathcal{D}(\mathcal{S}) := \left\{ \boldsymbol{\theta} \in \mathbb{R}^d : \mathbf{R}_0 \boldsymbol{\theta} = \mathbf{0},\ \mathbf{R}\boldsymbol{\theta} \geq \mathbf{0} \right\},$$
$$\mathbf{R} = \begin{bmatrix} \mathbf{R}_s' & \mathbf{R}_\Delta' \end{bmatrix}'.$$

This set is the B-model analogue (Kilian & Lütkepohl 2017), of the rotation-based identified set

$$\mathcal{Q}(\boldsymbol{\phi} \mid S) = \{ \mathbf{Q} \in \mathcal{O}(r) : S(\boldsymbol{\phi}, \mathbf{Q}) \geq \mathbf{0} \}$$

defined in Read & Zhu (2025). Because the imposed restrictions generally identify a set rather than a unique structural parameter vector, all elements of $\mathcal{D}(\mathcal{S})$ are observationally equivalent: they generate the same likelihood value and cannot be distinguished using the sample information alone.

2.3.2 Soft Parameter Restrictions

The inequality restrictions above are commonly imposed on the sampler through an accept–reject step (Rubio-Ramírez et al. 2010). An alternative is proposed by Baumeister & Hamilton (2015), who implement sign restrictions via prior information on the structural parameters $\boldsymbol{\theta}$, so that the identifying information is stated on economically meaningful objects (contemporaneous effects and spillovers) rather than on rotations. The conditional posterior is

$$p(\boldsymbol{\theta} \mid \mathbf{Y}, \mathbf{A}, \mathbf{B}) \propto \mathcal{L}(\mathbf{Y} \mid \boldsymbol{\theta}, \mathbf{A}, \mathbf{B})\, p(\boldsymbol{\theta}),$$

with $\mathcal{L}(\cdot)$ the matrix-(Gaussian) likelihood using the stacked observations in $\mathbf{Y}$, and $p(\boldsymbol{\theta})$ the prior. A *soft* sign restriction on θ_i is a prior that favors a sign but does not enforce a particular sign, e.g., a Gaussian $p(\theta_i) = \mathcal{N}(\theta_i;\, \mu_{\theta_i}, \sigma^2_{\theta_i})$ with $\mu_{\theta_i} > 0$ and $\sigma^2_{\theta_i}$ chosen so that $P(\theta_i > 0)$ is high (say 0.95) places most, but not all, of the prior mass on the restricted region. If the likelihood contradicts the restriction sharply enough, the posterior still assigns

probability to sign violations, so identification operates through probabilistic weighting of structural models rather than through deterministic exclusion.

2.3.3 Normalization Restrictions

It is well known that a Kronecker product is unique only up to multiplication by a non-zero scalar γ_ℓ. The likelihood function is unchanged if, for $\ell = 1, \ldots, p$, the pair $(\mathbf{A}_\ell, \mathbf{B}_\ell)$ is replaced by $(\gamma_\ell \mathbf{A}_\ell, \gamma_\ell^{-1} \mathbf{B}_\ell)$, and likewise if $(\mathbf{B}_c, \mathbf{B}_r)$ is replaced by $(\gamma_\ell \mathbf{B}_c, \gamma_\ell^{-1} \mathbf{B}_r)$. We follow a similar approach as Chan & Qi (2026) and remove this indeterminacy by fixing the $(1,1)$ element of $\mathbf{B}_\ell$ at one for $\ell = 1, \ldots, p$ and the $(1,1)$ element of $\mathbf{B}_c$ at one. This fixes the scale and determines how the prior mean of these elements is set.

2.4 Bayesian Estimation

In this section we develop the Gibbs sampler to draw from the joint posterior distribution of the elements in the model. First, we introduce the prior distributions. Then, we present the algorithms to sample from the posterior distributions.

2.4.1 Prior Specification

To simplify computations we follow Chan & Qi (2026) and assume a natural conjugate prior onlPl $\mathbf{A}$ and $\mathbf{B}$. For the unrestricted elements in $\mathbf{B}_r$ and $\mathbf{B}_c$ we assume independent Gaussian distributions.

Prior Information on A We assume a Gaussian prior on $\mathbf{A}$, $p(\text{vec}(\mathbf{A}) \mid \mathbf{\Sigma}_r) = \mathcal{N}\big(\text{vec}(\underline{\mathbf{A}}), \mathbf{\Sigma}_r \otimes \mathbf{V}_\mathbf{A}\big)$ with density function:

$$p(\mathbf{A}) \propto |\mathbf{V}_\mathbf{A}|^{-\frac{n}{2}} |\mathbf{\Sigma}_r|^{-\frac{np}{2}} \, \mathrm{e}^{-\frac{1}{2}\,\mathrm{tr}\left(\mathbf{\Sigma}_r^{-1}(\mathbf{A}-\underline{\mathbf{A}})'\mathbf{V}_\mathbf{A}^{-1}(\mathbf{A}-\underline{\mathbf{A}})\right)}, \tag{6}$$

where $\underline{\mathbf{A}}$ is the prior mean, $\mathbf{V}_\mathbf{A}$ is the $np \times np$ prior covariance matrix, and $\mathrm{tr}(\cdot)$ denotes the trace operator. We specify both to encode Minnesota-type beliefs (Litterman 1986)

adapted to the matrix setting by Chan & Qi (2026). The choice of the prior mean depends on the data specification. In our application, the time series in the data set are in yearly growth rates. Thus, we shrink towards white noise and set $\underline{\mathbf{A}} = \mathbf{0}$. The matrix $\mathbf{V}_{\mathbf{A}}$ is diagonal and is the object through which the shrinkage hyperparameter $\kappa_{\mathbf{A}}$ acts. Recall that $i \in \{1, \ldots, n\}$ indexes the rows of $\mathbf{Y}_t$, that is the economic variables, and $j \in \{1, \ldots, k\}$ indexes its columns, that is the countries. Let $m_A = (\ell - 1)n + i$ denote the row of $\mathbf{A}$ that carries the coefficients on row variable i at lag $\ell \in \{1, \ldots, p\}$. To fix the scale, let $\hat{s}^2_{i,j}$ be the residual variance of an AR(4) regression fitted to the single series $y_{t,ij}$, and let $\hat{s}^2_{i,\bullet} = \frac{1}{k}\sum_{j=1}^{k} \hat{s}^2_{i,j}$ denote its average across the k countries in row i. The corresponding diagonal element of $\mathbf{V}_{\mathbf{A}}$ then is

$$v_{\mathbf{A},m_A m_A} = \frac{\kappa_{\mathbf{A}}}{\ell^2 \, \hat{s}^2_{i,\bullet}}, \qquad m_A = (\ell - 1)n + i. \tag{7}$$

The prior variance thus factors into three pieces. The term ℓ^{-2} shrinks higher-order lags towards zero more aggressively, reflecting the belief that distant lags carry little information. The term $\hat{s}^{-2}_{i,\bullet}$ makes the prior invariant to the units in which each row of $\mathbf{Y}_t$ is measured. Finally, $\kappa_{\mathbf{A}}$ scales the whole matrix and therefore governs the *overall* shrinkage strength: as $\kappa_{\mathbf{A}} \to 0$ the prior collapses onto $\underline{\mathbf{A}}$, while a large $\kappa_{\mathbf{A}}$ leaves $\mathbf{A}$ essentially unrestricted. Equivalently, $\mathbf{V}_{\mathbf{A}} = \kappa_{\mathbf{A}} \mathbf{C}_{\mathbf{A}}$, where $\mathbf{C}_{\mathbf{A}}$ is a (fixed) diagonal matrix determined by the data and the lag structure alone.

Rather than fixing $\kappa_{\mathbf{A}}$ at some default value, we treat it as unknown and let the data determine its value, following Giannone et al. (2015). We place a hierarchical gamma prior $\kappa_{\mathbf{A}} \sim \mathcal{G}(c_{\mathbf{A},1}, c_{\mathbf{A},2})$ and set $c_{\mathbf{A},1} = c_{\mathbf{A},2} = 5$, a weakly informative choice centered at one.

Prior Information for B The prior on $\mathbf{B}$ is close in spirit, the density function given by:

$$p(\mathbf{B}) \propto |\mathbf{V}_{\mathbf{B}}|^{-\frac{k}{2}} |\mathbf{\Sigma}_c|^{-\frac{kp}{2}} \, \mathrm{e}^{-\frac{1}{2}\operatorname{tr}\left(\mathbf{\Sigma}_c^{-1}(\mathbf{B}-\underline{\mathbf{B}})' \mathbf{V}_{\mathbf{B}}^{-1}(\mathbf{B}-\underline{\mathbf{B}})\right)}. \tag{8}$$

We set the prior mean as $\underline{\mathbf{B}} = (\mathbf{I}_k, \ldots, \mathbf{I}_k)'$ to represent no spillovers and country-homogeneous beliefs on $\mathbf{B}$. The matrix $\mathbf{V_B}$ is built symmetrically to (7): let $m_B = (\ell - 1)k + j$ denote the row of $\mathbf{B}$ that carries the coefficients on column country j at lag ℓ, and let $\hat{s}^2_{\bullet,j} = n^{-1} \sum_{i=1}^{n} \hat{s}^2_{i,j}$ denote the average AR(4) residual variance down column j. We set

$$v_{\mathbf{B},m_B m_B} = \frac{\kappa_{\mathbf{B}}}{\ell^2 \, \hat{s}^2_{\bullet,j}}, \qquad m_B = (\ell - 1)k + j,$$

so that $\kappa_{\mathbf{B}}$ determines how informative the prior is: a small $\kappa_{\mathbf{B}}$ pins $\mathbf{B}_\ell$ at the identity, a large one results in an uninformative prior, and again $\mathbf{V_B} = \kappa_{\mathbf{B}} \mathbf{C_B}$ with $\mathbf{C_B}$ fixed and diagonal. We use a hierarchical gamma prior $\kappa_{\mathbf{B}} \sim \mathcal{G}(c_{\mathbf{B},1}, c_{\mathbf{B},2})$ with $c_{\mathbf{B},1} = c_{\mathbf{B},2} = 5$ and estimate it from the data.

Priors on $\mathbf{B}_r$ and $\mathbf{B}_c$ We specify *Gaussian* priors on the free elements of $\mathbf{B}_r$ and $\mathbf{B}_c$,

$$\theta_i \sim \mathcal{N}\left(\mu_{\theta_i}, \, \sigma^2_{\theta_i}\right), \qquad \theta_i \in \boldsymbol{\theta},$$

with $\boldsymbol{\mu}_\theta$ and the diagonal $\boldsymbol{\Sigma}_\theta = \mathrm{diag}(\sigma^2_{\theta_i})$ collecting the prior means and variances discussed in Section 2.3.2. The prior on the covariance blocks is therefore *induced* by the prior on the structural parameters rather than elicited separately, which is what allows the identifying restrictions to be stated on economically interpretable objects.

2.4.2 Conditional Posteriors

Sampling A and B.

Conditional on $\mathbf{B}_r$ and $\mathbf{B}_c$, the two covariance matrices are known, and the coefficient blocks retain their conjugate form. The conditional posteriors for $\mathbf{A}$ and $\mathbf{B}$ are

$$\mathrm{vec}(\mathbf{A}) \mid \mathbf{Y}, \boldsymbol{\Sigma}_r, \mathbf{B}, \boldsymbol{\Sigma}_c, \kappa_{\mathbf{A}}, \kappa_{\mathbf{B}} \sim \mathcal{N}\Big(\mathrm{vec}(\widehat{\mathbf{A}}), \boldsymbol{\Sigma}_r \otimes \mathbf{K}_{\mathbf{A}}^{-1} \Big), \tag{9}$$

with

$$\mathbf{K_A} = \mathbf{V}_{\mathbf{A}}^{-1} + \sum_{t=1}^{T} \mathbf{X}_t \mathbf{B} \boldsymbol{\Sigma}_c^{-1} \mathbf{B}' \mathbf{X}_t', \qquad \widehat{\mathbf{A}} = \mathbf{K}_{\mathbf{A}}^{-1} \left(\mathbf{V}_{\mathbf{A}}^{-1} \underline{\mathbf{A}} + \sum_{t=1}^{T} \mathbf{X}_t \mathbf{B} \boldsymbol{\Sigma}_c^{-1} \mathbf{Y}_t' \right).$$

Similarly, the conditional posterior of $\mathbf{B}$ is given by

$$\text{vec}(\mathbf{B}) \mid \mathbf{Y}, \boldsymbol{\Sigma}_c, \mathbf{A}, \boldsymbol{\Sigma}_r, \kappa_{\mathbf{A}}, \kappa_{\mathbf{B}} \sim \mathcal{N}\Big(\text{vec}(\widehat{\mathbf{B}}), \boldsymbol{\Sigma}_c \otimes \mathbf{K}_{\mathbf{B}}^{-1}\Big), \tag{10}$$

where

$$\mathbf{K}_{\mathbf{B}} = \mathbf{V}_{\mathbf{B}}^{-1} + \sum_{t=1}^{T} \mathbf{X}_t' \mathbf{A} \boldsymbol{\Sigma}_r^{-1} \mathbf{A}' \mathbf{X}_t, \qquad \widehat{\mathbf{B}} = \mathbf{K}_{\mathbf{B}}^{-1} \left(\mathbf{V}_{\mathbf{B}}^{-1} \underline{\mathbf{B}} + \sum_{t=1}^{T} \mathbf{X}_t' \mathbf{A} \boldsymbol{\Sigma}_r^{-1} \mathbf{Y}_t \right).$$

One of the key advantages of this approach is its computational efficiency. In particular, drawing the autoregressive component of the model costs $\mathcal{O}(n^3 p^3)$ for $\mathbf{A}$ and $\mathcal{O}(k^3 p^3)$ for $\mathbf{B}$, rather than the $\mathcal{O}\big((nk)^3 p^3\big)$ required by a corresponding unrestricted VAR; see Chan & Qi (2026) for a detailed discussion. For the normalization restrictions we implement the algorithm by Rue & Held (2005).

Sampling $\kappa_{\mathbf{A}}$ and $\kappa_{\mathbf{B}}$. Because the shrinkage parameters enter the prior variances multiplicatively through $\mathbf{V}_{\mathbf{A}} = \kappa_{\mathbf{A}} \mathbf{C}_{\mathbf{A}}$ and $\mathbf{V}_{\mathbf{B}} = \kappa_{\mathbf{B}} \mathbf{C}_{\mathbf{B}}$, combining the gamma priors with the coefficient priors yields generalized inverse Gaussian (GIG) conditional posteriors,

$$\kappa_{\mathbf{A}} \mid \mathbf{A}, \boldsymbol{\Sigma}_r \sim \mathcal{GIG}\Big(c_{\mathbf{A},1} - \tfrac{n^2 p}{2},\ 2c_{\mathbf{A},2},\ \text{tr}\Big(\boldsymbol{\Sigma}_r^{-1} (\mathbf{A} - \underline{\mathbf{A}})' \mathbf{C}_{\mathbf{A}}^{-1} (\mathbf{A} - \underline{\mathbf{A}})\Big)\Big),$$
$$\kappa_{\mathbf{B}} \mid \mathbf{B}, \boldsymbol{\Sigma}_c \sim \mathcal{GIG}\Big(c_{\mathbf{B},1} - \tfrac{k^2 p}{2},\ 2c_{\mathbf{B},2},\ \text{tr}\Big(\boldsymbol{\Sigma}_c^{-1} (\mathbf{B} - \underline{\mathbf{B}})' \mathbf{C}_{\mathbf{B}}^{-1} (\mathbf{B} - \underline{\mathbf{B}})\Big)\Big).$$

To draw from this density we use the algorithm by Devroye (2014). The covariance matrices are then rebuilt from the new draws.

Sampling $\mathbf{B}_r$ and $\mathbf{B}_c$. The structural factors do not inherit the conjugacy of the coefficient blocks. Conditional on $(\mathbf{A}, \mathbf{B})$ the reduced-form innovations are known,

$$\mathbf{U}_t(\mathbf{A}, \mathbf{B}) \;=\; \mathbf{Y}_t - \sum_{\ell=1}^{p} \mathbf{A}_\ell \mathbf{Y}_{t-\ell} \mathbf{B}_\ell', \qquad \mathbf{u}_t \;=\; \text{vec}\left(\mathbf{U}_t\right) \;=\; \mathbf{B}_0\, \mathbf{e}_t, \qquad \mathbf{B}_0 = \mathbf{B}_c \otimes \mathbf{B}_r,$$

with $\mathbf{e}_t = \text{vec}(\mathbf{E}_t) \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_r)$. Gathering the free elements of $\mathbf{B}_r$ and $\mathbf{B}_c$ in the vector $\boldsymbol{\theta}$, the conditional log-likelihood is

$$\log \mathcal{L}(\mathbf{Y} \mid \boldsymbol{\theta}, \mathbf{A}, \mathbf{B}) = \alpha - T\Big(k \log \Big|\det \mathbf{B}_r\Big| + n \log \Big|\det \mathbf{B}_c\Big|\Big) - \frac{1}{2} \sum_{t=1}^{T} \text{tr}\Big(\boldsymbol{\Sigma}_r^{-1} \mathbf{U}_t \boldsymbol{\Sigma}_c^{-1} \mathbf{U}_t'\Big), \tag{11}$$

where the log-determinant term uses $\log|\det \mathbf{B}_0| = k\log|\det \mathbf{B}_r| + n\log|\det \mathbf{B}_c|$, the Jacobian of the map from shocks to innovations. Combining (11) with the Gaussian prior $\boldsymbol{\theta} \sim \mathcal{N}(\boldsymbol{\mu}_\theta, \boldsymbol{\Sigma}_\theta)$ of the previous subsection and the identified set $\mathcal{D}(\mathcal{S})$ gives the conditional posterior

$$p(\boldsymbol{\theta} \mid \mathbf{Y}, \mathbf{A}, \mathbf{B}) \;\propto\; \mathcal{L}(\mathbf{Y} \mid \boldsymbol{\theta}, \mathbf{A}, \mathbf{B})\, p(\boldsymbol{\theta} \mid \boldsymbol{\mu}_\theta, \boldsymbol{\Sigma}_\theta)\, \mathbf{1}\{\boldsymbol{\theta} \in \mathcal{D}(\mathcal{S})\}\,. \tag{12}$$

Because $\boldsymbol{\theta}$ enters the likelihood through both the determinant and the inverse of $\mathbf{B}_0$, (12) belongs to no known family: there is no closed-form draw and no conjugate update. We sample it with an elliptical slice sampler, which we describe next.

2.4.3 Elliptical Slice Sampling (ESS)

To sample from the conditional posterior (12) of the structural parameters we implement the elliptical slice sampling algorithm of Murray et al. (2010). The method is tailored to target a posterior constructed from a Gaussian prior multiplied by an arbitrary likelihood and it requires neither a Metropolis accept-reject step nor any stepsize tuning.[5]

The closest framework to ours is by Arias et al. (2025), who embed the elliptical slice sampler in a Gibbs sampler for sign-identified SVARs and show that it remains fast in large dimensions where accept-reject sampling becomes infeasible. Our sampler shares that logic and differs in what the slice sampler moves. They work in the orthogonal reduced-form parameterization of Arias et al. (2018) and update each block in turn. The proposals from their algorithm come from the posterior distributions and satisfy the restrictions needed by the orthogonal matrices. We instead parameterize the structural impact matrix directly through $\boldsymbol{\theta}$ and move its free elements, so that the ellipse carries the Gaussian prior $\mathcal{N}(\boldsymbol{\mu}_\theta, \boldsymbol{\Sigma}_\theta)$ and the likelihood (11) enters the slice threshold. In our setting the restrictions are enforced

[5]We have also experimented with the DREAM algorithm by Vrugt (2016), an efficient Metropolis–Hastings scheme that, unlike the elliptical slice sampler, accommodates arbitrary (non-Gaussian) priors on the structural parameters.

as a feasibility check that shrinks the bracket, while their algorithm only accepts draws where the orthogonal matrix fulfills the conditions.

The Proposal. Let $\boldsymbol{\theta}^{(m-1)}$ denote the current state. The sampler draws an auxiliary direction from the prior, $\boldsymbol{\nu} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_\theta)$, with candidate draws along the ellipse through $\boldsymbol{\theta}^{(m-1)}$ and $\boldsymbol{\nu}$:

$$\boldsymbol{\theta}^*(\alpha) = \boldsymbol{\mu}_\theta + \left(\boldsymbol{\theta}^{(m-1)} - \boldsymbol{\mu}_\theta\right)\cos\alpha + \boldsymbol{\nu}\sin\alpha, \qquad \alpha \in [0, 2\pi]. \tag{13}$$

The construction is such that $\boldsymbol{\theta}^*(\alpha)$ is marginally distributed as $\mathcal{N}(\boldsymbol{\mu}_\theta, \boldsymbol{\Sigma}_\theta)$ for *every* angle. Two special cases are worth noting. At $\alpha = 0$ we recover the current state, $\boldsymbol{\theta}^*(0) = \boldsymbol{\theta}^{(m-1)}$; at $\alpha = \pi/2$ we obtain an independent prior draw, $\boldsymbol{\theta}^*(\pi/2) = \boldsymbol{\mu}_\theta + \boldsymbol{\nu}$. The angle therefore interpolates between a null move and a bold, prior-sized move, and it is the only quantity that has to be chosen.

The Slice. The move along (13) is a rotation of the augmented pair $(\boldsymbol{\theta}, \boldsymbol{\nu})$ that leaves its joint Gaussian density unchanged. The prior therefore cancels.[6] The proposal for $\boldsymbol{\theta}$ is evaluated on the likelihood. We then compare it with:

$$\log y = \log \mathcal{L}(\mathbf{Y}|\boldsymbol{\theta}^{(m-1)}, \mathbf{A}, \mathbf{B}) + \log\omega, \qquad \omega \sim \text{Unif}(0, 1). \tag{14}$$

A proposal is accepted whenever its log-likelihood (11) lies above this threshold and when it satisfies the restrictions set by the researcher. We discuss ours in the empirical section. The angle in (13) is located by shrinking the bracket rather than by tuning:

1. Draw $\alpha \sim \text{Unif}[0, 2\pi)$ and initialize the bracket $[\alpha_{\min}, \alpha_{\max}] = [\alpha - 2\pi,\ \alpha]$.

2. Form $\boldsymbol{\theta}^*(\alpha)$ from (13) and evaluate the likelihood (11).

[6] The update rotates the whitened pair $(\mathbf{z}, \boldsymbol{\nu})$ into $(\mathbf{z}\cos\alpha + \boldsymbol{\nu}\sin\alpha,\ \boldsymbol{\nu}\cos\alpha - \mathbf{z}\sin\alpha)$, which preserves $\|\mathbf{z}\|^2 + \|\boldsymbol{\nu}\|^2$. It is this joint invariance, not constancy of the prior along the ellipse, that removes the prior from the acceptance condition (Murray et al. 2010).

3. If $\log \mathcal{L}(\mathbf{Y} \mid \boldsymbol{\theta}^*(\alpha), \mathbf{A}, \mathbf{B}) \geq \log y$ and $\boldsymbol{\theta}^*(\alpha)$ fulfills the restrictions, accept: set $\boldsymbol{\theta}^{(m)} = \boldsymbol{\theta}^*(\alpha)$ and stop.

4. Otherwise shrink the bracket towards the current state, setting $\alpha_{\min} = \alpha$ if $\alpha < 0$ and $\alpha_{\max} = \alpha$ otherwise, draw a new $\alpha \sim \text{Unif}(\alpha_{\min}, \alpha_{\max})$, and return to step 2.

The bracket always contains $\alpha = 0$, at which the proposal coincides with the current state; the latter is feasible and, because $\log \omega < 0$, strictly above the threshold. The update is rejection-free in the sense that every sweep returns a draw, and it needs no stepsize: the bracket adapts itself. In practice we work in the whitened coordinates $\mathbf{z} = \mathbf{D}^{-1}(\boldsymbol{\theta} - \boldsymbol{\mu}_\theta)$, with $\mathbf{D} = \boldsymbol{\Sigma}_\theta^{1/2}$ diagonal, so that $\boldsymbol{\nu}$ is a standard normal draw; this is equivalent to (13) and numerically more stable. The scheme scales well with the dimension of $\boldsymbol{\theta}$, which matters here because the structural block grows with both the number of countries and the number of variables.

Enforcing the Identified Set. The three classes of restrictions defining $\mathcal{D}(\mathcal{S})$ are enforced at different stages of the sampling update. Zero (equality) restrictions are imposed by construction: the slice sampler operates only on the free elements of $\boldsymbol{\theta}$ while leaving the others constant. Sign restrictions are incorporated into the feasibility check in Step 3. Consequently, any proposal that violates a sign restriction is treated in the same manner as one falling below the slice level, causing the sampling bracket to shrink until a feasible proposal is found[7]. Finally, magnitude (two-sided inequality) restrictions are enforced through an accept–reject step applied to the normalized draw. A proposal is retained only if

[7]Imposing the sign restrictions as a hard constraint at this point does not affect the reversibility of the slice sampler with respect to (12). The argument of Murray et al. (2010) holds for the Gaussian prior multiplied by an arbitrary non-negative function, and the indicator $\mathbf{1}\{\boldsymbol{\theta} \in \mathcal{D}(\mathcal{S})\}$ is one such function; no continuity of the likelihood is required. The hard sign restriction thus amounts to truncating the Gaussian prior to $\mathcal{D}(\mathcal{S})$ rather than to modifying the transition kernel.

Algorithm 1 MCMC Sampler for the BSMAR Model

Initialize $\mathbf{A}^{(0)}, \mathbf{B}^{(0)}, \kappa_{\mathbf{A}}^{(0)}, \kappa_{\mathbf{B}}^{(0)}, \boldsymbol{\theta}^{(0)}$
Compute initial structural matrices $\mathbf{B}_r^{(0)}, \mathbf{B}_c^{(0)}$ from $\boldsymbol{\theta}^{(0)}$ and covariance components $\boldsymbol{\Sigma}_r^{(0)} = \mathbf{B}_r^{(0)}\mathbf{B}_r^{(0)\prime}, \boldsymbol{\Sigma}_c^{(0)} = \mathbf{B}_c^{(0)}\mathbf{B}_c^{(0)\prime}$
for $m = 1$ **to** M **do**
    Step 1: Autoregressive Dynamics
    Compute $\mathbf{K}_{\mathbf{A}}$ and $\widehat{\mathbf{A}}$ using $\mathbf{B}^{(m-1)}$ and $\boldsymbol{\Sigma}_c^{(m-1)}$
    Draw $\text{vec}(\mathbf{A}^{(m)}) \sim \mathcal{N}(\text{vec}(\widehat{\mathbf{A}}), \boldsymbol{\Sigma}_r^{(m-1)} \otimes \mathbf{K}_{\mathbf{A}}^{-1})$
    Step 2: Cross-Sectional Dependencies
    Compute $\mathbf{K}_{\mathbf{B}}$ and $\widehat{\mathbf{B}}$ using $\mathbf{A}^{(m)}$ and $\boldsymbol{\Sigma}_r^{(m-1)}$
    Draw $\text{vec}(\mathbf{B}^{(m)}) \sim \mathcal{N}(\text{vec}(\widehat{\mathbf{B}}), \boldsymbol{\Sigma}_c^{(m-1)} \otimes \mathbf{K}_{\mathbf{B}}^{-1})$
    Step 3: Structural Parameters (Elliptical Slice Sampling)
    Draw candidate auxiliary direction $\boldsymbol{\nu} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_\theta)$
    Generate slice log-threshold $\log y = \log \mathcal{L}(\mathbf{Y} \mid \boldsymbol{\theta}^{(m-1)}, \mathbf{A}^{(m)}, \mathbf{B}^{(m)}) + \log \omega,\ \omega \sim \text{Unif}(0, 1)$
    Repeat until acceptable α is found:
        Propose $\boldsymbol{\theta}^*(\alpha) = \boldsymbol{\mu}_\theta + (\boldsymbol{\theta}^{(m-1)} - \boldsymbol{\mu}_\theta)\cos\alpha + \boldsymbol{\nu}\sin\alpha$
        Evaluate log-likelihood $\log \mathcal{L}^* = \log \mathcal{L}(\mathbf{Y} \mid \boldsymbol{\theta}^*(\alpha), \mathbf{A}^{(m)}, \mathbf{B}^{(m)})$
        Shrink bracket $[\alpha_{\min}, \alpha_{\max}]$ if $\log \mathcal{L}^* < \log y$ or a sign restriction is violated
    Reshape the candidate $\boldsymbol{\theta}^*(\alpha^\star)$ into $\mathbf{B}_r^*, \mathbf{B}_c^*$ and normalize
    If the magnitude restrictions hold at $\mathbf{B}_c^*$, set $\boldsymbol{\theta}^{(m)}, \mathbf{B}_r^{(m)}, \mathbf{B}_c^{(m)}$ to the candidate; otherwise retain $\boldsymbol{\theta}^{(m-1)}, \mathbf{B}_r^{(m-1)}, \mathbf{B}_c^{(m-1)}$
    Update $\boldsymbol{\Sigma}_r^{(m)} = \mathbf{B}_r^{(m)}\mathbf{B}_r^{(m)\prime}$ and $\boldsymbol{\Sigma}_c^{(m)} = \mathbf{B}_c^{(m)}\mathbf{B}_c^{(m)\prime}$
    Step 4: Shrinkage Parameters
    Draw $\kappa_{\mathbf{A}}^{(m)}$ and $\kappa_{\mathbf{B}}^{(m)}$ from their $\mathcal{GIG}$ conditional posteriors
    Update prior covariance matrices $\mathbf{V}_{\mathbf{A}}$ and $\mathbf{V}_{\mathbf{B}}$
end for

all inequality constraints are satisfied; otherwise, the previous draw is preserved[8]. Together, these steps ensure that every retained draw belongs to the identified set $\mathcal{D}(\mathcal{S})$. We provide an illustration of the sampler's capabilities in the Online Appendix, Section B.

[8]One can also impose the magnitude restrictions in the ESS step. In our implementation, we only implement sign restrictions in the feasibility check of Step 3 given that the number of strict sign restrictions is smaller, leading to fewer tries of the ESS sampler. Rejections at this step are rare; the sampler records the magnitude acceptance rate for every run.

3 Studying High-Dimensional International Spillovers

There has been long-standing and renewed interest in understanding the drivers of inflation, particularly whether demand or supply shocks have played a dominant role in its recent surge. This question has been addressed using single-country SVAR models identified through sign restrictions (e.g., Giannone & Primiceri 2025, Bergholt et al. 2026, Chang et al. 2023). Spillovers during the pandemic period have been analyzed by Di Giovanni et al. (2023) in a multi-country, multi-sector New Keynesian model.

We contribute to this literature by providing a multi-country perspective that allows for spillovers across countries. In this section, we apply our proposed framework to a panel of 17 economies. We identify country-specific supply and demand shocks and study their spillovers to other countries over time. We begin by describing the data, then discuss our identification strategy and conclude by presenting our empirical findings.

3.1 Data and Specification

We use quarterly data for a panel of 17 economies spanning North America, Europe, and the Asia–Pacific region. The panel has the United States (USA) ordered first followed by Australia (AUS), Belgium (BEL), Canada (CAN), China (CHN), Germany (DEU), Spain (ESP), Finland (FIN), France (FRA), United Kingdom (GBR), Italy (ITA), Japan (JPN), South Korea (KOR), Netherlands (NLD), Norway (NOR), Portugal (PRT), and Sweden (SWE). For each country, we observe real gross domestic product (GDP) and the consumer price index (CPI). GDP and CPI are taken from the OECD Quarterly National Accounts and Consumer Price Indices databases, both downloaded from the OECD Data Explorer; the series are listed in the Online Appendix, Section A.

The estimation sample covers the period from 1998Q1 to 2019Q4. To construct the HDs, we use an out-of-sample evaluation period from 2020Q1 to 2024Q4, encompassing both

the COVID-19 pandemic and the subsequent global inflation episode. Prior to estimation, GDP and CPI are transformed into yearly growth rates and standardized. We estimate the model using $p = 4$ quarterly lags. The choice of the lag length is standard in SVAR studies using quarterly data.

3.2 Identification of Supply and Demand in a Multi-Country Setup

We now turn to the identification of supply and demand shocks in the multi-country BSMAR model and present the full set of restrictions used in our baseline empirical specification. Recall that from (5) the contemporaneous structural impact matrix is given by the bilinear decomposition $\mathbf{B}_0 = \mathbf{B}_c \otimes \mathbf{B}_r$. With output and inflation as the variables in each country block, this structure allows us to identify country-specific supply and demand shocks while accounting for contemporaneous international spillovers.

Our identifying restrictions are summarized in Table 1. We motivate these assumptions as follows. First, the sign restrictions imposed on $\mathbf{B}_r$ follow the standard economic interpretation of supply and demand shocks (see, e.g., Giannone & Primiceri (2025)). Both shocks are restricted to be expansionary for economic activity: a positive demand shock raises activity and inflation, whereas a positive (favorable) supply shock raises activity while lowering inflation. The two shocks are therefore distinguished by the sign of the inflation response. Own-country effects are positive, consistent with the expected direction of domestic shocks. This is not a separate restriction but a consequence of the magnitude restriction introduced below, which is stated on the own effect including its sign. For the off-diagonal elements, we assume that international spillovers are likely to operate in the same direction as domestic effects, which we implement using soft sign restrictions (Section 2.3.2).[9]

[9] In the Online Appendix, Section D, we provide a robustness study of the flexibilization of our prior

To separate country-specific shocks from each other, we impose magnitude restrictions based on the assumption that country-specific shocks have their largest impact on the country in which they originate. The zero restrictions impose belief of the prominent role of the US in the global economy. Namely, we assume that the US is not contemporaneously affected by shocks originating in other countries.

Table 1: Identifying restrictions in the baseline empirical specification.

Family	**Restriction**	**Reading**
Sign, hard	$\text{sign}(\mathbf{B}_r) = \left(\begin{smallmatrix} + & + \\ - & + \end{smallmatrix}\right)$	Supply raises output, lowers prices; demand raises both
Soft sign	$P(b_c^{j_1 j_2} > 0) = 0.95$	Spillovers keep the domestic sign
Magnitude	$b_c^{j_2 j_2} \geq \lvert b_c^{j_1 j_2} \rvert$	A shock is largest at home
Zero	$b_c^{1 j_2} = 0,\ j_2 \neq 1$	The US is contemporaneously exogenous
Normalization	$b_c^{11} = 1$	Fixes the scale

Note: j_2 denotes the country in which the shock originates and j_1 the responding country. The hard sign restrictions are imposed inside the elliptical slice sampler, as a feasibility check that shrinks the angular bracket, and the magnitude restriction by an accept–reject step applied to the normalized draw, the zero restrictions as the exact equality $\mathbf{R}_0\boldsymbol{\theta} = \mathbf{0}$, and the soft sign restriction through the Gaussian prior $b_{c,0}^{j_1 j_2} = 0.5$, $v_{c,0}^{j_1 j_2} \approx 0.0924$, so that $0.5/\sqrt{v_{c,0}^{j_1 j_2}} = z_{0.95}$ and the data can override it; the own effects carry the tight prior $b_{c,0}^{jj} = 1$, $v_{c,0}^{jj} = 0.1$.

Taken together, these restrictions define the identified set $\mathcal{D}(\mathcal{S})$ of the baseline specification and allow us to recover, for each country, its supply and demand shocks as well as the cross-border spillovers between countries.

3.3 Empirical Results

We organize the discussion around four sets of objects: SIRFs, FEVDs, out-of-sample HDs, and connectedness measures. The formulas for the results presented in this Section are

assumptions. Namely, we relax the prior covariance of the off-diagonal entries. Our findings imply that loosening the prior leads to a reduction in the identification of some spillovers. Nonetheless, we find evidence of demand and supply spillovers between countries, with the United States as an important transmitter.

described in the Online Appendix, Section C. Diagnostics on the mixing are presented in the Online Appendix, Section G.1.4. The posterior traces show no drift, and the autoregressive blocks mix well as shown by the inefficiency factors in every specification. The structural block $\mathbf{B}_c$ is the slowest, as the only one drawn subject to weaker identifying restrictions.

Structural Impulse Responses. Figure 1 reports the structural impulse responses of GDP and CPI growth to the identified US aggregate supply (AS) and aggregate demand (AD) shocks, with the GDP responses in the left block and the CPI responses in the right. The responses to the demand shock show homogeneity in the sign and dynamics, a by-product of the sign restrictions on $\mathbf{B}_r$. More interestingly, the spillovers also maintain the same pattern abroad. That is, this shock leads to an increase in prices and output. The domestic US responses are the largest and the ones with least uncertainty. The responses are largest on impact and die out within about six quarters, after a small overshoot. The responses are mostly similar in shape across countries, while their magnitude and posterior uncertainty differ. Given the dominant role of the US, almost all responses exclude zero from their posterior credible bands. For shocks originating in other countries, we find, as expected, wider posterior bands, since smaller economies play a lesser role in driving the system. Not all responses are as precise as the ones from the US. Whether the responses to other shocks exclude the zero line varies case by case.

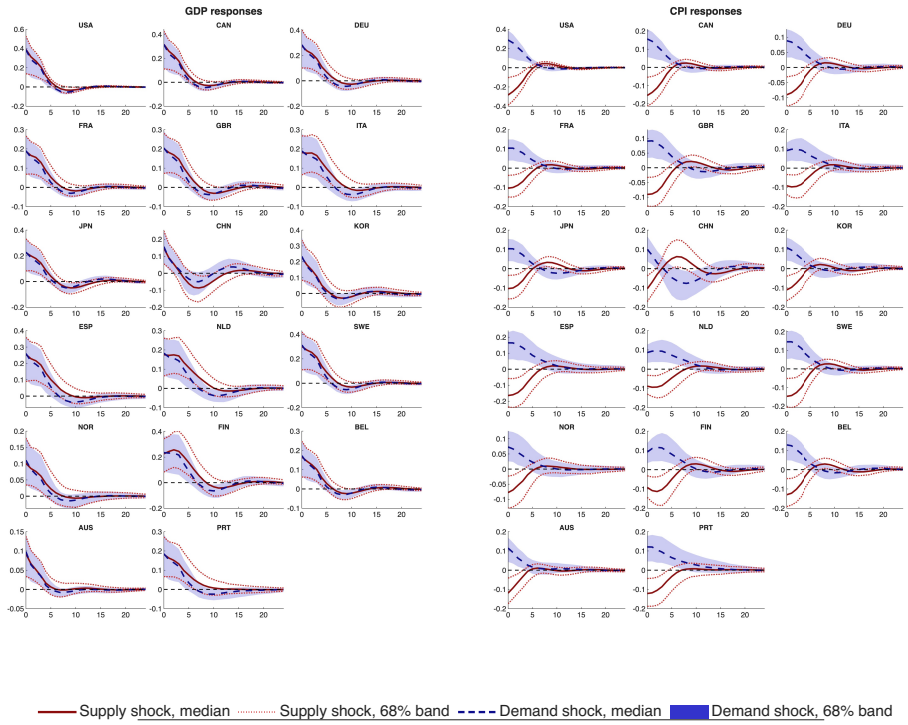


Figure 1: Structural impulse responses of GDP growth (left block) and CPI growth (right block) across the panel to the identified US aggregate supply (AS) shock (solid line, red band) and aggregate demand (AD) shock (dashed line, blue band), baseline specification. Posterior median and 68% credible band.

Forecast Error Variance Decompositions. The FEVD decomposes the shock contributions by the region the shock originates from. The results from the sampler for this measure are reported in the Online Appendix in Section G.1.1. The groups are: domestic, North America, Western Europe, and East Asia and the Pacific at each forecast horizon $h = 1, \ldots, 24$. The US decomposition is entirely domestic on impact, as the exogeneity restriction requires. At longer horizons foreign shocks, mostly from Western Europe and East Asia, account for a growing share, which is larger for CPI than for GDP. For every other country, shocks originating in Western Europe, the region with the most countries in the panel, account for a larger share of the forecast error variance than domestic shocks. Domestic shocks matter most for Norway, Japan, China and Australia. The shares change

little with the horizon, so that the spillovers are already in place on impact. Canada loads on both North America and Western Europe, and the East Asian and Pacific economies retain a visible own-region block alongside their exposure to Europe. The same section reports the decomposition by shock type, which splits each of these blocks into supply and demand. We find that supply and demand shocks contribute roughly equal shares. Supply shocks account for slightly more of the variance of GDP at longer horizons, and demand shocks for slightly more of the variance of CPI. Results from other countries are also presented there.

Historical Decompositions. We follow Giannone & Primiceri (2025) by calculating out-of-sample HDs for the period 2020Q1–2024Q4 to analyze whether the recent increase in inflation was demand or supply driven. Figure 2 reports the decomposition of UK CPI in the left column and US CPI in the right. The figure has four panels: the upper row splits each CPI series into domestic and foreign supply and demand shocks, while the lower row disaggregates the same series by region of origin. In the US, demand shocks clearly dominate the development of CPI during the pandemic. Domestic demand shocks drive both the fall in inflation in 2020 and most of its rise in 2021 and 2022, when domestic supply shocks add a smaller positive contribution. The origin of the shocks is mostly domestic, with minor contributions from Western Europe, East Asia, and the Pacific. The UK, by contrast, shows a more balanced contribution of supply and demand shocks over the COVID period, and almost all of it originates abroad. North American shocks account for most of both the fall in UK inflation in 2020 and 2021 and its rise in 2022, while East Asian shocks pull it down early in the period and push it up after 2023. Placing the two economies side by side reinforces this contrast. US inflation is mainly demand-driven and domestic in origin, whereas UK inflation is driven by foreign shocks, with supply shocks playing a larger role. Overall, and in line with Giannone & Primiceri (2025), these decompositions point to

demand shocks as the main explanatory force behind the surge of US inflation during the pandemic. The HDs for the remaining countries and variables are presented in the Online Appendix, Section G.1.2. In them, the foreign contributions come mostly from Western Europe, the largest group of economies in the panel, with North American and East Asian shocks contributing smaller and similar shares.

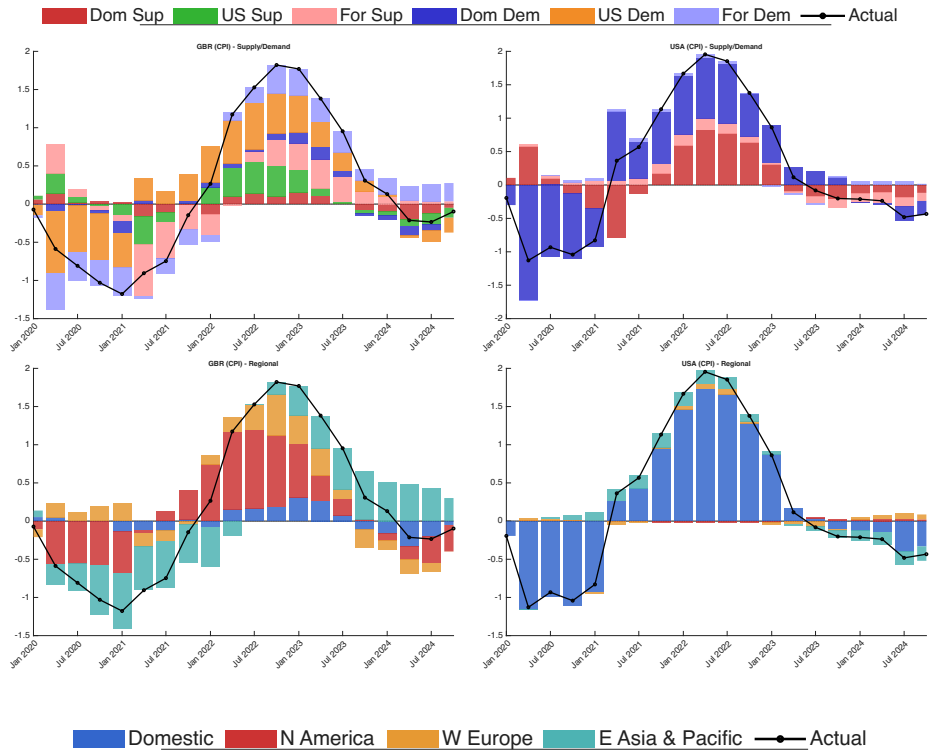


Figure 2: Out-of-sample historical decomposition of CPI for the UK (left column) and the US (right column), baseline specification. Top row: supply/demand decomposition; bottom row: regional decomposition. The supply/demand legend is shown above; the regional legend below.

Connectedness. Another indicator to summarize the spillover effects is the connectedness measure of Diebold & Yılmaz (2014). Figure 3 presents the network implied by this result, showing the spillovers between countries in both magnitude and direction: the thickness of each edge represents the strength of the directional spillover, while the color of each

node indicates whether the country is a net transmitter (red) or a net receiver (blue) and its size the magnitude of that net position. The US is by far the largest net transmitter, with its largest spillover to Canada and visible spillovers to most European economies. Australia is a small net transmitter, while the remaining countries are net receivers. The ranking of spillovers changes with the horizon: short horizons display fewer connections, because the contemporaneous exogeneity of the US binds directly on the impact matrix. We report results for more horizons and under alternative assumptions in the Online Appendix, Sections G.1.3 and G.2.3.[10]

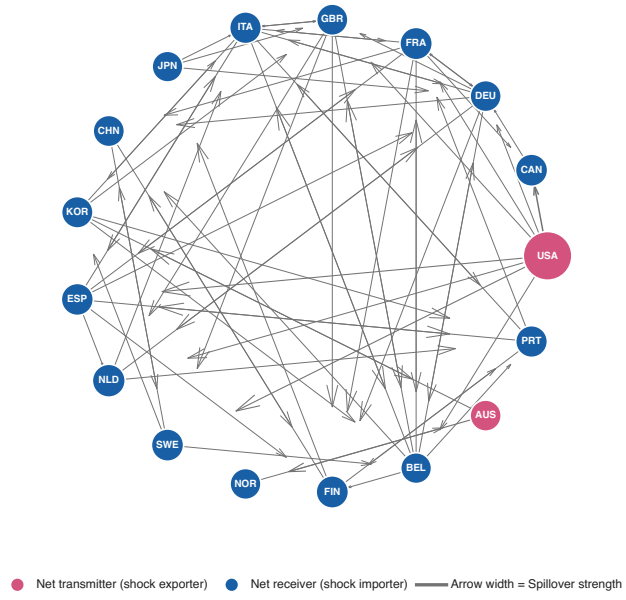


Figure 3: Connectedness network of the spillover effects between countries on impact ($H = 1$), from connectedness tables computed for each posterior draw and averaged over the draws.

[10]Another interesting measure of connectedness is presented in Buchwalter et al. (2026). The application of this new measure is left as an extension for future work.

4 Summary

This paper develops a Bayesian Structural Matrix Autoregression (BSMAR) framework for identifying country-specific macroeconomic shocks and their contemporaneous transmission across countries in large international systems. By exploiting the matrix structure of the data and decomposing the structural impact matrix into variable- and country-specific components, the framework substantially reduces the dimensionality of the model while allowing for rich cross-country spillovers. The proposed Bayesian sampler accommodates a broad class of identification restrictions, providing a flexible framework for structural analysis in high-dimensional systems.

Applying the model to 17 economies, we find a vast number of spillovers between countries. Our results also, partly, validate the findings of the literature by establishing the dominance of demand shocks during the pandemic period. The United States emerges as an important transmitter of shocks, generating sizable spillovers to the rest of the countries in the panel. Compared with the benchmark model in the literature, we find that our model identifies a large number of spillovers under the sign restrictions placed on the behavior of the data, while allowing for violations of these assumptions. Overall, our findings underscore the importance of jointly accounting for country-specific shocks and contemporaneous spillovers when studying the international transmission of macroeconomic disturbances. An interesting direction for future research would be the identification of common shocks that affect all countries at once, e.g., oil or commodity price shocks. This would allow the model to separate global and country-specific shocks. We leave this as a path for future research.

Declarations

Acknowledgments. We would like to thank Chenghan Hou, Bowen Fu, Wei Zhang, Boris Blagov, participants from the ESOBE 2026, Statistical Week and the TRR 391 conference for comments.

Funding. This research was (partially) funded in the course of TRR 391 Spatio-temporal Statistics for the Transition of Energy and Transport (520388526) by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation). Jan Prüser gratefully acknowledges the support of the German Research Foundation (DFG, 468814087).

Disclosure Statement. The authors report there are no competing interests to declare.

Data Availability Statement. The data are publicly available. The Online Appendix presents the origin and transformation of the data.

Supplementary Material. Additional results, extensions and benchmark comparison are provided in the Online Appendix.

AI Usage. Claude Fable 5, Opus 5.5 and 5.4 were used for proof-reading, language editing, and coding assistance.

References

Arias, J. E., Rubio-Ramírez, J. F. & Shin, M. (2025), Large SVARs, Working paper, Federal Reserve Bank of Philadelphia.

Arias, J. E., Rubio-Ramírez, J. F. & Waggoner, D. F. (2018), 'Inference based on structural vector autoregressions identified with sign and zero restrictions: Theory and applications', *Econometrica* **86**(2), 685–720.

Baumeister, C. & Hamilton, J. D. (2015), 'Sign restrictions, structural vector autoregressions, and useful prior information', *Econometrica* **83**(5), 1963–1999.

Bergholt, D., Canova, F., Furlanetto, F., Maffei-Faccioli, N. & Ulvedal, P. (2026), 'What drives the recent surge in inflation? The historical decomposition roller coaster', *American Economic Journal: Macroeconomics* **18**(4), 34–63.

Blanchard, O. J. & Quah, D. (1989), 'The dynamic effects of aggregate demand and supply disturbances', *American Economic Review* **79**(4), 655–673.

Bucci, A. (2025), 'A smooth transition autoregressive model for matrix-variate time series', *Computational Economics* **65**(1), 429–458.

Buchwalter, B., Diebold, F. X. & Yilmaz, K. (2026), 'Clustered network connectedness: A new measurement framework with application to global equity markets', *Journal of Econometrics* p. 106243.

Camehl, A. & von Schweinitz, G. (2026), 'What explains international interest rate co-movement?', *Journal of Applied Econometrics* **41**(4), 343–359.

Celani, A. & Pagnottoni, P. (2024), 'Matrix autoregressive models: Generalization and Bayesian estimation', *Studies in Nonlinear Dynamics & Econometrics* **28**(2), 227–248.

Celani, A., Pagnottoni, P. & Jones, G. (2024), 'Bayesian variable selection for matrix autoregressive models', *Statistics and Computing* **34**(2), 91.

Chan, J. C. & Qi, Y. (2026), 'Large Bayesian matrix autoregressions', *Journal of Econometrics* **256**, 105955.

Chan, J. C. & Zhang, W. (2024), 'Bayesian dynamic factor models for high-dimensional matrix-valued time series'. arXiv preprint 2409.08354 [econ.EM]. https://arxiv.org/abs/2409.08354

Chang, J.-C. D., Jansen, D. W. & Pagliacci, C. (2023), 'Inflation and real GDP growth in the U.S.—demand or supply driven?', *Economics Letters* **231**, 111274.

Charnavoki, V. & Dolado, J. J. (2014), 'The effects of global shocks on small commodity-exporting economies: Lessons from Canada', *American Economic Journal: Macroeconomics* **6**(2), 207–237.

Chen, R., Xiao, H. & Yang, D. (2021), 'Autoregressive models for matrix-valued time series', *Journal of Econometrics* **222**(1), 539–560.

Crespo Cuaresma, J., Doppelhofer, G., Feldkircher, M. & Huber, F. (2019), 'Spillovers from US monetary policy: evidence from a time varying parameter global vector autoregressive model', *Journal of the Royal Statistical Society Series A: Statistics in Society* **182**(3), 831–861.

Dees, S., di Mauro, F., Pesaran, M. H. & Smith, L. V. (2007), 'Exploring the international linkages of the euro area: A global VAR analysis', *Journal of Applied Econometrics* **22**(1), 1–38.

Devroye, L. (2014), 'Random variate generation for the generalized inverse gaussian distribution', *Statistics and Computing* **24**(2), 239–246.

Di Giovanni, J., Kalemli-Özcan, Ş., Silva, A. & Yildirim, M. A. (2023), Pandemic-era inflation drivers and global spillovers, NBER Working Paper 31887, National Bureau of Economic Research.

Diebold, F. X. & Yılmaz, K. (2014), 'On the network topology of variance decompositions: Measuring the connectedness of financial firms', *Journal of Econometrics* **182**(1), 119–134.

Feldkircher, M. & Huber, F. (2016), 'The international transmission of US shocks—evidence from Bayesian global vector autoregressions', *European Economic Review* **81**, 167–188.

Forni, M. & Lippi, M. (2001), 'The generalized dynamic factor model: representation theory', *Econometric Theory* **17**(6), 1113–1141.

Giannone, D., Lenza, M. & Primiceri, G. E. (2015), 'Prior selection for vector autoregressions', *The Review of Economics and Statistics* **97**(2), 436–451.

Giannone, D. & Primiceri, G. E. (2025), Demand-driven inflation. Working paper. https://faculty.wcas.northwestern.edu/gep575/DemandDrivenInflation_4-3.pdf

Hou, C. (2024), 'Large Bayesian SVARs with linear restrictions', *Journal of Econometrics* **244**(1), 105850.

Kilian, L. & Lütkepohl, H. (2017), *Structural vector autoregressive analysis*, Cambridge University Press.

Kong, X., Liu, Y., Yu, L. & Zhao, P. (2026), 'Matrix quantile factor model', *Journal of Business & Economic Statistics* . Forthcoming.

Litterman, R. B. (1986), 'Forecasting with Bayesian vector autoregressions: Five years of experience', *Journal of Business & Economic Statistics* **4**(1), 25–38.

Mumtaz, H. & Surico, P. (2009), 'The transmission of international shocks: a factor-augmented VAR approach', *Journal of Money, Credit and Banking* **41**(s1), 71–100.

Murray, I., Adams, R. P. & MacKay, D. J. C. (2010), Elliptical slice sampling, *in* Y. W. Teh & M. Titterington, eds, 'Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics', Vol. 9 of *Proceedings of Machine Learning Research*, PMLR, pp. 541–548. https://proceedings.mlr.press/v9/murray10a.html

Pesaran, M. H., Schuermann, T. & Weiner, S. M. (2004), 'Modeling regional interdepen-

dencies using a global error-correcting macroeconometric model', *Journal of Business & Economic Statistics* **22**(2), 129–162.

Pfarrhofer, M. & Stelzer, A. (2025), 'High-frequency and heteroskedasticity identification in multicountry models: Revisiting spillovers of monetary shocks', *Macroeconomic Dynamics* **29**, e122.

Read, M. & Zhu, D. (2025), Fast posterior sampling in tightly identified SVARs using 'soft' sign restrictions, Research Discussion Paper 2025-03, Reserve Bank of Australia.

Rubio-Ramírez, J. F., Waggoner, D. F. & Zha, T. (2010), 'Structural vector autoregressions: Theory of identification and algorithms for inference', *Review of Economic Studies* **77**(2), 665–696.

Rue, H. & Held, L. (2005), *Gaussian Markov Random Fields: Theory and Applications*, Chapman & Hall/CRC, Boca Raton.

Samadi, S. Y. & Billard, L. (2025), 'On a matrix-valued autoregressive model', *Journal of Time Series Analysis* **46**(1), 3–32.

Sims, C. A. (1980), 'Macroeconomics and reality', *Econometrica* **48**(1), 1–48.

Stock, J. H. & Watson, M. W. (2005), Implications of dynamic factor models for VAR analysis, NBER Working Paper 11467, National Bureau of Economic Research.

Vrugt, J. A. (2016), 'Markov chain Monte Carlo simulation using the DREAM software package: Theory, concepts, and MATLAB implementation', *Environmental Modelling & Software* **75**, 273–316.

Ye, S., Zheng, T., Xiao, H. & Hong, Y. (2026), 'Multi-matrix autoregressive models with an application to multi-modal network', *Journal of Business & Economic Statistics* . Forthcoming.