

Article

Inference of Impulse Responses via Bayesian Graphical Structural VAR Models

Daniel Felix Ahelegbey 

School of Mathematics, Statistics, and Actuarial Science, University of Essex, Colchester CO4 3SQ, UK;
d.f.ahelegbey@essex.ac.uk

Abstract: Impulse response functions (IRFs) are crucial for analyzing the dynamic interactions of macroeconomic variables in vector autoregressive (VAR) models. However, traditional IRF estimation methods often have limitations with assumptions on variable ordering and restrictive identification constraints. This paper applies the Bayesian graphical structural vector autoregressive (BGSVAR) model, which integrates structural learning to capture both temporal and contemporaneous dependencies for more accurate impulse response estimation. The BGSVAR framework provides a more efficient and interpretable method for estimating IRFs, which can enhance both forecasting performance and structural inferences in economic modelling. Through extensive simulations across various data-generating processes, we evaluate BGSVAR's effectiveness in modelling dynamic interactions among US macroeconomic variables. Our results demonstrate that BGSVAR outperforms traditional methods, such as LASSO and Bayesian VAR (BVAR), by delivering more precise impulse response estimates and better capturing the structural dynamics of VAR-based models.

Keywords: Bayesian graphical model; Bayesian VAR; impulse response; orthogonalized; generalized; LASSO

JEL Classification: C11; C15; C32; C53; E24; E27



Academic Editor: Deborah Gefang

Received: 27 January 2025

Revised: 16 March 2025

Accepted: 24 March 2025

Published: 2 April 2025

Citation: Ahelegbey, D. F. (2025). Inference of Impulse Responses via Bayesian Graphical Structural VAR Models. *Econometrics*, 13(2), 15. <https://doi.org/10.3390/econometrics13020015>

Copyright: © 2025 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The primary goal of structural vector autoregressive (SVAR) models is structural identification, which aims to determine the dynamic relationships between interrelated time series variables. While a reduced-form VAR model captures serial correlation and cross-correlation among variables, it treats contemporaneous relationships as part of the shocks. VAR models have become standard tools in econometrics for analyzing and forecasting system dynamics. A key feature of these models is the impulse response function (IRF), which quantifies the impact of structural shocks on the system. Common methods for analyzing IRFs include the orthogonalized IRF (OIRF; Sims, 1980), generalized IRF (GIRF; H. H. Pesaran & Shin, 1998), and local projection IRF (LP-IRF; Jordà, 2005).

The OIRF estimates the impact of structural shocks using the Cholesky decomposition of the VAR error covariance matrix. This decomposition creates a unique variable ordering with a lower triangular matrix. Identification is typically achieved by imposing exclusion restrictions (Bernanke & Mihov, 1998; Sims, 1980), long-run restrictions (Blanchard & Quah, 1989; King et al., 1991), or sign restrictions on variable responses to shocks (Canova & De Nicrolo, 2002; Faust, 1998; Uhlig, 2005) or by categorizing variables into slow- and fast-moving groups (Banbura et al., 2010; Christiano et al., 1999; Eichenbaum & Evans,

1995; Kilian, 2009). These restrictions are often based on “strongly held convictions” or the structure of a specific economic model. However, a major drawback is that changing the variable ordering can significantly affect the impulse response estimates, making the results highly sensitive to the underlying theoretical assumptions (Kilian, 2013). This makes such methods challenging to apply in situations where there is no clear theory or convictions for identification.

To address the issue of variable ordering in the OIRF, H. H. Pesaran and Shin (1998) builds on Koop et al. (1996) to propose the GIRF, which is order-invariant. Unlike the OIRF, which uses a lower triangular matrix, the GIRF assumes a fully symmetric covariance structure. This makes it a widely used alternative, as it avoids imposing identification restrictions (F. Diebold & Yilmaz, 2014; F. X. Diebold & Yilmaz, 2012; M. H. Pesaran et al., 2004). However, while the GIRF offers advantages over the OIRF, the assumption of a fully connected system of endogenous variables—where all variables react instantaneously to shocks—can be problematic. If only a few variables are contemporaneously related, the GIRF may produce inaccurate impulse response estimates.

Jordà (2005) introduced local projection (LP) as an alternative to SVAR for measuring impulse responses. LP avoids the invertibility problem of VAR-IRFs by using linear projections, where forecasting models are re-estimated at each horizon using ordinary least squares. This method projects the dependent variable onto the space generated by its past values at each forecast horizon. Jordà demonstrates that LP-IRFs can outperform estimates from a misspecified VAR model. Despite its simplicity and growing popularity, LP-IRF faces theoretical and practical issues (see Choi & Chudik, 2019; Jordà, 2023; Jordà & Taylor, 2025; Li et al., 2024; Plagborg-Møller & Wolf, 2021; Stock & Watson, 2018). Stock and Watson (2018) and Plagborg-Møller and Wolf (2021) compare LP-IRFs and VAR-IRFs, with and without external instruments. Although there are differences, both methods are conceptually similar. Kilian and Kim (2011) shows that LP-IRF confidence intervals are often less accurate and wider than those derived from well-specified VAR models.

It is well-known that VAR-IRF estimates are non-linear functions of reduced-form slope coefficient matrices and the contemporaneous shock matrix. In contrast, LP-IRF derives structural impulse estimates from the reduced-form impulse response coefficients and the structural impact multiplier matrix. Both methods rely on the slope coefficient matrices and the structural shock matrix to assess the propagation of shocks. A significant limitation of these conventional methods is how the coefficient matrices are estimated. Rather than relying on theoretical restrictions to impose contemporaneous and lagged relationships, it would be more appropriate to let the data determine the structure of these matrices. Additionally, it is well established that some sets of variables are better modelled with information from a few other variables, while others may require a broader set. Therefore, using an unrestricted VAR with many variables may not effectively capture the shock effects in a dynamic system. Conversely, reducing the model to a small set of variables may lead to misspecification, omitting important variables and dynamics, which can distort the results.

To avoid over-parameterization and overfitting in VAR models, several techniques have been discussed in the literature, including Bayesian inference with suitable prior distributions (Banbura et al., 2010; Bassetti et al., 2014; Canova & Ciccarelli, 2004; Kalli & Griffin, 2018; Karlsson, 2013; Koop & Korobilis, 2016; Korobilis, 2013), dynamic factor and compression models (Bai & Ng, 2002; Bernanke et al., 2005; Forni et al., 2005; Koop et al., 2019; Lopes & Carvalho, 2007; Stock & Watson, 2002), selection and shrinkage methods (Gefang, 2014; Medeiros & Mendes, 2016; Tibshirani, 1996; Zou & Hastie, 2005), and graphical models for time series data (Ahelegbey et al., 2016a, 2016b; Casarin et al., 2020; Corander & Villani, 2006; Demiralp & Hoover, 2003; Eichler, 2007). Sparse VAR models

help address the issue of impulse response estimation by selecting relevant predictors and eliminating redundant or irrelevant variables. This selection process aims to identify the most informative variables to improve model fit and accuracy. Recent advancements in variable selection techniques for VAR models include (Ahelegbey et al., 2016a, 2016b, 2021; Basu & Michailidis, 2015; Billio et al., 2019; Celani et al., 2024; Davis et al., 2016; Kock & Callot, 2015).

This paper extends the BGSVAR approach of Ahelegbey et al. (2016a) to estimate IRFs. The BGSVAR model views SVAR as a graphical (network) model, where the coefficients represent conditional dependencies among variables. This technique treats both the structural identification of SVAR and model selection in VAR as a joint problem, solved by learning the contemporaneous and temporal dependencies among the variables. The BGSVAR model's identification approach uses a Bayesian procedure, enabling the incorporation of prior information to avoid including uninterpretable relationships in the final model. Ahelegbey et al. (2016a) demonstrated that BGSVAR produces sparse, parsimonious models that improve forecasting performance and yield more meaningful structural inferences.

The contribution of this work is twofold. First, we propose a novel approach to impulse response analysis by leveraging the BGSVAR model, which provides a more accurate estimation of both lagged and contemporaneous coefficient matrices necessary for impulse response computation. We demonstrate that the BGSVAR estimator is more robust compared to established methods in the literature. Second, we apply our approach to well-known macroeconomic data sets, specifically modelling the impulse response of shocks to key US macroeconomic time series within a moderate-dimension VAR framework, as in Cimadomo et al. (2022). By comparing our method with traditional techniques such as BVAR and LASSO-VAR, we show that BGSVAR yields more reliable impulse response estimates.

This paper is organized as follows: Section 2 reviews the SVAR model and discusses various techniques for impulse response analysis. Section 3 outlines the Bayesian inference process for estimating the BGSVAR model parameters. Section 4 presents the simulation results based on real macroeconomic data. Finally, Section 5 offers concluding remarks.

2. SVAR and Impulse Response Analysis

This section presents the SVAR model, outlines the orthogonalized and generalized IRF, and describes the Bayesian graphical approach to estimating the IRF.

2.1. SVAR Model

The dynamic interactions between interrelated endogenous variables that exhibit time-lagged and contemporaneous dependencies can be modelled as the SVAR(p) process:

$$Y_t = B_0 Y_t + B_1 Y_{t-1} + \dots + B_p Y_{t-p} + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \Sigma_\varepsilon) \quad (1)$$

where Y_t is $(n \times 1)$ vector of observed variables; B_0 is $(n \times n)$ zero diagonal matrix of contemporaneous coefficients; $B_k, k = 1, \dots, p$ are $n \times n$ matrices of autoregressive coefficients; and ε_t is $(n \times 1)$ vector of structural shocks, independent and identically normal with the diagonal covariance matrix Σ_ε which, sometimes, is normalized to be an identity matrix. Under the assumption that $(I_n - B_0)$ is invertible, (1) can be expressed in a reduced form as

$$Y_t = B_1^* Y_{t-1} + \dots + B_p^* Y_{t-p} + U_t, \quad U_t \sim \mathcal{N}(0, \Sigma_u) \quad (2)$$

where $A_0 = (I_n - B_0)^{-1}$, $B_k^* = A_0 B_k$, $U_t = A_0 \varepsilon_t$ is a vector of reduced-form VAR innovations, with the variance-covariance matrix $\Sigma_u = A_0 \Sigma_\varepsilon A_0'$.

2.2. Impulse Response Analysis via VAR

This dynamic mechanism of analyzing shock propagation over time in a given system is referred to as impulse response analysis. IRFs measure the effect of a shock to $Y_{j,t}$ on $Y_{i,t+h}$, $h \geq 0$, assuming that no other shocks occur between t and $t+h$. This can be defined as the partial derivative of $Y_{i,t+h}$ for some $h \geq 0$ for $\varepsilon_{j,t}$. To derive the IRF, (2) can be re-written as a vector moving average where Y_t is expressed as a linear function of current and past structural errors, U_t (or ε_t):

$$Y_t = \sum_{i=0}^{\infty} \Psi_i U_{t-i} = \sum_{i=0}^{\infty} \Psi_i A_0 \varepsilon_{t-i} \quad (3)$$

where $\Psi_i = B_1^* \Psi_{i-1} + \dots + B_p^* \Psi_{i-p}$, with $\Psi_0 = I_n$. The above expressions show that the response of Y_t to shocks to ε_{t-i} depends on non-linear functions of B_k^* 's and A_0 . In practice, traditional techniques estimate unrestricted B_k^* 's from reduced-form VARs which implies that all variables are serially correlated and cross-correlated. The next step is to identify A_0 from Σ_u by imposing a theory-driven structure on the contemporaneous dependence between variables.

2.2.1. Orthogonalized Impulse Response—OIR

In the standard application of VAR, the adoption of orthogonal shocks is used to achieve identification. The commonest approach is Cholesky decomposition, which assumes that the model is fully recursive so that B_0 is lower triangular with zero diagonals. The orthogonalized impulse response of $Y_{i,t+h}$ to a shock to $Y_{j,t}$ is as follows (OIR; Sims, 1980):

$$OIR_{ij}(h) = \frac{\partial Y_{i,t+h}}{\partial \varepsilon_{j,t}} = [\Psi_h A_0]_{ij} = \mathbf{e}_i' \Psi_h A_0 \mathbf{e}_j, \quad h = 0, 1, 2, \dots \quad (4)$$

where $\mathbf{e}_j$ is an $n \times 1$ vector with unity as its j -th element and zeros elsewhere, and A_0 is a lower triangular Cholesky factorization of Σ_u . The i -th row and j -th column of $\Psi_h A_0$ indicate the response of $Y_{i,t+h}$ to a unit change in $\varepsilon_{j,t}$ holding all other innovations at all dates constant. The lower triangular structure of A_0 imposes a recursive causal order where the first variable responds to its exogenous shock, while the second responds to shocks of the first variable plus its shock, and so on.

Recovering A_0 from Σ_u via Cholesky factorization is usually carried out by imposing short-run, long-run, or sign restrictions based on some economic model or theory. Much research achieves this by dividing the variables into two categories: slow- and fast-moving (see Banbura et al., 2010; Christiano et al., 1999; Eichenbaum & Evans, 1995; Kilian, 2009; Stock & Watson, 2018). In many cases, the empirical results are only as credible as the underlying theory (see Kilian, 2013). In some instances, there are not enough credible exclusion restrictions to achieve identification. It can be seen that changing the variable ordering of the model changes the structure of A_0 and affects the results dramatically.

2.2.2. Generalized Impulse Response—GIR

The GIR is an alternative to the OIR to reduce the reliance on theory for variable ordering. It is an order-invariant technique that assumes a full symmetric covariance structure instead of a lower triangular matrix. The generalized impulse response of $Y_{i,t+h}$ to a shock to $Y_{j,t}$ is given by the following (H. H. Pesaran & Shin, 1998):

$$GIR_{ij}(h) = \sigma_{jj}^{-1/2} (\mathbf{e}_i' \Psi_h \Sigma_u \mathbf{e}_j), \quad h = 0, 1, 2, \dots \quad (5)$$

where σ_{jj} is the j -th diagonal element of Σ_u . Here, when one variable is shocked, the other variables also vary following the structure of Σ_u . When Σ_u is diagonal, then A_0 is also diagonal, which means that the GIR is the same as the OIR. Both the GIRF and OIRF yield the same result for the shock to the first variable in Y_t .

The GIR is deeply embedded in the assumption that all the variables are contemporaneously related, in the sense that a shock to one variable instantaneously affects the rest of the system. If only a few of these variables (or none) are contemporaneously related, the GIR might produce propagation effects that can be misleading. It would be more reasonable to use identifying restrictions that are described by the data instead of assuming that all variables are serially, contemporaneously, and lagged-related.

2.3. Graphical Structural VAR (GSVAR)

This section presents the methodology employed in the analysis. SVAR models have been recently extended to graphical SVAR (GSVAR) where the lagged and contemporaneous dependence structure is encoded by a graph (network) (e.g., see [Ahelegbey et al., 2016a, 2016b, 2021](#); [Casarin et al., 2020](#); [Corander & Villani, 2006](#)). This is based on viewing SVAR as a graphical (network) model where the variables are represented as nodes and the links denote the coefficients that encoded the weights of the conditional (in)dependent relationships among the variables. In this paper, we extend [Ahelegbey et al. \(2016a\)](#) to analyse the impulse effect of structural shocks in dynamic systems.

Let Y_t be a vector of n units of sequences of endogenous variables observed at the time t . We model Y_t as GSVAR by $\mathcal{M} = \{V, G_{0:p}, \Phi_{0:p}\}$, where $V = (V_1, \dots, V_n)$ is a set of vertices which represents the variables in Y_t ; each $G_k(\Phi_k)$, $k = 0, 1, \dots, p$, is an $n \times n$ binary (coefficient) matrix that depicts the connectivity between variables at different times points such that the (i, j) -th entry of Φ_k is given by

$$\Phi_{k,ij} = \begin{cases} 0 & \text{if } G_{k,ij} = 0 \implies Y_{j,t-k} \not\rightarrow Y_{i,t} \\ \beta_{k,ij} \in \mathbb{R} & \text{if } G_{k,ij} = 1 \implies Y_{j,t-k} \rightarrow Y_{i,t} \end{cases} \quad (6)$$

Following [Ahelegbey et al. \(2016a\)](#), GSVAR can be re-written as

$$Y_t = \sum_{k=1}^p \Phi_k Y_{t-k} + U_t \quad (7)$$

$$U_t = \Phi_0 U_t + \varepsilon_t \quad (8)$$

where c is a constant term, and the structural error terms ε_t are i.i.d. with the covariance matrix Σ . The $(n \times n)$ -dimensional matrices Φ_k , $k = 1, \dots, p$ contain the autoregressive coefficients, and the matrix Φ_0 is full (non-symmetric) with zeroes on the main diagonal that records the contemporaneous dependence between the endogenous variables. The expressions in (7) and (8) can be written in a more compact form as

$$Y_t = \Phi_+ X_t + (I - \Phi_0)^{-1} \varepsilon_t \quad (9)$$

where $\Phi_+ = (\Phi_1, \dots, \Phi_p)$, and $X_t = (Y'_{t-1}, \dots, Y'_{t-p})'$. It can be shown that the matrix $(I - \Phi_0)^{-1}$ records the (in)direct contemporaneous effect of ε_t on Y_t . A shock to Y_{jt} can only affect Y_{it} if there is a contemporaneous link from Y_{jt} to Y_{it} .

Definition 1. Let ε_t denote a vector of orthogonal shocks with an identity covariance matrix (i.e., $\Sigma_\varepsilon = I_n$). Given that $\Phi_+ = (\Phi_1, \dots, \Phi_p)$ as the GSVAR slope coefficient matrices are conditional

on G_+ , with Φ_0 as the estimated contemporaneous coefficients matrix conditional on G_0 , we define the GSVAR impulse response estimator as

$$\text{GSVAR-IR}_{ij}(h) = \mathbf{e}_i' \tilde{\Psi}_h (I_n - \Phi_0)^{-1} \mathbf{e}_j, \quad h = 0, 1, 2, \dots \quad (10)$$

where $\tilde{\Psi}_i = \Phi_1 \tilde{\Psi}_{i-1} + \dots + \Phi_p \tilde{\Psi}_{i-p}$, with $\tilde{\Psi}_0 = I_n$.

Let $A_\phi = (I_n - \Phi_0)^{-1}$. If Φ_+ is dense or a full matrix and A_ϕ is a lower triangular structure, then the GSVAR-IR will produce estimates (approximately) equal to the OIR. On the other hand, if both Φ_+ and A_ϕ are full matrices, then the GSVAR-IR will produce (approximately) equal estimates to the GIR. Thus, the GSVAR-IR will be more robust compared to the OIR and the GIR.

3. Bayesian Graphical SVAR Model Estimation

Selecting appropriate variables for inclusion in each equation is the most important and challenging part of building the graphical SVAR model. We implement the Bayesian procedure of Geiger and Heckerman (2002) and Ahelegbey et al. (2016a, 2021) that allows a researcher to incorporate expert or prior information where necessary, to avoid including uninterpretable relationships among variables in the final model. The selection is based on the probabilistic concept of Suppes (1970) and is carefully undertaken to establish a causal relationship between the variables, which are then used to build a predictive forecasting model.

Let $X_t = (Y'_{t-1}, \dots, Y'_{t-p})'$ be an $np \times 1$ vector of lagged observations and $Z_t = (Y'_t, X'_t)'$ be $m \times 1$, where $m = n(p+1)$. Denote with $Z = (Z'_1, \dots, Z'_N)'$ the $N \times m$ matrix collection of Z_t . Under the assumption that $Z \sim \mathcal{N}(0, \Sigma)$, the likelihood function is given by

$$P(Z|\Omega, G) = (2\pi)^{-\frac{mN}{2}} |\Omega|^{\frac{N}{2}} \exp \left\{ -\frac{1}{2} \langle \Omega, \hat{S} \rangle \right\} \quad (11)$$

where $\langle A, B \rangle = \text{tr}(A'B)$ denotes the trace inner product and $\hat{S} = Z'Z$.

3.1. Prior on Network Parameters

We consider the inclusion of a link in G a Bernoulli trial, $G_{ij} \sim \text{Ber}(\gamma_{ij})$, with the density

$$P(G) = \prod_{i \neq j} \gamma_{ij}^{G_{ij}} (1 - \gamma_{ij})^{(1-G_{ij})} \quad (12)$$

where $\gamma_{ij} \in (0, 1)$ is the probability of a directed link from nodes j to i . We assign to each variable inclusion a prior probability, $\gamma_{ij} = 1/2, \forall i, j$, which is equivalent to assigning an equal prior probability to all network structures, i.e., $P(G) \propto 1$.

Following Geiger and Heckerman (2002), we assume that unconstrained Ω , given a complete network graph, is Wishart distributed $\Omega|G \sim \mathcal{W}(S_0^{-1}, \nu)$, with the density

$$P(\Omega|G) = \frac{1}{K_m(\nu, S_0)} |\Omega|^{\frac{(\nu-m-1)}{2}} \exp \left\{ -\frac{1}{2} \langle \Omega, S_0 \rangle \right\} \quad (13)$$

where S_0 is the prior scale matrix, $\nu > m + 1$ is the degree-of-freedom parameter, and $K_m(\nu, S_0)$ is the normalizing constant.

3.2. Posterior Inference of Network Parameters

Under the Bayesian framework of [Geiger and Heckerman \(2002\)](#), Ω can be integrated out analytically to obtain a marginal likelihood function with a closed-form local expression given by the following ([Ahelegbey et al., 2016a, 2016b](#)):

$$P(Z_{y_i}|Z_{x_i}) = \frac{\pi^{-\frac{1}{2}N} \nu^{\frac{1}{2}\nu}}{(\nu + N)^{\frac{1}{2}(\nu+N)}} \frac{\Gamma\left(\frac{\nu+N-n_f}{2}\right)}{\Gamma\left(\frac{\nu-n_f}{2}\right)} \left(\frac{|Z'_{x_i}Z_{x_i} + \nu I_{n_x}|}{|Z'_{f_i}Z_{f_i} + \nu I_{n_f}|} \right)^{\frac{1}{2}(\nu+N)} \quad (14)$$

where Z_{y_i} is the i -th dependent variable, Z_{x_i} is a set of predictor variables of Z_{y_i} , $Z_{f_i} = (Z_{y_i}, Z_{x_i})$, n_x is the number of nodes in Z_{x_i} , and $n_f = n_x + 1$ is the number of nodes in Z_{f_i} .

The MCMC algorithm devised by [Ahelegbey et al. \(2016a\)](#) and extended by [Ahelegbey and Giudici \(2022\)](#) samples from the conditional posterior distributions of (G_0, G_+) and the parameters, $(\Phi_+, \Phi_0, \Sigma_\varepsilon, \Sigma_u)$, of the GSVAR in each iteration. Each iteration of the algorithm samples from the following conditional posteriors:

1. Sample via a Metropolis-within-Gibbs $[G_0, G_+|Z]$ by the following:
 - (a) Sampling G_+ from the marginal distribution: $[G_+|Z]$.
 - (b) Sampling G_0 from the conditional distribution: $[G_0|Z, G_+]$.
2. Estimate $(\Phi_+, \Phi_0, \Sigma_\varepsilon, \Sigma_u)$ given $(\hat{G}_0, \hat{G}_+)$ and Z following the conventional normal-inverse Wishart prior distribution for a seemingly unrelated regression (SUR). The posterior expectations of $(\Phi_+, \Phi_0, \Sigma_u)$ and Σ_ε are as follows:

$$\begin{aligned} \hat{\Phi}_{+, y_i x_i} &= (I_{|x_i|} + \hat{\sigma}_{u_i}^{-2} Z'_{x_i} Z_{x_i})^{-1} \hat{\sigma}_{u_i}^{-2} Z'_{x_i} Z_{y_i} \\ \hat{U}_{y_i} &= Z_{y_i} - Z_{x_i} \hat{\Phi}'_{+, y_i x_i} \\ \hat{\Phi}_{0, y_i \pi_x} &= (I_{|\pi_x|} + \hat{\sigma}_{\varepsilon_i}^{-2} \hat{U}'_{\pi_x} \hat{U}_{\pi_x})^{-1} \hat{\sigma}_{\varepsilon_i}^{-2} \hat{U}'_{\pi_x} \hat{U}_{y_i} \\ \hat{\Sigma}_u &= \frac{1}{\delta_0 + N} [S_0 + \hat{U}' \hat{U}] \\ \hat{\Sigma}_\varepsilon &= (I_n - \hat{\Phi}_0) \hat{\Sigma}_u (I_n - \hat{\Phi}_0)' \end{aligned}$$

where $\hat{\Phi}_{+, y_i x_i}$ is the y_i -th row and x_i -columns of $\hat{\Phi}_+$, $|x_i|$ is that number of elements in x_i , $Z_{y_i} \in Z$ corresponds the dependent variable, and Y_i and $Z_{x_i} \in Z$ are the lag predictors of Y_i according to G_+ . $\hat{U}_{y_i}$ is the estimate of the reduced-form error for the y_i -th equation. $\hat{\Phi}_{0, y_i \pi_x}$ is the y_i -th row and π_x -columns of $\hat{\Phi}_0$, $\hat{U}_{\pi_x} \in \hat{U}$ corresponds to the predictors of $\hat{U}_{y_i}$ according to G_0 , and $V_{B, \pi}$ is a diagonal matrix of the prior covariance. The posterior of Σ_u^{-1} is Wishart distributed with $\delta + N$ degrees of freedom.

The convergence of the Gibbs sampler is assessed using standard diagnostics such as trace plots, the Gelman–Rubin diagnostic ([Gelman & Rubin, 1992](#)), and the effective sample size. These tools help to ensure that the Markov Chain has reached its stationary distribution and that the samples are reliable. Additionally, we validated the model using out-of-sample testing and cross-validation techniques, which assessed how well the model generalized to unseen data and helped prevent overfitting. For the practical implementation of the BGSVAR model, we used Matlab 2024b statistical software. Given the computational challenges that arise with high-dimensional data, we addressed these issues by optimizing the code and utilizing parallel processing where applicable. This ensured the convergence of the Gibbs sampler while effectively managing the complexity of the model.

4. Modeling Macroeconomic Time Series

This section presents a simulation study based on real macroeconomic data. The objective was to estimate impulse response for a set of US macroeconomic variables. Our approach followed the simulation framework of Korobilis (2022), where we generated multivariate time series from a VAR model using real macroeconomic data.

4.1. Data

The dataset comprised monthly observations of 15 US macroeconomic variables sourced from the Federal Reserve Economic Database (FRED) of St. Louis (<https://fred.stlouisfed.org/>, accessed on 28 January 2023). Most of these variables were also used in Cimadomo et al. (2022). The sample period spanned from January 2000 (2000M1) to December 2022 (2022M12).

To ensure stationarity, most variables were transformed into log differences, while others were expressed in first differences. A complete list of the variables included in the simulation study is provided in Table 1. The variables were ordered based on their release timing within the calendar month, following Cimadomo et al. (2022), except for total reserves and the real M1 money stock.

Table 1. Macroeconomic variables and transformation.

No.	Variable	Transformation	FRED Id	Code
1	Economic policy uncertainty	Δx_t	USEPUINDXM	EPU
2	Employment	$\Delta \log(x_t)$	PAYEMS	EMP
3	Unemployment rate	Δx_t	UNRATE	UNR
4	Avg. weekly hours	$\Delta \log(x_t)$	AWHNONAG	HRS
5	Industrial production	$\Delta \log(x_t)$	INDPRO	IP
6	CPI inflation	$\Delta \log(x_t)$	CPIAUCSL	CPI
7	Loans	$\Delta \log(x_t)$	BUSLOANS	LOA
8	Housing starts	$\Delta \log(x_t)$	HOUST	HS
9	Private consumption	$\Delta \log(x_t)$	PCE	PC
10	PCE price index	$\Delta \log(x_t)$	PCEPI	PCE
11	Real disp. personal income	$\Delta \log(x_t)$	DSPIC96	DPI
12	Fed. funds rate	Δx_t	FEDFUNDS	FF
13	Credit spread	Δx_t	BAA10YM	BAA
14	Total reserves	$\Delta \log(x_t)$	TOTRESNS	TRS
15	Real M1 money stock	$\Delta \log(x_t)$	M1REAL	M1

4.2. Model Specification

We considered both sparse and dense data-generating processes (DGPs) of the following form:

$$Y_t = \alpha + B_+ X_t + A_0 \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, I_n) \quad (15)$$

where $A_0 = (I_n - B_0)^{-1}$, $X_t = (Y'_{t-1}, \dots, Y'_{t-p})'$, α and B_+ are the intercept and slope of the regression model, and ε_t is a vector of orthogonal shocks.

To set the coefficients of our DGP, we followed these steps:

1. Estimate B_+ : Regress Y_t on X_t using OLS to obtain B_+ , and compute the residuals $U = Y - \alpha - B_+ X_t$.
2. Estimate A_0 and B_0 : Regress each residual component U_i on the remaining residuals U_{-i} (i.e., all U excluding U_i) to obtain B_0 and A_0 .

We simulated a VAR(3) process with 396 observations and estimated all competing models using the true lag length. Impulse response (IR) coefficients were computed for

24 forecast horizons ($h = 24$). To ensure robustness, we replicated the simulation 1000 times using Monte Carlo simulations with random draws from the DGP.

To evaluate the impact of sparsity and density in the coefficient matrices, we considered four different scenarios:

- (i) DGP-1: Sparse (B_+)–Sparse (A_0);
- (ii) DGP-2: Sparse (B_+)–Dense (A_0);
- (iii) DGP-3: Dense (B_+)–Sparse (A_0);
- (iv) DGP-4: Dense (B_+)–Dense (A_0).

The sparsity in B_+ and B_0 was determined by testing whether each estimated coefficient was significantly different from zero at the 5% level. Coefficients that failed this test were set to zero, ensuring that only statistically significant relationships remained in the model.

4.3. Methods Compared and Evaluation Criteria

To assess the performance of BGSVAR, we compared its impulse response estimates with those obtained from widely used alternatives in empirical macroeconomic research. Specifically, we considered BVAR and LASSO-VAR (Tibshirani, 1996). Table 2 summarizes these competing approaches.

Table 2. Summary of competing impulse response estimation methods.

#	Model	Description	Structural Assumptions
1	BGSVAR-IR	BGSVAR + IR	Sparse (B_+)–Sparse (A_0)
2	BVAR-OIR	BVAR + OIRF	Dense (B_+)–Cholesky (A_0)
3	BVAR-GIR	BVAR + GIRF	Dense (B_+)–Dense (A_0)
4	LASSO-IR	LASSO-VAR + IRF	Sparse (B_+)–Sparse (A_0)

We conducted all estimations using 20,000 Gibbs iterations to sample the network graph structure in the BGSVAR model. The BGSVAR coefficients were estimated with a normal–Wishart prior, followed by seemingly unrelated regression (SUR) estimation. The BVAR model was estimated using normal–Wishart conjugate priors, while the LASSO method selected its regularization parameter via tenfold cross-validation, minimizing mean squared cross-validation errors. All models included an intercept.

To evaluate impulse response estimates, we compared the BGSVAR, BVAR, and LASSO methods using two key metrics, the average interval length and root mean squared error (RMSE), following Kilian and Kim (2011). The average interval length measures the precision of confidence intervals, with shorter intervals indicating greater accuracy. We also computed the relative average interval length, defined as the ratio of a method's interval length to the minimum obtained across methods, where a value of 1.0 represents the best-performing model. Similarly, we report relative RMSE values to assess overall estimation accuracy.

4.4. Simulation Results

4.4.1. VAR Slope Coefficient Matrix

Before analyzing the impulse response performance of the competing methods, we first compared the VAR slope coefficient matrix (B_1) for DGP-1 with the estimates from BGSVAR, BVAR, and LASSO. Figure 1 presents these comparisons, with response variables listed along the rows and explanatory variables along the columns. The coefficients are colour-coded: red represents negative values, green indicates positive values, and white denotes zero (or near-zero) coefficients.




Figure 1. VAR (B_1) of DGP-1: Sparse (B_+)–Sparse (A_0). The response variables are shown along the rows, while explanatory variables are shown along the columns. The light green (dark green) signifies mild (very high) positive effects, while light red (dark red) denotes mild (very high) negative effects. White indicates zero or near-zero coefficients.

The structure of DGP-1 (B_1) in Figure 1 reveals that the temporal relationships among the selected macroeconomic variables are sparse, unlike the dense structure seen in BVAR (B_1). This finding supports the well-known observation that some variable sets are better modelled with information from only a few variables, while others require more variables. As expected, both the BGSVAR and LASSO methods captured the sparse structure of the VAR slope coefficients in DGP-1, whereas BVAR assumed temporal dependence among all variables, which could introduce spurious relationships with limited predictive value.

The variable selection methods, BGSVAR and LASSO, identified a smaller set of temporal relationships that closely matched the DGP. Since impulse responses were derived from VAR slope coefficients, inaccuracies in capturing the temporal dependence structure could significantly impact the accuracy of impulse response estimates.

4.4.2. Shock Matrix

Figure 2 compares the shock matrix A_0 of DGP-1 with those estimated from BGSVAR, BVAR-OIR, and LASSO. The shock matrix for BVAR-GIR is excluded from Figure 2 due to its fully dense structure. As expected, when the shock matrix in the DGP was sparse, using an orthogonalized shock matrix (such as BVAR-OIR) failed to accurately capture the true contemporaneous relationships among the variables. In contrast, both BGSVAR and LASSO estimated a shock matrix that more closely aligned with the sparse structure of DGP-1.




Figure 2. Shock matrix of DGP-1: Sparse (B_+)–Sparse (A_0). Response (explanatory) variables are in the rows (columns). The light green (dark green) signifies mild (very high) positive effects, while light red (dark red) denotes mild (very high) negative effects. White indicates zero or near-zero coefficients.

A detailed comparison of the BGSVAR and LASSO matrices shows that the BGSVAR provided a more sparse structure, which better matches the DGP. For example, the DGP reveals that a shock to unemployment (UNR) only significantly affects loans (LOA) and has no contemporaneous effect on the other variables. However, the BGSVAR matrix in Figure 2b suggests that a shock to the UNR also impacts CPI, LOA, PCE, and M1. Meanwhile, the LASSO matrix in Figure 2d shows that the UNR has a significant contemporaneous effect on almost all other variables, except EPU and BAA, which is similar to the BVAR-OIR shock matrix in Figure 2c (with the inclusion of EMP). While the LASSO method shares some characteristics with BVAR-OIR, the BGSVAR model produces a more accurate representation of the DGP's structure, particularly in its more sparse shock matrix.

4.4.3. Impulse Response Performance

In this section, we evaluate the impulse response performance of the competing methods by comparing their results to the true impulse responses (IRs) generated by the DGP. For brevity, we focus on the responses of the unemployment rate (UNR), industrial production (IP), and inflation (CPI) to a shock to the federal funds rate (FF).

Figure 3 presents the dynamic responses of unemployment (UNR), industrial production (IP), and inflation (CPI) to a contractionary monetary policy shock (FF) using different competing VAR-based models. The black dashed lines represent the true impulse responses from the DGP. After the monetary tightening shock, we observe that the unemployment rate decreases across most of the horizon, while industrial production and inflation increase. Specifically, the unemployment rate responds significantly to the contractionary shock, while industrial production and inflation show notable increases during the short- and medium-term horizons.




Figure 3. The response of the UNR, IP, and CPI to FF based on 1000 trials from DGP-1: Sparse (B_+)–Sparse (A_0). The x-axis represents the time horizon (lags after the shock), and the y-axis shows the response magnitude of the variable. The black dashed lines represent the true DGP IRs, the green solid lines represent the median IR estimates, and the shaded areas represent the 95% confidence bands. The red line marks the zero level.

Figures 4–6 present similar plots for DGP-2, DGP-3, and DGP-4, respectively. A one-unit shock to FF corresponds to a one-standard deviation shock in monetary policy. In the plots, the green solid lines indicate the median estimates, the black dashed lines represent the true values from the DGP, and the shaded areas represent the 95% confidence bands. A standard interpretation is that if the 95% confidence bands always encompass the true impulse responses, estimation precision is considered satisfactory. As seen in Figures 3 and 4, all methods produce relatively precise estimates. However, in Figures 5 and 6, both the BVAR-GIR and LASSO-IR demonstrate unsatisfactory performance, particularly at shorter horizons.

Although the competing methods perform well in Figure 3, some methods—particularly the BVAR-GIR and LASSO-IR—show noticeable deviations from the true values, especially at shorter horizons. These deviations highlight the importance of correct temporal dependence structures for accurate impulse response estimation.




Figure 4. The response of the UNR, IP, and CPI to FF based on 1000 trials from DGP-2: Sparse (B_+)–Dense (A_0). Note: The DGP IRs are shown in black dashed lines, the green solid lines are the median IRs, and the shaded areas are the confidence bands. The red line marks the zero level.


Figure 7 reports the average interval length of the joint pointwise impulse response of the competing methods for the different DGPs. The figures show that the LASSO-IR records the minimum average interval length which means that it has much thinner confidence bands than the rest of the models. This is followed by the BGSVAR-IR. We, however, notice that in the case of DGP-3, the BGSVAR-IR reports the second narrowest confidence band for short ($h < 6$) and medium horizons ($h \leq 12$). In all the cases shown by the figure, the BVAR-OIR and BVAR-GIR appear to record much wider confidence bands than their counterparts.

Table 3 shows the numeric summary of Figure 7 in terms of the relative average interval for $h = 6$ (short), $h = 12$ (medium), and $h = 24$ (long horizon). The table confirms the above observation that the LASSO-IR consistently recorded the thinnest confidence bands of the impulse response estimates, followed closely by the BGSVAR-IR.

Figure 8 reports the IR RMSEs of the competing methods under the different DGPs. The figure shows that the BGSVAR-IR estimators are generally more efficient than the competing methods. Across all the estimands, the smallest RMSEs are typically achieved by the BGSVAR-IR, followed by the LASSO-IR, and then the BVAR-OIR. The BVAR-GIR is the least preferred. Table 4 presents the numeric summary of the plots in Figure 8 in terms of the relative RMSE for $h = 6$ (short), $h = 12$ (medium), and $h = 24$ (long horizon). The table shows that the BGSVAR-IR is more accurate in producing pointwise impulse response at short to medium horizons. It is closely followed by the BVAR-IR and the LASSO-IR.

Table 3. Relative average interval length for $h = 6$ (short), $h = 12$ (medium), and $h = 24$ (long horizon).

	h	BGSVAR-IR	BVAR-OIR	BVAR-GIR	LASSO-IR
DGP-1	$h = 6$	1.4758	2.2833	2.0968	1
	$h = 12$	1.7044	2.6109	2.5530	1
	$h = 24$	2.4368	3.4333	3.5384	1
DGP-2	$h = 6$	1.2719	2.0683	1.8724	1
	$h = 12$	1.2974	2.122	2.1174	1
	$h = 24$	1.5215	2.3536	2.6283	1
DGP-3	$h = 6$	1.3558	2.0467	1.6591	1
	$h = 12$	1.8660	1.9078	1.9866	1
	$h = 24$	4.4813	1.8479	2.0159	1
DGP-4	$h = 6$	1.2959	1.9693	1.6462	1
	$h = 12$	1.5176	2.0208	1.8672	1
	$h = 24$	1.7328	1.8186	1.8118	1

**Figure 5.** The response of the UNR, IP, and CPI to FF based on 1000 trials from DGP-3: Dense (B_+)–Sparse (A_0). Note: The DGP IRs are shown in black dashed lines, the green solid lines are the median IRs, and the shaded areas are the confidence bands. The red line marks the zero level.




Figure 6. The response of the UNR, IP, and CPI to FF based on 1000 trials from DGP-4: Dense (α, B_+) –Dense (A_0) . Note: The DGP IRs are shown in black dashed lines, the green solid lines are the median IRs, and the shaded areas are the confidence bands. The red line marks the zero level.




Figure 7. Average interval lengths of joint pointwise impulse response of competing methods.




Figure 8. RMSE values for each method across the 1000 Monte Carlo iterations.

Table 4. Relative RMSE for $h = 6$ (short horizon), $h = 12$ (medium horizon), and $h = 24$ (long horizon).

	h	BGSVAR-IR	BVAR-OIR	BVAR-GIR	LASSO-IR
DGP-1	$h = 6$	1	1.3950	1.7341	1.3277
	$h = 12$	1	1.3108	1.7150	1.4481
	$h = 24$	1	1.0909	1.2723	1.5377
DGP-2	$h = 6$	1	1.0721	1.7214	1.2878
	$h = 12$	1.1327	1	1.6255	1.1929
	$h = 24$	1.0967	1	1.0162	1.2814
DGP-3	$h = 6$	1	2.4003	5.1577	2.6716
	$h = 12$	1	1.9571	4.2454	2.2892
	$h = 24$	1.1822	1	1.4845	1.5540
DGP-4	$h = 6$	1	1.7935	2.9006	2.2331
	$h = 12$	1	1.4542	2.5316	1.7729
	$h = 24$	1.0865	1	1.2464	1.2447

In summary, we find evidence that BGSVAR achieves higher pointwise impulse response predictive accuracy and fits the structural dynamics of VAR-based models better than the LASSO and BVAR.

5. Conclusions

Impulse response functions (IRFs) in VAR models have long been used in the literature to measure the direct and indirect propagation effects of shocks. It is well-established that impulse response estimates are non-linear functions of both autoregressive and contemporaneous coefficient matrices. Thus, accurate estimation of these matrices is essential for efficient impulse response estimation. This study focused on identifying structural relationships among endogenous variables using a Bayesian graphical SVAR (BGSVAR) framework. By inferring both lagged and contemporaneous dependencies among the variables, we were able to estimate the coefficient matrices required for accurate impulse response analysis.

Through empirical simulation experiments on a well-known macroeconomic dataset, we demonstrated that the BGSVAR impulse response estimator is more robust compared to traditional methods such as orthogonalized impulse response (OIR), generalized impulse response (GIR), and LASSO impulse response (LASSO-IR). Specifically, we applied the method to model the impulse responses of shocks to key US macroeconomic time series in a moderate-dimensional VAR model, as presented in [Cimadomo et al. \(2022\)](#). Our results

showed that LASSO-IR tended to generate less accurate pointwise confidence intervals for impulse responses on average compared to the BVAR-OIR, while BGSVAR outperformed all the competing methods by providing more precise and reliable estimates.

The results indicate that the BGSVAR model outperforms traditional methods like LASSO and BVAR in estimating impulse responses, particularly in terms of accuracy and robustness. This improvement can be attributed to the model's ability to incorporate both temporal and contemporaneous causal structures, which traditional models often fail to adequately address. As such, BGSVAR offers a more reliable framework for understanding the dynamic interactions among macroeconomic variables, with implications for better policy analysis.

The empirical findings underscore a significant pitfall in using conventional VAR impulse response techniques. Specifically, inferences made using unrestricted VAR slope coefficient matrices, combined with orthogonalized or generalized shock matrices, may yield pointwise impulse response confidence intervals that are often less accurate than those from sparse VAR-based models. This result is not surprising, as assuming unrestricted temporal dependence among all variables in a moderate-to-large VAR model is often unrealistic in practice. Additionally, the method of imposing restrictions on contemporaneous shock matrices based on "strongly held convictions" is subjective and only valid when such convictions are substantiated by data. Our analysis further confirms that the generalized impulse response approach, which assumes full contemporaneous and lagged dependence, may not be suitable in many real-world settings.

Limitations and Future Work

While the BGSVAR model offers notable advantages in estimating impulse responses, several limitations should be considered. First, the model's performance is dependent on the prior distribution and hyperparameter selection. Although we used a uniform prior for the network structure to allow the data to inform the relationships, this approach may not always be the most suitable for cases where prior knowledge or expert judgment could significantly improve the model. Future work could investigate the impact of different prior distributions, particularly in situations where some prior information is available, such as the relationships between key macroeconomic variables.

Second, the computational complexity of the MCMC-based sampling approach used in BGSVAR can be a limitation, especially when working with large datasets. Although we conducted convergence diagnostics to ensure the reliability of the estimates, the computational burden remains significant, particularly for high-dimensional models. Future research could explore more efficient sampling algorithms, such as Hamiltonian Monte Carlo (HMC) or Variational Bayes, to improve the computational efficiency without sacrificing model accuracy.

Finally, while we focused on impulse response analysis within the context of a moderate-dimensional VAR model, the scalability of BGSVAR to larger datasets and higher-dimensional models remains an area for further exploration. As datasets grow, the model may face challenges related to memory usage and computational time, which could limit its practical applicability in high-frequency or large-scale macroeconomic forecasting. Future studies could investigate ways to optimize the model's performance for larger datasets, such as leveraging parallel computing or more scalable variational techniques.

Funding: This research received no external funding.

Data Availability Statement: The data used in the paper is publicly available at <https://fred.stlouisfed.org/>.

Conflicts of Interest: The authors declare no conflicts of interest.

References

- Ahelegbey, D. F., Billio, M., & Casarin, R. (2016a). Bayesian graphical models for structural vector autoregressive processes. *Journal of Applied Econometrics*, 31(2), 357–386. [\[CrossRef\]](#)
- Ahelegbey, D. F., Billio, M., & Casarin, R. (2016b). Sparse graphical vector autoregression: A bayesian approach. *Annals of Economics and Statistics*, 123/124, 333–361. [\[CrossRef\]](#)
- Ahelegbey, D. F., & Giudici, P. (2022). NetVIX—A network volatility index of financial markets. *Physica A: Statistical Mechanics and Its Applications*, 594, 127017. [\[CrossRef\]](#)
- Ahelegbey, D. F., Giudici, P., & Hashem, S. Q. (2021). Network VAR models to measure financial contagion. *The North American Journal of Economics and Finance*, 55, 101318. [\[CrossRef\]](#)
- Bai, J., & Ng, S. (2002). Determining the number of factors in approximate factor models. *Econometrica*, 70(1), 191–221. [\[CrossRef\]](#)
- Banbura, M., Giannone, D., & Reichlin, L. (2010). Large bayesian vector autoregressions. *Journal of Applied Econometrics*, 25, 71–92. [\[CrossRef\]](#)
- Bassetti, F., Casarin, R., & Leisen, F. (2014). Beta-product dependent pitman–Yor processes for bayesian inference. *Journal of Econometrics*, 180(1), 49–72. [\[CrossRef\]](#)
- Basu, S., & Michailidis, G. (2015). Regularized estimation in sparse high-dimensional time series models. *The Annals of Statistics*, 43(4), 1535–1567.
- Bernanke, B. S., Boivin, J., & Elias, P. S. (2005). Measuring the effects of monetary policy: A factor-augmented vector autoregressive (FAVAR) approach. *The Quarterly Journal of Economics*, 120(1), 387–422. [\[CrossRef\]](#)
- Bernanke, B. S., & Mihov, I. (1998). Measuring monetary policy. *The Quarterly Journal of Economics*, 113(3), 869–902.
- Billio, M., Casarin, R., & Rossini, L. (2019). Bayesian nonparametric sparse var models. *Journal of Econometrics*, 212(1), 97–115. [\[CrossRef\]](#)
- Blanchard, O. J., & Quah, D. (1989). The dynamic effects of aggregate demand and supply disturbances. *The American Economic Review*, 79(4), 655–673.
- Canova, F., & Ciccarelli, M. (2004). Forecasting and turning point predictions in a bayesian panel VAR model. *Journal of Econometrics*, 120(2), 327–359. [\[CrossRef\]](#)
- Canova, F., & De Nicolo, G. (2002). Monetary disturbances matter for business fluctuations in the G-7. *Journal of Monetary Economics*, 49(6), 1131–1159. [\[CrossRef\]](#)
- Casarin, R., Iacopini, M., Molina, G., Ter Horst, E., Espinasa, R., Sucre, C., & Rigobon, R. (2020). Multilayer network analysis of oil linkages. *The Econometrics Journal*, 23(2), 269–296. [\[CrossRef\]](#)
- Celani, A., Pagnottoni, P., & Jones, G. (2024). Bayesian variable selection for matrix autoregressive models. *Statistics and Computing*, 34(2), 91. [\[CrossRef\]](#)
- Choi, C.-Y., & Chudik, A. (2019). Estimating impulse response functions when the shock series is observed. *Economics Letters*, 180, 71–75. [\[CrossRef\]](#)
- Christiano, L. J., Eichenbaum, M., & Evans, C. L. (1999). Monetary policy shocks: What have we learned and to what end? In J. B. Taylor, & M. Woodford (Eds.), *Handbook of macroeconomics* (Vol. 1, pp. 65–148). Elsevier.
- Cimadomo, J., Giannone, D., Lenza, M., Monti, F., & Sokol, A. (2022). Nowcasting with large bayesian vector autoregressions. *Journal of Econometrics*, 231(2), 500–519. [\[CrossRef\]](#)
- Corander, J., & Villani, M. (2006). A Bayesian approach to modelling graphical vector autoregressions. *Journal of Time Series Analysis*, 27(1), 141–156. [\[CrossRef\]](#)
- Davis, R. A., Zang, P., & Zheng, T. (2016). Sparse vector autoregressive modeling. *Journal of Computational and Graphical Statistics*, 25(4), 1077–1096. [\[CrossRef\]](#)
- Demiralp, S., & Hoover, K. D. (2003). Searching for the causal structure of a vector autoregression. *Oxford Bulletin of Economics and Statistics*, 65, 745–767. [\[CrossRef\]](#)
- Diebold, F., & Yilmaz, K. (2014). On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics*, 182(1), 119–134. [\[CrossRef\]](#)
- Diebold, F. X., & Yilmaz, K. (2012). Better to give than to receive: Predictive directional measurement of volatility spillovers. *International Journal of Forecasting*, 28(1), 57–66. [\[CrossRef\]](#)
- Eichenbaum, M., & Evans, C. L. (1995). Some empirical evidence on the effects of shocks to monetary policy on exchange rates. *The Quarterly Journal of Economics*, 110(4), 975–1009. [\[CrossRef\]](#)
- Eichler, M. (2007). Granger causality and path diagrams for multivariate time series. *Journal of Econometrics*, 137(2), 334–353. [\[CrossRef\]](#)
- Faust, J. (1998). The robustness of identified VAR conclusions about money. In *Carnegie-rochester conference series on public policy* (Vol. 49, pp. 207–244). North-Holland.
- Forni, M., Hallin, M., Lippi, M., & Reichlin, L. (2005). The generalized dynamic factor model: One-Sided estimation and forecasting. *Journal of the American Statistical Association*, 100(471), 830. [\[CrossRef\]](#)

- Gefang, D. (2014). Bayesian doubly adaptive elastic-net lasso for VAR shrinkage. *International Journal of Forecasting*, 30(1), 1–11. [\[CrossRef\]](#)
- Geiger, D., & Heckerman, D. (2002). Parameter priors for directed acyclic graphical models and the characterization of several probability distributions. *Annals of Statistics*, 30(5), 1412–1440. [\[CrossRef\]](#)
- Gelman, A., & Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences, (with discussion). *Statistical Science*, 7, 457–511. [\[CrossRef\]](#)
- Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1), 161–182. [\[CrossRef\]](#)
- Jordà, Ò. (2023). Local projections for applied economics. *Annual Review of Economics*, 15(1), 607–631. [\[CrossRef\]](#)
- Jordà, Ò., & Taylor, A. M. (2025). Local projections. *Journal of Economic Literature*, 63(1), 59–110.
- Kalli, M., & Griffin, J. E. (2018). Bayesian nonparametric vector autoregressive models. *Journal of Econometrics*, 203(2), 267–282. [\[CrossRef\]](#)
- Karlsson, S. (2013). Forecasting with bayesian vector autoregressions. In G. Elliott, & A. Timmermann (Eds.), *Handbook of economic forecasting* (Vol. 2, pp. 689–1324). Elsevier.
- Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3), 1053–1069. [\[CrossRef\]](#)
- Kilian, L. (2013). Structural vector autoregressions. In N. Hashimzade, & M. Thornton (Eds.), *Handbook of research methods and applications in empirical macroeconomics* (pp. 515–554). Edward Elgar Publishing.
- Kilian, L., & Kim, Y. J. (2011). How reliable are local projection estimators of impulse responses? *Review of Economics and Statistics*, 93(4), 1460–1466. [\[CrossRef\]](#)
- King, R. G., Plosser, C. I., Stock, J. H., & Watson, M. W. (1991). Stochastic trends and economic fluctuations. *American Economic Review*, 81, 819–840.
- Kock, A. B., & Callot, L. (2015). Oracle inequalities for high dimensional vector autoregressions. *Journal of Econometrics*, 186(2), 325–344. [\[CrossRef\]](#)
- Koop, G., & Korobilis, D. (2016). Model uncertainty in panel vector autoregressive models. *European Economic Review*, 81, 115–131. [\[CrossRef\]](#)
- Koop, G., Korobilis, D., & Pettenuzzo, D. (2019). Bayesian compressed vector autoregressions. *Journal of Econometrics*, 210(1), 135–154. [\[CrossRef\]](#)
- Koop, G., Pesaran, M. H., & Potter, S. M. (1996). Impulse response analysis in nonlinear multivariate models. *Journal of Econometrics*, 74(1), 119–147. [\[CrossRef\]](#)
- Korobilis, D. (2013). VAR forecasting using bayesian variable selection. *Journal of Applied Econometrics*, 28(2), 204–230. [\[CrossRef\]](#)
- Korobilis, D. (2022). A new algorithm for structural restrictions in bayesian vector autoregressions. *European Economic Review*, 148, 104241. [\[CrossRef\]](#)
- Li, D., Plagborg-Møller, M., & Wolf, C. K. (2024). Local projections vs. VARs: Lessons from thousands of DGPs. *Journal of Econometrics*, 244(2), 105722. [\[CrossRef\]](#)
- Lopes, H. F., & Carvalho, C. M. (2007). Factor stochastic volatility with time varying loadings and markov switching regimes. *Journal of Statistical Planning and Inference*, 137(10), 3082–3091. [\[CrossRef\]](#)
- Medeiros, M. C., & Mendes, E. F. (2016). L1-regularization of high-dimensional time-series models with non-gaussian and heteroskedastic errors. *Journal of Econometrics*, 191(1), 255–271. [\[CrossRef\]](#)
- Pesaran, H. H., & Shin, Y. (1998). Generalized impulse response analysis in linear multivariate models. *Economics Letters*, 58(1), 17–29. [\[CrossRef\]](#)
- Pesaran, M. H., Schuermann, T., & Weiner, S. M. (2004). Modelling regional interdependencies using a global error-correcting macroeconomic model. *Journal of Business and Economic Statistics*, 22(2), 129–162. [\[CrossRef\]](#)
- Plagborg-Møller, M., & Wolf, C. K. (2021). Local projections and VARs estimate the same impulse responses. *Econometrica*, 89(2), 955–980. [\[CrossRef\]](#)
- Sims, C. A. (1980). Macroeconomics and reality. *Econometrica, Econometric Society*, 48(1), 1–48. [\[CrossRef\]](#)
- Stock, J. H., & Watson, M. W. (2002). Forecasting using principal components from a large number of predictors. *Journal of the American Statistical Association*, 97(460), 1167–1179. [\[CrossRef\]](#)
- Stock, J. H., & Watson, M. W. (2018). Identification and estimation of dynamic causal effects in macroeconomics using external instruments. *The Economic Journal*, 128(610), 917–948. [\[CrossRef\]](#)
- Suppes, P. (1970). *A probabilistic theory of causality*. North-Holland Pub. Co.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1), 267–288. [\[CrossRef\]](#)

- Uhlig, H. (2005). What are the effects of monetary policy on output? Results from an agnostic identification procedure. *Journal of Monetary Economics*, 52(2), 381–419. [[CrossRef](#)]
- Zou, H., & Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2), 301–320. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.