



Enforcing Stationarity through the Prior in Vector Autoregressions

Sarah E. Heaps

To cite this article: Sarah E. Heaps (2023) Enforcing Stationarity through the Prior in Vector Autoregressions, Journal of Computational and Graphical Statistics, 32:1, 74-83, DOI: 10.1080/10618600.2022.2079648

To link to this article: <https://doi.org/10.1080/10618600.2022.2079648>



© 2022 The Author(s). Published with license by Taylor & Francis Group, LLC.



[View supplementary material](#)



Published online: 23 Jun 2022.



[Submit your article to this journal](#)



Article views: 771



[View related articles](#)



[View Crossmark data](#)



Citing articles: 1 [View citing articles](#)

Enforcing Stationarity through the Prior in Vector Autoregressions

Sarah E. Heaps

Department of Mathematical Sciences, Durham University, Durham, UK

ABSTRACT

Stationarity is a very common assumption in time series analysis. A vector autoregressive process is stable if and only if the roots of its characteristic equation lie outside the unit circle, constraining the autoregressive coefficient matrices to lie in the stationary region. However, the stationary region has a highly complex geometry which impedes specification of a prior distribution. In this work, an unconstrained reparameterization of a stationary vector autoregression is presented. The new parameters are partial autocorrelation matrices, which are interpretable, and can be transformed bijectively to the space of unconstrained square matrices through a simple mapping of their singular values. This transformation preserves various structural forms of the partial autocorrelation matrices and readily facilitates specification of a prior. Properties of this prior are described along with an important special case which is exchangeable with respect to the order of the elements in the observation vector. Posterior inference and computation are described and implemented using Hamiltonian Monte Carlo via Stan. The prior and inferential procedures are illustrated with an application to a macroeconomic time series which highlights the benefits of enforcing stationarity and encouraging shrinkage toward a sensible parametric structure. Supplementary materials for this article are available online.

ARTICLE HISTORY

Received July 2021
Accepted May 2022

KEYWORDS

Partial autocorrelation matrix; Stan; Unconstrained reparameterization; Vector autoregressive model

1. Introduction

Denote by $\{y_t\}$ a time series of equally spaced m -variate observations. A stochastic process is said to be strictly stationary if its properties are unaffected by a shift in the time origin and weakly stationary if the mean $E(y_t) = \mu$ remains constant over time and the autocovariance function $\Gamma_i = \text{cov}(y_t, y_{t+i}) = E\{(y_t - \mu)(y_{t+i} - \mu)^T\}$ depends only on the lag i ($i = 0, 1, \dots$), with $\Gamma_{-i} = \Gamma_i^T$. For Gaussian processes, the two are equivalent, and we simply refer to a process as being stationary. As raw time series often exhibit periodic variation and systematic changes in the mean, stationarity is generally assumed only for the residuals of a detrended series, the variables of a differenced process, or those latent components of a state space model that are believed to be mean-reverting. In these cases, stationarity prevents the predictive variance of the transformed process from growing without bound as the forecast horizon increases, moving either forward or backward in time. This is a very reasonable assumption in many applications, for instance where the inferential objective is long-term forecasting or characterizing the long-run behavior of linear dynamic systems.

All stationary Gaussian processes can be approximated arbitrarily well by a finite-order, vector autoregressive moving average (VARMA) model (Neusser 2016, chap. 12). Although we focus on the more commonly used subclass of vector autoregressive (VAR) models, we discuss the extension to the more general case in Section 6. Without loss of generality, assume that the time series $\{y_t\}$ can be modeled as a zero-mean, order-

p vector autoregressive, or $\text{VAR}_m(p)$, process, $y_t = \phi_1 y_{t-1} + \dots + \phi_p y_{t-p} + \epsilon_t$, in which the errors $\{\epsilon_t\}$ form a sequence of uncorrelated, zero-mean multivariate normal random vectors, $\epsilon_t \sim N_m(\mathbf{0}, \Sigma)$. The parameters of the model therefore, comprise the autoregressive coefficient matrices $\phi_i \in M_{m \times m}(\mathbb{R})$ ($i = 1, \dots, p$), and the error variance matrix $\Sigma \in \mathcal{S}_m^+$, where $M_{m \times n}(V)$ denotes the set of $m \times n$ matrices with entries in V and \mathcal{S}_m^+ denotes the set of $m \times m$ symmetric, positive definite matrices. Henceforth, we refer to the collection of ϕ_i as Φ . VAR processes are a widely used class of time-series model which find application in a broad range of fields, such as macroeconomics (Koop and Korobilis 2009), neuroscience (Chiang et al. 2017), and genomics (Abegaz and Wit 2013). In most cases, priors for (Φ, Σ) are based on either the conjugate matrix normal inverse Wishart distribution (Bańbura, Giannone, and Reichlin 2010) or the semi-conjugate variant in which Φ and Σ are independent a priori (Karlsson 2013). Variations include priors that fix the error variance at an estimate, such as the Minnesota prior (Doan, Litterman, and Sims 1984), or hierarchical priors which are typically designed to allow data-informed shrinkage toward a sparse parameterization (Ghosh, Khare, and Michailidis 2019). Another class of priors arises in the context of graphical VARs where the idea is to learn the contemporaneous and lagged relationships over time (Corander and Villani 2005; Paci and Consonni 2020). In this case, a prior is first placed on the graph underpinning the process and then a (typically vague) prior is placed on the parameters that are unrestricted after conditioning on that graph. In this literature, although

stationarity is often stated as an assumption to derive properties of the process, it is not enforced as a constraint.

It is common to write the $\text{VAR}_m(p)$ process as $\epsilon_t = (I_m - \phi_1 B - \dots - \phi_p B^p)y_t = \phi(B)y_t$ where B is the backshift operator, that is $By_t = y_{t-1}$, I_m is the $m \times m$ identity matrix and $\phi(u) = I_m - \phi_1 u - \dots - \phi_p u^p$, $u \in \mathbb{C}$, is termed the characteristic polynomial. The process is stable if and only if all the roots of $\det\{\phi(u)\} = 0$ lie outside the unit circle. Because all stable processes are stationary and unstable stationary processes are not generally of interest, this subset of $M_{m \times m}(\mathbb{R})^p$ is often referred to as the stationary region, denoted by $\mathcal{C}_{p,m}$.

When $m = 1$ and $p = 1$ or $p = 2$, the stationary region is simple, with $\mathcal{C}_{1,1}$ representing the interval $(-1, 1)$, and $\mathcal{C}_{2,1}$ a triangle in the (ϕ_1, ϕ_2) -plane. However, increasing m or p further increases the complexity of the polynomial equation $\det\{\phi(u)\} = 0$ and hence the geometry of the stationary region. This causes two main problems for Bayesian inference. First, because there are no standard distributions on $\mathcal{C}_{p,m}$, it is not generally possible to directly specify a prior over this space which encodes genuine beliefs. Second, designing efficient Markov chain Monte Carlo (MCMC) samplers with state space constrained to $\mathcal{C}_{p,m}$ is very challenging. In the univariate ($m = 1$) case, Chib (1993) addresses the latter difficulty by assigning a multivariate normal prior to Φ , truncated to the region $\mathcal{C}_{p,1}$. For $p > 1$, the normalizing constant cannot be evaluated in closed form and so the prior density is known only up to proportionality. A sampler is described which updates Φ in a block in which the proposal is the full conditional distribution when the constraints are ignored; proposals that fall outside of $\mathcal{C}_{p,1}$ are rejected in the Metropolis acceptance step. However, this will be efficient only if the proposal density is concentrated over $\mathcal{C}_{p,1}$. In the univariate case, Piccolo (1982) calculates the volume of the stationary region for the parameters of autoregressive (AR) models. In the special case where the autoregressive coefficient matrices ϕ_i are diagonal, a trivial corollary of this result is that the volume of the stationary region is equal to U_p^m , in which $U_p = (M_1 M_3 \dots M_{p-1})^2$ for p even and $U_p = U_{p-1} M_p$ for p odd, where $M_i = 2^i [(i-1)/2]! / i!$. This region becomes vanishingly small as p increases and is likely to render an inferential scheme that is not tailored to the geometry of the problem highly inefficient. A natural way around both of these problems is to find an interpretable reparameterization of the $\text{VAR}_m(p)$ model which maps $\Phi \in \mathcal{C}_{p,m}$ to a space which has simpler constraints and allows prior specification to be carried out in a meaningful way.

As a result of their ubiquity in time series analysis, there is a large literature on reparameterizations of univariate AR models. Barndorff-Nielsen and Schou (1973) establish a bijection between $\Phi \in \mathcal{C}_{p,1}$ and the first p partial autocorrelations $\rho = (\rho_1, \dots, \rho_p)$ of a stationary, p th order autoregression. Monahan (1984) provides an alternative derivation of the mapping and explicit recursive formulas for its inverse. The new parameterization has the advantage that the partial autocorrelations ρ are interpretable and only constrained to lie in the Cartesian product space $(-1, 1)^p$. Marriott et al. (1996) and Barnett, Kohn, and Sheather (1996) present prior distributions for the partial autocorrelations and MCMC algorithms for computational inference. The latter suggest a uniform prior over $\mathcal{C}_{p,1}$ which

induces a closed form for the density of ρ due to the analytical tractability of the Jacobian term. The former use spike-and-slab priors for the ρ_i in which the slab is uniform over $(-1, 1)$ and the spike is an atom of probability at zero, which allows for uncertainty in the model order. Marriott and Smith (1992) and Huerta and West (1999) describe reparameterizations based on a representation of the characteristic equation in factorized form, that is, $\phi(u) = \prod_{i=1}^p (1 - \eta_i u)$. In this case, the condition for stationarity reduces to $|\eta_i| < 1$ ($i = 1, \dots, p$). Huerta and West (1999) also allow uncertainty in the model order p , and the balance of real and complex (reciprocal) roots η_i , by placing priors on the real roots and the moduli of the pairs of complex roots with atoms of probability at zero.

Extensions of these reparameterizations of univariate AR models to the general vector case, especially with a focus on prior specification, are surprisingly scarce in the literature. Huerta and Prado (2006) extend the ideas of Huerta and West (1999), but only in the special case where the ϕ_i are diagonal. He, She, and Wu (2013) describe a sparsity-inducing penalized maximum likelihood algorithm for frequentist model-fitting. However, it is not fully flexible, constraining inference to a subset of the stationary region where the spectral norm, rather than spectral radius, of the companion matrix is less than one. Morf, Vieira, and Kailath (1978) generalize the results of Barndorff-Nielsen and Schou (1973); for every error variance matrix Σ , a bijection is established between $\Phi \in \mathcal{C}_{p,m}$ and the first p partial autocorrelation matrices P_1, \dots, P_p of a stationary $\text{VAR}_m(p)$ process. Denoting by \mathcal{V}_m the subset of matrices in $M_{m \times m}(\mathbb{R})$ whose singular values are all less than one, in terms of the new parameters, the stationary region reduces to a simple Cartesian product space $(P_1, \dots, P_p) \in \mathcal{V}_m^p$. Ansley and Kohn (1986) build on the earlier work in Ansley and Newbold (1979) by generalizing the construction of Monahan (1984) and explicitly providing recursive formulas for the inverse mapping. A second bijective mapping between $(P_1, \dots, P_p) \in \mathcal{V}_m^p$ and $(A_1, \dots, A_p) \in M_{m \times m}(\mathbb{R})^p$ is described along with a maximum likelihood estimation procedure. Although the new parameters A_1, \dots, A_p are unconstrained, the reparameterization is not immediately amenable to Bayesian inference because the A_i are difficult to interpret. Roy, McElroy, and Linton (2019) derive an alternative reparameterization of stationary and invertible VARMA models. The new parameters comprise a set of p symmetric, positive definite matrices and a set of p orthogonal matrices. From an inferential perspective, the space in which the new parameters lie is almost as problematic as the stationary region $\mathcal{C}_{p,m}$. As we discuss further in Section 3.5, prior specification is difficult because the orthogonal matrices lack a meaningful interpretation, whilst computational inference remains challenging because the constraints that define the space of orthogonal matrices complicate sampling by MCMC. In this article we propose a different reparameterization and prior for the parameters of a stationary VAR process. Like in Ansley and Kohn (1986), the new parameters are constructed through two mappings: first to a set of (interpretable) partial autocorrelation matrices P_1, \dots, P_p and then to a set of unconstrained square matrices A_1, \dots, A_p . Our reparameterization differs, however, by facilitating the preservation of certain symmetries in the first mapping and the interpretability of

the partial autocorrelation matrices in the second mapping. This allows construction of prior distributions for (A_1, \dots, A_p) which encourage shrinkage toward meaningful parametric structures. In particular, we describe a prior that is exchangeable with respect to the order of the elements in the observation vector. This is likely to be a useful representation of prior beliefs in a variety of applications where the modeler does not have information, a priori, to distinguish between the m time series. We also present an inferential scheme that allows computational inference to be carried out using Euclidean Hamiltonian Monte Carlo, implemented by the probabilistic programming software Stan (Carpenter et al. 2017). This facilitates straightforward extension and adaptation by users in the wide-variety of fields in which vector autoregressions are used. The main contributions of the article are therefore, 2-fold: first, a reparameterization and prior for stationary VARs that allows the incorporation of meaningful initial beliefs; and second, routine implementation of computational inference using standard probabilistic programming software.

2. Reparameterization Over the Stationary Region

2.1. Reparameterization via Partial Autocorrelation Matrices

Ansley and Kohn (1986) establish a one-to-one correspondence between the parameters of a stationary $\text{VAR}_m(p)$ process $(\Sigma, \Phi) \in \mathcal{S}_m^+ \times \mathcal{C}_{p,m}$ and the parameter set $\{\Sigma, (P_1, \dots, P_p)\} \in \mathcal{S}_m^+ \times \mathcal{V}_m^p$ in which P_s denotes the s th partial autocorrelation matrix. In essence, the $(s+1)$ th partial autocorrelation matrix P_{s+1} is a conditional cross-correlation matrix between \mathbf{y}_{t+1} and \mathbf{y}_{t-s} given $\mathbf{y}_t, \dots, \mathbf{y}_{t-s+1}$. More precisely, the matrices P_1, \dots, P_p are defined as follows. For each $s = 1, \dots, p$ let

$$\mathbf{y}_{t+1} = \sum_{i=1}^s \phi_{si} \mathbf{y}_{t-i+1} + \boldsymbol{\epsilon}_{s,t+1}, \quad \mathbf{y}_{t-s} = \sum_{i=1}^s \phi_{si}^* \mathbf{y}_{t-s+i} + \boldsymbol{\epsilon}_{s,t-s}^* \quad (1)$$

in which the $m \times m$ matrices ϕ_{si} and ϕ_{si}^* are the coefficients of the i th terms \mathbf{y}_{t-i+1} and \mathbf{y}_{t-s+i} , respectively, in the conditional expectations $E(\mathbf{y}_{t+1} \mid \mathbf{y}_t, \dots, \mathbf{y}_{t-s+1})$ and $E(\mathbf{y}_{t-s} \mid \mathbf{y}_{t-s+1}, \dots, \mathbf{y}_t)$. It follows that $\phi_{pi} = \phi_i$ ($i = 1, \dots, p$) and $\Sigma_p = \Sigma$. Equivalently, because the multivariate normal distribution is defined by its first two moments, the ϕ_{si} and ϕ_{si}^* are the values of the coefficients, say α_i and α_i^* , in the autoregression of \mathbf{y}_{t+1} on its s predecessors or successors, respectively, that minimize the mean squared error $E\{(\mathbf{y}_{t+1} - \sum_{i=1}^s \alpha_i \mathbf{y}_{t-i+1})^T (\mathbf{y}_{t+1} - \sum_{i=1}^s \alpha_i \mathbf{y}_{t-i+1})\}$ or $E\{(\mathbf{y}_{t-s} - \sum_{i=1}^s \alpha_i^* \mathbf{y}_{t-s+i})^T (\mathbf{y}_{t-s} - \sum_{i=1}^s \alpha_i^* \mathbf{y}_{t-s+i})\}$; see, for example, chap. 3 of Christensen (1991). We define the corresponding conditional variances as $\Sigma_s = \text{var}(\boldsymbol{\epsilon}_{s,t+1}) = \text{var}(\mathbf{y}_{t+1} \mid \mathbf{y}_t, \dots, \mathbf{y}_{t-s+1})$ and $\Sigma_s^* = \text{var}(\boldsymbol{\epsilon}_{s,t-s}^*) = \text{var}(\mathbf{y}_{t-s} \mid \mathbf{y}_{t-s+1}, \dots, \mathbf{y}_t)$ ($s = 1, \dots, p$) and let $\Sigma_0 = \Sigma_0^* = \Gamma_0$ where $\Gamma_i = \text{cov}(\mathbf{y}_t, \mathbf{y}_{t+i}) = E(\mathbf{y}_t \mathbf{y}_{t+i}^T)$ is the i th autocovariance of \mathbf{y}_t . Now, express the conditional variance matrices through a matrix-square-root decomposition, $\Sigma_s = S_s S_s^T$ and $\Sigma_s^* = S_s^* S_s^{*T}$. Although any unique matrix-square-root could be used, we take the symmetric matrix-square-root factorization in this paper, and so $S_s = \Sigma_s^{1/2}$ and $S_s^* = \Sigma_s^{*1/2}$ are symmetric and positive definite. A different reparameterization, defined by Ansley and Kohn (1986), is based on the Cholesky factorization,

in which case S_s and S_s^* are lower triangular. However, as we discuss in Section 3.2, we prefer the symmetric factorization as it facilitates construction of a prior that is closed under orthogonal transformation of the observation vectors. Finally, let $\mathbf{z}_{0,t+1} = S_0^{-1} \mathbf{y}_{t+1}$ and $\mathbf{z}_{0,t}^* = S_0^{*-1} \mathbf{y}_t$ be standardized versions of the forward and reverse time series and, for each $s = 1, \dots, p-1$, let $\mathbf{z}_{s,t+1} = S_s^{-1} \boldsymbol{\epsilon}_{s,t+1}$ and $\mathbf{z}_{s,t-s}^* = S_s^{*-1} \boldsymbol{\epsilon}_{s,t-s}^*$ be standardized versions of the forward and reverse error series. We can now define the partial autocorrelation matrix P_{s+1} , ($s = 0, \dots, p-1$), as

$$\begin{aligned} P_{s+1} &= \text{cov}(\mathbf{z}_{s,t+1}, \mathbf{z}_{s,t-s}^*) \\ &= S_s^{-1} \text{cov}(\mathbf{y}_{t+1}, \mathbf{y}_{t-s} \mid \mathbf{y}_t, \dots, \mathbf{y}_{t-s+1}) (S_s^*)^{-1T} \\ &= S_s^{-1} \phi_{s+1,s+1} S_s^* \end{aligned}$$

in which the final equality is demonstrated in the supplementary materials. This simplifies to the well-known result $\rho_{s+1} = \phi_{s+1,s+1}$ in the univariate case.

The (recursive) mapping from $(\Sigma, \Phi) \in \mathcal{S}_m^+ \times \mathcal{C}_{p,m}$ to $\{\Sigma, (P_1, \dots, P_p)\} \in \mathcal{S}_m^+ \times \mathcal{V}_m^p$ and its inverse are described in the Appendix with proofs in the supplementary materials. Use of symmetric matrix-square-roots, rather than Cholesky factors, complicates the reverse map and precludes use of Lemma 2.3 from Ansley and Kohn (1986) to perform the calculation. This is remedied through a novel recursion which allows computation of the stationary variance matrix Γ_0 from the new set of parameters.

The constraints on the p -fold Cartesian product space \mathcal{V}_m^p are substantially simpler than those on the stationary region $\mathcal{C}_{p,m}$. Indeed, for each $P_s \in \mathcal{V}_m$, the constraint can be expressed as an inequality for the spectral norm, $\|P_s\|_2 < 1$, and so it may be possible to apply the spherical augmentation technique (Lan and Shahbaba 2016), designed for handling norm constraints, to sample from probability distributions over \mathcal{V}_m . However, though the partial autocorrelation matrices are a very natural set of quantities about which to elicit prior beliefs, there are no standard distributions defined over \mathcal{V}_m . Further, any distribution on $M_{m \times m}(\mathbb{R})$ that was simply truncated to \mathcal{V}_m would have an intractable normalizing constant, making its properties difficult to assess. This makes direct specification of a meaningful prior over \mathcal{V}_m^p untenable. We therefore, choose to apply a second reparameterization which maps the partial autocorrelations to unconstrained Euclidean space.

2.2. Reparameterization via Unconstrained Square Matrices

In the second reparameterization, each $P \in \mathcal{V}_m$ is mapped to an unconstrained square matrix $A \in M_{m \times m}(\mathbb{R})$. Generalizing the one-to-one mapping defined by Ansley and Kohn (1986) so that it involves arbitrary matrix-square-roots, the forward mapping is defined as follows. Let

$$B^{-1} B^{-1T} = I_m - P P^T \quad (2)$$

be a matrix-square-root factorization of $I_m - P P^T$. Then write $A = B P$. Similarly, for the inverse mapping, let

$$B B^T = I_m + A A^T \quad (3)$$

be a matrix-square-root factorization of $I_m + A A^T$, then write $P = B^{-1} A$. Although Ansley and Kohn define the mapping

in terms of the Cholesky factorization, we, instead, propose use of symmetric matrix-square-roots because it gives the new parameters a more natural interpretation. Specifically, denote the singular value decomposition of P by $P = URV^T$ in which the diagonal matrix $R = \text{diag}(r_1, \dots, r_m)$ contains the m singular values. These satisfy $1 > r_1 \geq r_2 \geq \dots \geq r_m \geq 0$. It is straightforward to show that if symmetric square-roots are used, the corresponding singular value decomposition of A is $A = U\tilde{R}V^T$ in which $\tilde{R} = (I_m - R^2)^{-1/2}R$ is a diagonal matrix whose i th diagonal element is $\tilde{r}_i = r_i/(1 - r_i^2)^{1/2} \geq 0$. It follows that P and A share the same singular vectors and that the singular values of A are a strictly increasing function of the singular values of P . Clearly, the same functional relationship connects their spectral norms, $r_1 = \|P\|_2$ and $\tilde{r}_1 = \|A\|_2$ and so the relative sizes of $\|A_s\|_2$ across lags s can be interpreted as indicating the relative magnitudes of the partial autocorrelations at each lag. If the symmetric matrix-square-root factorization is used in (2) and (3) we can therefore think of this second reparameterization as an orientation-preserving transformation from P to A which simply maps the singular values from $[0, 1)$ to the positive real line. As a direct consequence, the transformation preserves various meaningful parametric forms whose singular vectors only depend on the elements of the matrix through the requisite ordering of the singular values. Specific examples are detailed below.

From the singular value decomposition of a diagonal matrix it is clear that $P \in \mathcal{V}_m$ is diagonal, with j th diagonal element $p_{jj} = p_j \in [0, 1)$, if and only if A is diagonal, with j th diagonal element $a_{jj} = p_j/(1 - p_j^2)^{1/2} \in \mathbb{R}$ ($j = 1, \dots, m$). A corollary is that P is a scaled identity matrix if and only if the same is true of A . As a special case, when this scaling constant is equal to zero, $P = 0_m$ if and only if $A = 0_m$, where 0_m denotes an $m \times m$ matrix of zeros. This is a particularly useful theoretical result because it follows directly from the definition of the partial autocorrelation matrix that for $k < p$, $P_k \neq 0_m$ and $P_{k+s} = 0_m$ for $s = 1, \dots, p - k$ if and only if $\phi_k \neq 0_m$ and $\phi_{k+s} = 0_m$ for $s = 1, \dots, p - k$. The order of the VAR model is therefore, $k < p$ if and only if $A_k \neq 0_m$ and $A_{k+s} = 0_m$ for $s = 1, \dots, p - k$. We return to this point in Section 6.

Now, consider a two-parameter exchangeable matrix defined by $(b - c)I_m + cJ_m$ where $J_m = \mathbf{1}_m\mathbf{1}_m^T$ and $\mathbf{1}_m$ is an m -vector of 1s. This is the most general form for a $m \times m$ square matrix which is invariant under a common permutation of the rows and columns. It is straightforward to show that a matrix of this form has a singular value decomposition whose singular vectors depend only on m and the ordering of the singular values, $|b - c|$ and $|b + (m - 1)c|$, which have multiplicity $m - 1$ and 1, respectively. It follows that $P \in \mathcal{V}_m$ is a two-parameter exchangeable matrix if and only if the same is true of A . The necessary and sufficient condition for $P = (p_1 - p_2)I_m + p_2J_m$ to lie in \mathcal{V}_m can be expressed as $|p'_1| < \sqrt{2}/2$ and $|p'_2| < \sqrt{2}/m$ where

$$p'_1 = (p_1 - p_2)\sqrt{2}/2, \quad p'_2 = \{p_1 + (m - 1)p_2\} \sqrt{2}/m. \quad (4)$$

It is then straightforward to show that the corresponding unconstrained square matrix $A = (a_1 - a_2)I_m + a_2J_m$, with $a_1, a_2 \in \mathbb{R}$, is such that

$$a_i = a_i^* \left[m \left\{ (2 - m^2 p_2'^2) (1 - 2p_1'^2) \right\}^{1/2} \right]^{-1}, \quad (5)$$

where

$$a_i^* = \left\{ \sqrt{2mp_1'} (2 - m^2 p_2'^2)^{1/2} \right\} (2 - i) - \sqrt{2p_1'} (2 - m^2 p_2'^2)^{1/2} + mp_2' (1 - 2p_1'^2)^{1/2}.$$

3. Prior Distributions over the Unconstrained Space

3.1. General Form

Let $\text{vec}(\cdot)$ denote the vectorization operator. Conditional on a set of unknown hyperparameters, we construct a prior distribution with joint density

$$\pi(\Sigma, A_1, \dots, A_p) = \pi(\Sigma) \prod_{s=1}^p \pi\{\text{vec}(A_s^T)\} \quad (6)$$

in which Σ is assigned a distribution over \mathcal{S}_m^+ and $\text{vec}(A_s^T)$ ($s = 1, \dots, p$), is assigned a multivariate normal distribution. Predominantly, our focus in this article is specification of a prior for the latter.

3.2. Exchangeable Prior Distribution

Consider any $m \times m$ orthogonal matrix H . Assuming that symmetric matrix-square-roots are used in both parts of the reparameterization of a stationary VAR $_m(p)$ model for \mathbf{y}_t ($t = 1, 2, \dots$), we show in the supplementary materials that the parameters of the stationary VAR $_m(p)$ model for $\tilde{\mathbf{y}}_t = H\mathbf{y}_t$ are $\tilde{\Sigma} = H\Sigma H^T$ and $\tilde{A}_s = HA_s H^T$ ($s = 1, \dots, p$). It follows that if Σ and A_s ($s = 1, \dots, p$) are assigned a prior from the same distributional family as that of $H\Sigma H^T$ and $HA_s H^T$ ($s = 1, \dots, p$), then the prior induced for (Σ, Φ) over $\mathcal{S}_m^+ \times \mathcal{C}_{p,m}$ will be closed under orthogonal transformation of the observation vectors. Priors for Σ over \mathcal{S}_m^+ possessing this closure property include the inverse Wishart distribution and the multivariate normal distribution for the matrix logarithm, $\log \Sigma$ (Leonard and Hsu 1992). For $\text{vec}(A_s^T) \in \mathbb{R}^{m^2}$, a multivariate normal prior meets this requirement. It is important to note that if Cholesky factors were used in the first and second part of the reparameterization, such a prior would not be available because the partial autocorrelation matrices \tilde{P}_s and associated unconstrained \tilde{A}_s would not be orthogonal similarity transformations of P_s and A_s .

In the analysis of multivariate stochastic processes, we often do not have information, a priori, to distinguish between the m components of \mathbf{y}_t . In this case it is reasonable to assign (Σ, Φ) a prior which is exchangeable with respect to the ordering of the elements in the observation vector. Because of its closure under orthogonal transformation, we can obtain an exchangeable prior by restricting our attention to distributions for Σ and A_s that would be exactly the same as the distributions for $H\Sigma H^T$ and $HA_s H^T$ for any permutation matrix H . That is, distributions for Σ and A_s which are invariant under a common permutation of the rows and columns. Given certain choices of their hyperparameters, the (conjugate) inverse Wishart distribution and the multivariate normal distribution for the matrix-logarithm can yield an exchangeable prior for Σ . For instance, we could assign Σ an inverse Wishart prior with a (positive definite) two-parameter exchangeable scale matrix.

Now, suppose we wish to assign an exchangeable prior to $A_s = (a_{s,ij})$ ($s = 1, \dots, p$). Given the potential for the model to contain a very large number of parameters, suppose further that we want to specify a prior that allows borrowing of strength between the diagonal elements and between the off-diagonal elements of each A_s . To this end, we can adopt a prior in which the diagonal and off-diagonal elements are given hierarchical distributions. At the top-level, we choose

$$a_{s,ii} \mid \mu_{s1}, \omega_{s1} \sim \mathcal{N}(\mu_{s1}, \omega_{s1}^{-1}), \quad (i = 1, \dots, m), \quad (7)$$

$$a_{s,ij} \mid \mu_{s2}, \omega_{s2} \sim \mathcal{N}(\mu_{s2}, \omega_{s2}^{-1}), \quad (i, j = 1, \dots, m \text{ with } i \neq j). \quad (8)$$

The mean and precision at the bottom level of the hierarchy can then be assigned priors on \mathbb{R} and \mathbb{R}^+ , such as

$$\mu_{si} \sim \mathcal{N}(e_{si}, f_{si}^2), \quad \omega_{si} \sim \text{Gam}(g_{si}, h_{si}), \quad (i = 1, 2). \quad (9)$$

Marginally, $E(a_{s,ii}) = e_{s1}$, $\text{var}(a_{s,ii}) = f_{s1}^2 + h_{s1}/(g_{s1} - 1)$ (for $g_{s1} > 1$) and $\text{cor}(a_{s,ii}, a_{s,ij}) = f_{s1}^2(g_{s1} - 1)/\{f_{s1}^2(g_{s1} - 1) + h_{s1}\}$ with similar expressions for the moments of the off-diagonal elements. Therefore, given specifications for the common diagonal elements and off-diagonal elements in P_s , one can calculate corresponding values for the common diagonal and off-diagonal elements in A_s through (4)–(5). These can be taken as values for $e_{s1} = E(a_{s,ii})$ and $e_{s2} = E(a_{s,ij})$. Uncertainty in these central values, and the proportion of this which is shared among all the diagonal or all the off-diagonal elements, can be reflected through choices of the other hyperparameters. Clearly, specifications which make the marginal variances small and the marginal correlations large will shrink the posterior so that it is more concentrated over the space of two-parameter exchangeable structures.

3.3. Prior Distribution Centered on a Diagonal Matrix

Let $\mathbf{y}_{ij} = (\mathbf{y}_i^T, \dots, \mathbf{y}_j^T)^T$. For each $j = 1, \dots, m$, suppose it is believed a priori that once $\mathbf{y}_{t:(t-s+1)}$ is known, $\mathbf{y}_{t-s,j}$ is the only element in \mathbf{y}_{t-s} which provides further information about $\mathbf{y}_{t+1,j}$. This is tantamount to a conjecture that the partial autocorrelation matrix P_s , and hence A_s , is diagonal; see Section 2.2. To represent this belief we choose a prior in which $(a_{s,11}, \dots, a_{s,mm})^T \sim N_m(\mathbf{e}_s, F_s)$ for the diagonal elements. Alternatively, if there was nothing in our prior beliefs to distinguish among the diagonal elements, and we wanted to allow borrowing of strength between them, we might adopt the hierarchical prior in (7) and (9) when $i = 1$. For the off-diagonal elements $a_{s,ij}$ ($i \neq j$), we can center our prior around zero, whilst allowing the data to influence the degree of shrinkage toward zero, by adopting a special case of the hierarchical prior defined by (8) and (9) in which the distribution for μ_{s2} is a point mass at zero. Alternatively, as we discuss further in Section 3.4, the off-diagonal elements could be assigned a sparsity-inducing prior.

3.4. Prior Distribution Encouraging Sparse A_s Matrices

A $\text{VAR}_m(p)$ model is highly parameterized with $O(m^2)$ parameters. Indeed, particularly when m is large, it is entirely plausible for there to be fewer observations in the data than there are parameters in the model. In a Bayesian analysis, this can lead

to a diffuse posterior distribution for (Σ, Φ) , making predictive distributions imprecise and complicating model interpretation. When stationarity is not enforced, this issue is often addressed by inducing sparsity amongst the elements of the (unconstrained) autoregressive coefficient matrices ϕ_1, \dots, ϕ_p , either through graphical modeling or zero-mean shrinkage priors (George, Sun, and Ni 2008; Billio, Casarin, and Rossini 2019). The nonzero structure can then be associated with a directed graph representing a network of interactions because a zero in position (i, j) of ϕ_s implies conditional independence between $y_{t,i}$ and $y_{t-s,j}$ given $(y_{t-1}, \dots, y_{t-s,-j}, \dots, y_{t-p})$ where $y_{t-s,-j} = (y_{t-s,1}, \dots, y_{t-s,j-1}, y_{t-s,j+1}, \dots, y_{t-s,m})^T$. In principle, sparsity-inducing priors could also be chosen for the elements of A_1, \dots, A_p in our parameterization, where stationarity is enforced. Among the partial autocorrelation matrices, a zero in position (i, j) of P_s implies conditional independence between $y_{t,i}$ and $y_{t-s,j}$ given $\mathbf{y}_{(t-1):(t-s+1)}$. Transforming from P_s to A_s , an individual zero in position (i, j) of P_s does not give an individual zero in position (i, j) of A_s or vice versa and so sparsity in A_s does not have a clear structural interpretation as it would for P_s or ϕ_s . However, the overall size of A_s is strongly linked to the overall size of P_s through the strictly increasing relationship that connects their spectral norms; see Section 2.2. Therefore, although the justification for a sparsity-inducing prior for the A_s is weaker from an explanatory perspective, there is still an argument for their use as a means of regularizing the variance of predictive distributions.

A sparsity-inducing, zero-mean scale-mixture of normals prior would take the form $a_{s,ii} \mid \psi_{s,ii} \sim \mathcal{N}(0, \psi_{s,ii})$ with $\psi_{s,ii} \sim \mathcal{F}_1$, independently for $i = 1, \dots, m$, and $a_{s,ij} \mid \psi_{s,ij} \sim \mathcal{N}(0, \psi_{s,ij})$ with $\psi_{s,ij} \sim \mathcal{F}_2$, independently for $i, j = 1, \dots, m$ with $i \neq j$. Here \mathcal{F}_1 and \mathcal{F}_2 are mixing distributions which can either be discrete, as in spike-and-slab priors (Mitchell and Beauchamp 1988; George and McCulloch 1993), or continuous, as in the horseshoe (Carvalho, Polson, and Scott 2010; Piironen and Vehtari 2017).

3.5. Vague Prior via the Parameterization of Roy, McElroy, and Linton (2019)

Denote by $\mathcal{O}(m)$ the space of $m \times m$ orthogonal matrices. Roy, McElroy, and Linton (2019) establish a bijective mapping between the parameters of a stationary $\text{VAR}_m(p)$ process $(\Sigma, \Phi) \in \mathcal{S}_m^+ \times \mathcal{C}_{p,m}$ and the parameter set $\{\Sigma, (V_1, \dots, V_p), (Q_1, \dots, Q_p)\} \in \mathcal{S}_m^+ \times \mathcal{S}_m^{+p} \times \mathcal{O}(m)^p$ for any fixed choice of a pseudo error variance matrix $M \in \mathcal{S}_m^+$. It relies on a characterization in terms of positive definite block Toeplitz matrices, like that describing the variance of the joint stationary distribution of p consecutive time points. For the special case when $M = \Sigma$, we show in the supplementary materials that V_i represents the difference in conditional variances, $V_i = \text{var}(\mathbf{y}_t \mid \mathbf{y}_{t-1}, \dots, \mathbf{y}_{t-i+1}) - \text{var}(\mathbf{y}_t \mid \mathbf{y}_{t-1}, \dots, \mathbf{y}_{t-i})$ ($i = 1, \dots, p$), and the orthogonal matrix Q_i arises from the polar decomposition of an affine transformation of the partial autocorrelation matrix P_i . Unfortunately, this makes the orthogonal matrices difficult to interpret which impedes specification of a meaningful prior. Moreover, computational inference is challenging because the constraints that define $\mathcal{O}(m)$ complicate sampling by MCMC.

For example, attempts to provide a general reparameterization of $Q \in \mathcal{O}(m)$ in terms of unconstrained parameters, such as the Givens representation (Pourzanjani et al. 2021) or Cayley transform (Jauch, Hoff, and Dunson 2020), are typically frustrated by the pathological effects of mapping between two topologically distinct spaces. Indeed, as explained in the supplementary materials, the modified Cayley transform suggested by Roy, McElroy, and Linton is not bijective, which makes the posterior of their real-valued parameterization multimodal and can cause inefficient MCMC simulation. These geometric problems can be avoided by using Geodesic Monte Carlo (Byrne and Girolami 2013) which is able to sample efficiently from $\mathcal{O}(m)$ by tailoring the Hamiltonian Monte Carlo method to embedded manifolds. However, there is currently no methodology for automatic tuning of its parameters, nor any modular software for implementation. This makes it difficult to put into practice through bespoke MCMC programs and impossible to implement using probabilistic programming software. Fortunately, in the supplementary materials we present a simple reparameterization of the parameter set $\{(V_1, \dots, V_p), (Q_1, \dots, Q_p)\}$ in terms of p unconstrained square matrices which circumvents the sampling issue. We show that this is equivalent to assigning independent uniform distributions over $\mathcal{O}(m)$ to the Q_s and independent Wishart distributions to the V_s , with m degrees of freedom and identity scale, and so it can be regarded as a vague, stationary prior distribution. This might be attractive to some modelers as a default choice of prior.

3.6. Choice of Prior Variance

Using the simple example of a VAR₂(1) model, we show in the supplementary materials that the prior for the partial autocorrelation matrix P_1 can become multimodal when the prior variance for the elements of the unconstrained square matrix A_1 becomes too large. For most problems, a multimodal prior for a partial autocorrelation matrix P is unlikely to be representative of prior beliefs. To avoid this, care is clearly needed in the choice of prior variance for the elements of the unconstrained square matrices.

As discussed in Section 2.2, a partial autocorrelation matrix P is constructed from the corresponding unconstrained matrix $A = (a_{ij})$ through a simple mapping of its singular values from the positive real line to the unit interval. It is reasonable, therefore, to conjecture that the multimodality that can occur in the prior for the partial autocorrelations, but not in the multivariate normal prior for the unconstrained matrices, arises through this mapping of the singular values. For any m , under the simple prior $a_{ij} \sim N(0, s^2)$ ($i, j = 1, \dots, m$), we show in the supplementary materials that the singular values and right and left singular vectors of P are independent a priori. Moreover, we show that the singular vectors are distributed as independent, normalized Haar measures, and derive an analytic expression for the joint prior density of the singular values. Arguing that multimodality arises when the singular values have a local maximum in the interior of their parameter space, we find the smallest prior standard deviation s of the a_{ij} at which this occurs for various values of m . The conclusion is that a prior standard deviation of $s = 1$ should prevent multimodality for $m \geq 5$. Guidance on an upper limit for s when $m = 1, \dots, 4$ can be found in the supplementary materials.

4. Posterior Inference through MCMC

Consider observations $y_{1:n}$ modeled as realizations from a stationary VAR _{m} (p) process. The likelihood function can be expressed as

$$p(y_{1:n} | \Sigma, \Phi) = p(y_{1:p} | \Sigma, \Phi) \prod_{t=p+1}^n p(y_t | y_{(t-p):(t-1)}, \Sigma, \Phi)$$

in which $Y_t | y_{(t-p):(t-1)}, \Sigma, \Phi \sim N_m(\sum_{i=1}^p \phi_i y_{t-i}, \Sigma)$ and the initial distribution is $(Y_1^T, \dots, Y_p^T)^T | \Sigma, \Phi \sim N_{mp}(\mathbf{0}, G)$ where G is a positive definite block Toeplitz matrix with Γ_{j-i} as the block in rows $\{m(i-1)+1\}$ to mi and columns $\{m(j-1)+1\}$ to mj ($i, j = 1, \dots, p$) and $\Gamma_{-k} = \Gamma_k^T$ ($k = 1, \dots, p-1$). For the purposes of evaluating the likelihood, the stationary variance Γ_0 and covariances $\Gamma_1, \dots, \Gamma_{p-1}$ are available as a by-product of the reverse mapping detailed in the Appendix.

Treating the likelihood as a function of the new parameters and combining it with the prior (6) through Bayes' theorem yields the posterior distribution as

$$\begin{aligned} \pi(\Sigma, A_1, \dots, A_p, \vartheta | y_{1:n}) \\ \propto \pi(\Sigma)\pi(\vartheta) \prod_{i=1}^p \pi\{\text{vec}(A_i^T) | \vartheta\} p(y_{1:n} | \Sigma, A_1, \dots, A_p), \end{aligned}$$

in which ϑ denotes any unknown hyperparameters in the prior for the A_s . As each of the autoregressive coefficients ϕ_s is a complicated function of the complete set of new parameters, $\{\Sigma, (A_1, \dots, A_p)\}$, this posterior distribution neither has a standard form, nor admits any simple factorization that would arise from conditional independence amongst $\{\Sigma, (A_1, \dots, A_p)\}$. As a consequence, it is ill-suited to MCMC methods that are based on Gibbs sampling. Indeed, our experience suggests that Metropolis-within-Gibbs samplers, which iterate through random-walk updates of $(\Sigma | A_1, \dots, A_p)$, $(A_1 | A_2, \dots, A_p, \Sigma)$, \dots , $(A_p | A_1, \dots, A_{p-1}, \Sigma)$, perform very poorly as soon as $p > 1$. It is therefore, beneficial to use a sampler such as Hamiltonian Monte Carlo (HMC) (Girolami and Calderhead 2011; Neal 2011) which uses information on the slope of the logarithm of the posterior density to generate global proposals that update all parameters simultaneously. We use rstan (Stan Development Team 2020), the R interface to the Stan software, to implement the HMC algorithm. Stan requires users to write a program in the probabilistic Stan modeling language, the role of which is to provide instructions for computing the logarithm of the kernel of the posterior density function. The Stan software then automatically sets up a Markov chain simulation to sample from the resulting posterior. This includes calculation of the gradient of the logarithm of the posterior density, random initialization of the chains, and the tuning of the sampler.

5. Application

In this section, we illustrate use of our exchangeable prior and inferential methods by applying them to a quarterly time series of U.S. macroeconomic data. The complete dataset comprises measurements on 168 variables, running from quarter 1 of 1959

to quarter 4 of 2007. The variables are transformed to stationarity by differencing, sometimes after applying a log transformation, and then standardized, so it is reasonable to model the data as arising from a zero mean stationary process. A full description of the data and transformations can be found in Koop (2013). Following analyses of the same data in Koop (2013) and Koop and Korobilis (2009), interest lies primarily in forecasting the first three variables: real GDP, the consumer price index, and an interest rate (Federal funds). Like these original analyses, we consider a small VAR₃(4) model and two larger VAR₁₀(4) and VAR₂₀(4) models, where the original three variables are supplemented by an additional 7, then a further 10, which are thought to have forecasting value. A list of these variables can be found in the supplementary materials. In order to assess the forecasting properties of the models, we fitted the models using data $y_{1:n}$ where $n = 156$ and held back the last 40 observations $y_{(n+1):(n+40)}$ in all analyses. This allowed us to base our measures of forecasting performance on the posterior predictive distribution of the held-back data.

When it is reasonable to assume that a process is stationary, one of the problems of fitting an unconstrained VAR model is that some posterior mass often lies outside the stationary region. Typically, this is due to a combination of model misspecification and epistemic uncertainty in the parameter values that cannot be resolved by the data that have been observed. Therefore, to demonstrate the practical benefits of using our stationary prior, we additionally consider two commonly used priors that do not constrain inference to the stationary region. Specifically, we compare (i) the exchangeable, stationary prior from Section 3.2 to: (ii) a Minnesota prior; (iii) a semi-conjugate prior, which takes the form $\pi(\Phi, \Sigma) = \pi(\Phi)\pi(\Sigma)$, where $\text{vec}(\phi_k) \sim N(\mathbf{u}_{k3}, W_{k3})$ independently for $k = 1, \dots, 4$ and $\Sigma \sim IW(m + 4, \mathbf{I}_m)$. Clearly, for the latter two analyses, the time series cannot be initialized at the stationary distribution, and so we simply condition on the first $p = 4$ observations in the time series. For the exchangeable, stationary prior, we choose correlations of 0.7 between diagonal and between off-diagonal elements to facilitate borrowing of strength between elements. Under the Minnesota prior, the autoregressive coefficients are assigned a multivariate normal distribution, $\text{vec}(\phi_k) \sim N(\mathbf{u}_{k2}, W_{k2})$ independently for $k = 1, \dots, 4$, and the error variance matrix Σ is replaced by an estimate $\hat{\Sigma} = \text{diag}(s_1^2, \dots, s_m^2)$, where s_j^2 is the ordinary least-squares estimate of the error variance in the (univariate) autoregression for variable j . The prior mean \mathbf{u}_{k2} is chosen so that $E(\phi_{k,ij}) = 0$ if $i \neq j$ and the prior variance W_{k2} is taken to be diagonal; generally different variances are chosen for the $\phi_{k,ii}$ and the $\phi_{k,ij}$ for $i \neq j$ and the variances decrease as the lag k increases. The idea is to encourage shrinkage toward a simple set of low order univariate AR models with the objective of reducing the epistemic component of the forecast variance. The complete prior specifications are provided in the supplementary materials.

Even when inference is constrained to the stationary region, without appropriate borrowing of strength between parameters, posterior predictive distributions can still be overly diffuse, especially when m is large. We therefore demonstrate the useful regularization effect that our exchangeable prior can provide by considering two further analyses where stationarity is guaran-

teed, but where there is no shrinkage toward a sensible parametric structure. In particular, we consider: (iv) the modified reparameterization of Roy, McElroy, and Linton (2019), and its associated vague prior, described in Section 3.5; (v) a frequentist, maximum likelihood analysis using the partial autocorrelation reparameterization of Ansley and Kohn (1986) that is based on Cholesky factors rather than symmetric matrix-square-roots.

To fit the model under the exchangeable, stationary prior, the semi-conjugate prior and the vague, stationary prior we used HMC implemented by Stan. For all three datasets and all three priors, we used the `rstan` interface to the Stan software to run four chains, initialized at different starting points, for 2000 iterations, half of which were discarded as burn-in. The usual graphical and numerical diagnostics gave no evidence of any lack of convergence and, after pooling the chains, the effective sample size was at least 1262 for every parameter. The Minnesota prior is conjugate and so the posterior distribution can be computed analytically; see, for example, chap. 2 of Koop and Korobilis (2009) for its closed form. We generated 4000 independent draws from this distribution to allow calculation of the sample-based statistics described below. Model-fitting by maximum likelihood was also carried out using Stan which implements numerical maximization of the log-likelihood function by the quasi-Newton algorithm L-BFGS. We note that when $m = 20$ the algorithm repeatedly failed to converge when initialized randomly and convergence was only achieved after initializing at the posterior mean deduced from the analysis under the exchangeable, stationary prior. The four Stan programs are given in the supplementary materials.

To assess the forecasting performance of the various model-prior combinations we consider a variety of forecast horizons: $h = 1, 2, 4, 8$, ranging from short-term, one-quarter-ahead forecasting ($h = 1$) to longer-term, two-year-ahead forecasting ($h = 8$). For the comparison of h -step-ahead predictions, we use a number of proper scoring rules (Gneiting and Raftery 2007), and the posterior for the empirical mean squared forecast error (MSFE) for each variable of interest. The MSFE is designed to measure the accuracy and precision of point forecasts, being based on the mean squared deviation between y_t and its expectation given $y_{1:(t-h)}$ across the hold-out period, $t = n + h, \dots, n + 40$. Proper scoring rules compare the whole forecast distribution with the observation that arises; by assigning a numerical score, this allows competing forecast distributions to be ranked. At every $t = n + h, \dots, n + 40$ these scores are based on the h -step-ahead posterior predictive distribution at time t , which is then averaged across the hold-out period. We chose two widely used scores to assess forecasting performance for the three variables of interest, individually: the continuous rank probability score (CRPS) and the logarithmic score. We additionally compare joint forecasts of the three variables of interest by computing the energy score (ES), which is a multivariate generalization of the CRPS. In all cases, the scores are negatively oriented so that small values indicate better forecasting performance. Further details on the calculation of the MSFE and the proper scoring rules, along with their adaptation for the maximum likelihood analysis, can be found in the supplementary materials. For each value of m and each prior, the values of the one-step-ahead ($h = 1$) CRPS, ES and posterior mean MSFE are shown in Figure 1, along with approximations

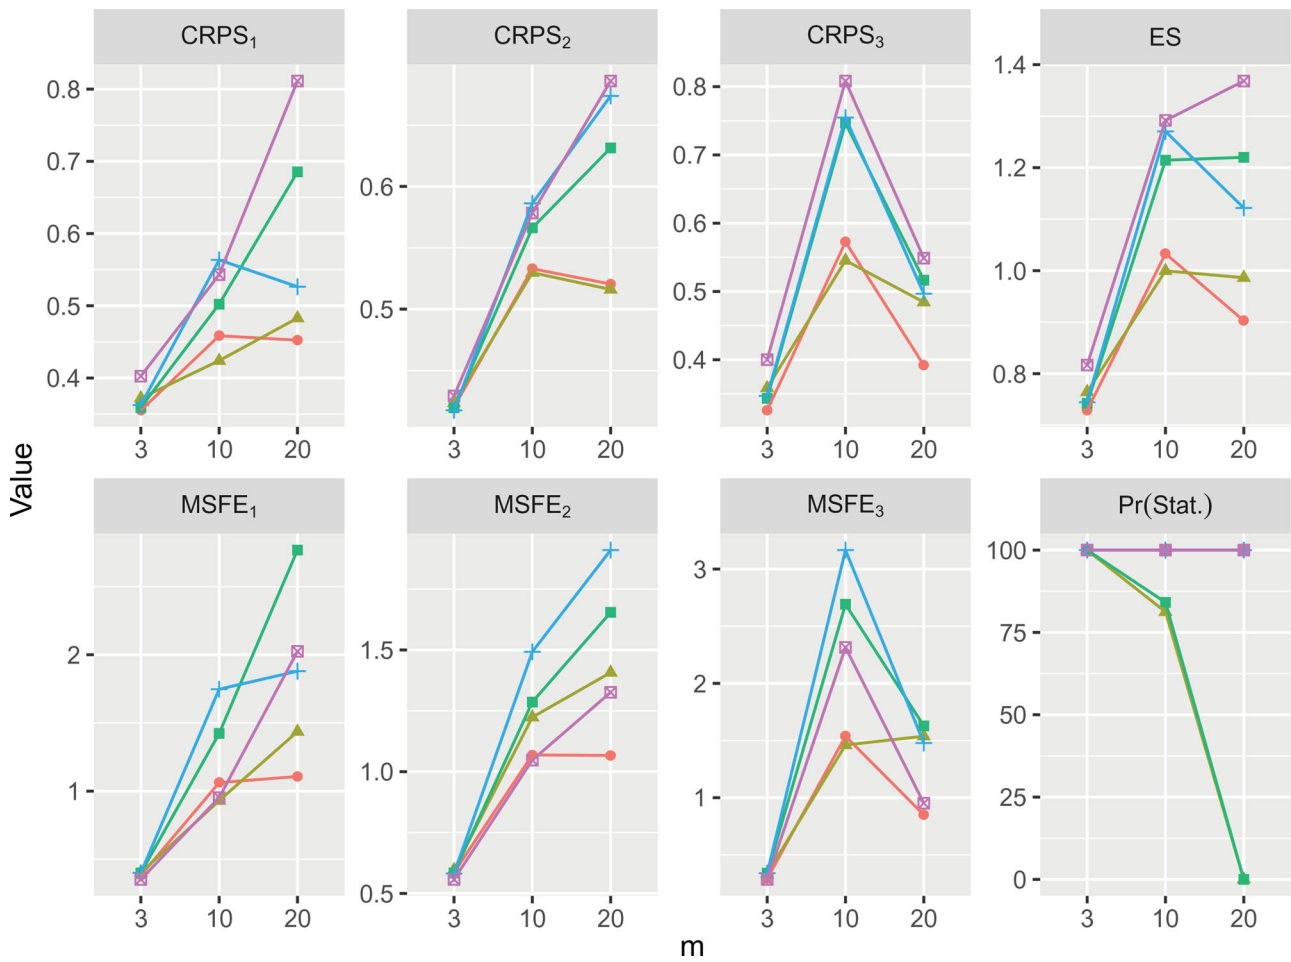


Figure 1. For each value of m and each prior: one-step-ahead CRPS for variable k ($CRPS_k$); one-step-ahead ES; posterior mean one-step-ahead MSFE for variable k ($MSFE_k$); posterior probability that $\Phi \in \mathcal{C}_{4,m}$ ($Pr(Stat.)$). The priors are: exchangeable and stationary (●); Minnesota (▲); semi-conjugate (■); vague and stationary prior (+). Also shown are analogous statistics based on the stationary MLE (Ansley and Kohn 1986) (⊠).

of the posterior probability that $\Phi \in \mathcal{C}_{4,m}$. Analogous plots for the one-step-ahead logarithmic scores, which showed similar patterns to the CRPS, and for all metrics at the other horizons $h = 2, 4, 8$ can be found in the supplementary materials.

For each value of m , and under all metrics, it is noticeable from Figure 1 that the exchangeable, stationary prior performs consistently well, especially when $m = 20$, where it is ranked as best by all criteria. The Minnesota prior also forecasts well in this one-step-ahead setting. The frequentist forecasts appear successful in terms of the MSFE, but their performance is ranked worst by the proper scoring rules. Since these forecasts ignore epistemic uncertainty in the parameter values, it is likely that this is because the forecast distributions are overly concentrated. The semi-conjugate prior and the vague, stationary prior perform poorly according to all metrics, which is likely because of the lack of structure imposed by modeling assumptions, in the former case, or encouraged by the prior, in the latter case, leading to very diffuse predictive distributions. Additional figures in the supplementary materials reveal how the performance of the various predictive distributions changes moving from one-quarter-ahead forecasts to up to two-year-ahead forecasts. Over longer time horizons h , the exchangeable, stationary prior is the only one whose performance is consistently strong according to all metrics, especially when $m = 10$ and $m = 20$ where

differences to other forecasts are most apparent. The performance of the frequentist forecasts continue to appear successful in terms of the MSFE, but less so under the proper scoring rules. In contrast, the performance of the Minnesota prior deteriorates rapidly as h increases, possibly due to its use of a fixed diagonal error variance matrix. As the forecast horizon increases, it is interesting to observe how the behavior of the forecasts under the priors which do and do not enforce stationarity change with m . Under the stationarity-enforcing prior distributions, performance is generally similar, and sometimes better, going from $m = 10$ to $m = 20$. However, under the priors which do not enforce stationarity, there is a marked deterioration in forecasting performance. As m increases from 10 to 20, the approximate posterior probability that Φ lies inside the stationary region decreases from 0.8125 to 0.0000 under the Minnesota prior and from 0.8413 to 0.0000 under the semi-conjugate prior. At least in part, this is likely to be due to epistemic uncertainty spilling outside the stationary region, whose volume becomes vanishingly small when p or, in this case, m become large; see Section 1. When the parameters of a $VAR_m(p)$ model lie outside the stationary region, the forecast variance increases without bound into the future. Therefore, one might reasonably conjecture that it is this greater concentration of posterior mass outside the stationary region that leads to the more

noticeable deterioration in forecasting performance when inference is not constrained to the stationary region. This suggests that using a constrained prior distribution may be even more important in problems where m or p are large. Moreover, the better performance of the exchangeable, stationary prior over the vague, stationary prior illustrates the benefits of adopting a prior which encourages shrinkage toward sensible parametric structures.

6. Discussion

It is straightforward to extend the reparameterization and prior presented here to VARMA models. Consider the model of order (p, q) , or the VARMA $_m(p, q)$ model, $\theta(B)\epsilon_t = \phi(B)y_t$, where $\theta(u) = I_m + \theta_1 u + \dots + \theta_q u^q$, $u \in \mathbb{C}$, is the characteristic moving average polynomial, in which $\theta_i \in M_{m \times m}(\mathbb{R})$ ($i = 1, \dots, q$). As for VAR $_m(p)$ models, the process is stable if and only if all the roots of $\det\{\phi(u)\} = 0$ lie outside the unit circle. It is invertible if and only if all the roots of $\det\{\theta(u)\} = 0$ lie outside the unit circle, and hence the invertible region is $C_{q,m}$. We can therefore, constrain inference to the stationary and invertible regions simultaneously by reparameterizing the model in terms of two sets of matrices with singular values less than one. In other words, we can apply the recursions described in the Appendix twice; once as if we had a pure VAR $_m(p)$ model with coefficients ϕ_1, \dots, ϕ_p and variance Σ to get P_1, \dots, P_p , and again as if we had a pure VAR $_m(q)$ model with coefficients $-\theta_1, \dots, -\theta_q$ and variance Σ to get, say, R_1, \dots, R_q . Unfortunately, interpretation of the new parameter sets is a little less clear; the P_s represent the partial autocorrelation matrices of the autoregressive part of the process, and the R_s represent a multivariate analogue of the inverse partial autocorrelation function (Bhansali 1983) for the moving average part of the process. In each case, the second transformation from Section 2.2 can be used to map the parameters to unconstrained Euclidean space. Details on inference are given in the supplementary materials.

The ideas discussed in this article can also be extended to VAR models of unknown order. This might be useful when there is little prior information to guide the choice of p . At least in theory, allowing for uncertainty in the model order is easy to handle in the Bayesian framework. Suppose we are prepared to consider models where the order does not exceed p_{\max} . We can then make the model order, say k , unknown and assign it a prior over $\{0, \dots, p_{\max}\}$. As discussed in Section 2.2, the vector autoregression is of order $k < p_{\max}$ if and only if $A_k \neq 0_m$ and $A_{k+s} = 0_m$ ($s = 1, \dots, p_{\max} - k$) and so the models of different orders are nested. Various transdimensional MCMC samplers have been developed for problems like these, where the dimension of the parameter space is itself unknown. In the context of univariate autoregressions, this includes reversible jump MCMC (Green 1995) and birth-death MCMC (Stephens 2000), with samplers that both do, and do not, enforce stationarity (e.g., Vermaak et al. 2004; Philippe 2006). The advantages of tackling this problem with our reparameterization of the VAR $_m(p)$ process is that it allows stationarity to be enforced whilst maintaining the nested structure of the models from the original parameterization. We defer further consideration of this interesting challenge to future work.

Supplementary Materials

Supporting information: Proofs of the forward and reverse mappings, additional plots, simulations and text illustrating properties of the prior, full details of the modified parameterization of Roy, McElroy, and Linton (2019) and its vague prior, a complete description of the extension of the reparameterization and prior to VARMA models and further details on the application to macroeconomic data. (.pdf file)

Data, R code and Stan programs: The macroeconomic data and R scripts for preprocessing and running Stan, four Stan programs for fitting the VAR $_m(p)$ models in the application, and a fifth for fitting a VARMA $_m(p, q)$ model with a prior that constrains inference to the stationary and invertible regions. (.zip archive file)

Acknowledgments

I would like to thank the referees and editors for their helpful comments and suggestions. I am also grateful to Gary Koop for providing access to the macroeconomic data and to Michael Betancourt, Malcolm Farrow, Tom Nye, Anindya Roy and Darren Wilkinson for conversations which have improved the manuscript.

Funding

The author gratefully acknowledges the financial support of the UK EPSRC grant EP/N510129/1 through the Alan Turing Institute project “Streaming data modeling for real-time monitoring and forecasting.”

References

- Abegaz, F., and Wit, E. (2013), “Sparse Time Series Chain Graphical Models for Reconstructing Genetic Networks,” *Biostatistics*, 14, 586–599. [74]
- Ansley, C. F., and Kohn, R. (1986), “A Note on Reparameterizing a Vector Autoregressive Moving Average Model to Enforce Stationarity,” *Journal of Statistical Computation and Simulation*, 24, 99–106. [75,76,80,81]
- Ansley, C. F., and Newbold, P. (1979), “Multivariate Partial Autocorrelations,” in *Proceedings of the Business and Economics Section*, pp. 349–353. American Statistical Association. [75]
- Bañbura, M., Giannone, D., and Reichlin, L. (2010), “Large Bayesian Vector Auto Regressions,” *Journal of Applied Econometrics*, 25, 71–92. [74]
- Barndorff-Nielsen, O., and Schou, G. (1973), “On the Parameterization of Autoregressive Models by Partial Autocorrelations,” *Journal of Multivariate Analysis* 3, 408–419. [75]
- Barnett, G., Kohn, R., and Sheather, S. (1996), “Bayesian Estimation of an Autoregressive Model Using Markov Chain Monte Carlo,” *Journal of Econometrics*, 74, 237–254. [75]
- Bhansali, R. J. (1983), “The Inverse Partial Correlation Function of a Time Series and its Applications,” *Journal of Multivariate Analysis*, 13, 310–327. [82]
- Billio, M., Casarin, R., and Rossini, L. (2019), “Bayesian Nonparametric Sparse VAR Models,” *Journal of Econometrics*, 212, 97–115. [78]
- Byrne, S., and Girolami, M. (2013), “Geodesic Monte Carlo on Embedded Manifolds,” *Scandinavian Journal of Statistics*, 40, 825–845. [79]
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M. A., Guo, J., Li, P., and Riddell, A. (2017), “Stan: A Probabilistic Programming Language,” *Journal of Statistical Software*, 76, 1–32. [76]
- Carvalho, C. M., Polson, N. G., and Scott, J. G. (2010), “The Horseshoe Estimator for Sparse Signals,” *Biometrika*, 97, 465–480. [78]
- Chiang, S., Guindani, M., Yeh, H. J., Haneef, Z., Stern, J. M., and Vannucci, M. (2017), “Bayesian Vector Autoregressive Model for Multi-Subject Effective Connectivity Inference Using Multi-Modal Neuroimaging Data,” *Human Brain Mapping*, 38, 1311–1332. [74]
- Chib, S. (1993), “Bayes Estimation of Regressions with Autoregressive Errors: A Gibbs Sampling Approach,” *Journal of Econometrics*, 58, 275–294. [75]

- Christensen, R. (1991), *Linear Models for Multivariate, Time Series, and Spatial Data*, New York: Springer. [76]
- Corander, J., and Villani, M. (2005), “A Bayesian Approach to Modeling Graphical Vector Autoregressions,” *Journal of Time Series Analysis*, 27, 141–156. [74]
- Doan, T., Litterman, R. B., and Sims, C. A. (1984), “Forecasting and Conditional Projection Using Realistic Prior Distributions,” *Econometric Reviews*, 3, 1–100. [74]
- George, E. I., and McCulloch, R. E. (1993), “Variable Selection via Gibbs Sampling,” *Journal of the American Statistical Association*, 88, 881–889. [78]
- George, E. I., Sun, D., and Ni, S. (2008), “Bayesian Stochastic Search for VAR Model Restrictions,” *Journal of Econometrics*, 142, 553–580. [78]
- Ghosh, S., Khare, K., and Michailidis, G. (2019), “High-Dimensional Posterior Consistency in Bayesian Vector Autoregressive Models,” *Journal of the American Statistical Association*, 114, 735–748. [74]
- Girolami, M., and Calderhead, B. (2011), “Riemann Manifold Langevin and Hamiltonian Monte Carlo Methods,” *Journal of the Royal Statistical Society, Series B*, 73, 123–214. [79]
- Gneiting, T., and Raftery, A. E. (2007), “Strictly Proper Scoring Rules, Prediction, and Estimation,” *Journal of the American Statistical Association*, 102, 359–378. [80]
- Green, P. (1995), “Reversible Jump Markov chain Monte Carlo Computation and Bayesian Model Determination,” *Biometrika*, 82, 711–732. [82]
- He, Y., She, Y., and Wu, D. (2013), “Stationary-Sparse Causality Network Learning,” *Journal of Machine Learning Research*, 14, 3073–3104. [75]
- Huerta, G., and Prado, R. (2006), “Structured Priors for Multivariate Time Series,” *Journal of Statistical Planning and Inference*, 136, 3802–3821. [75]
- Huerta, G., and West, M. (1999), “Priors and Component Structures in Autoregressive Time Series Models,” *Journal of the Royal Statistical Society, Series B*, 61, 881–899. [75]
- Jauch, M., Hoff, P. D., and Dunson, D. B. (2020), “Random Orthogonal Matrices and the Cayley Transform,” *Bernoulli*, 26, 1560–1586. [79]
- Karlsson, S. (2013), “Forecasting with Bayesian Vector Autoregression,” in *Handbook of Economic Forecasting* (Vol. 2B), eds. G. Elliott and A. Timmermann, pp. 791–897, Amsterdam: North Holland. [74]
- Koop, G., and Korobilis, D. (2009), “Bayesian Multivariate Time Series Methods for Empirical Macroeconomics,” *Foundations and Trends in Econometrics*, 3, 267–358. [74,80]
- Koop, G. M. (2013), “Forecasting with Medium and Large Bayesian VARs,” *Journal of Applied Econometrics*, 28, 177–203. [80]
- Lan, S., and Shahbaba, B. (2016), “Sampling Constrained Probability Distributions Using Spherical Augmentation,” in *Algorithmic Advances in Riemannian Geometry and Applications*, Advances in Computer Vision and Pattern Recognition, eds. H. Q. Minh and V. Murino, pp. 25–71, Cham: Springer. [76]
- Leonard, T., and Hsu, J. S. J. (1992), “Bayesian Inference for a Covariance Matrix,” *The Annals of Statistics*, 20, 1669–1696. [77]
- Marriott, J., Ravishanker, N., Gelfand, A., and Pai, J. (1996), “Bayesian Analysis of ARMA Processes: Complete Sampling-Based Inference Under Exact Likelihoods,” in *Bayesian Analysis in Statistics and Econometrics*, eds. D. A. Berry, K. M. Chaloner, and J. K. Geweke, pp. 243–256, New York: Wiley. [75]
- Marriott, J. M., and Smith, A. F. M. (1992), “Reparameterization Aspects of Numerical Bayesian Methodology for Autoregressive Moving-Average Models,” *Journal of Time Series Analysis*, 13, 327–343. [75]
- Mitchell, T. J., and Beauchamp, J. J. (1988), “Bayesian Variable Selection in Linear Regression,” *Journal of the American Statistical Association*, 832, 1023–1032. [78]
- Monahan, J. F. (1984), “A Note on Enforcing Stationarity in Autoregressive-Moving Average Models,” *Biometrika*, 71, 403–404. [75]
- Morf, M., Vieira, A., and Kailath, T. (1978), “Covariance Characterization by Partial Autocorrelation Matrices,” *The Annals of Statistics*, 6, 643–648. [75]
- Neal, R. M. (2011), “MCMC using Hamiltonian Dynamics,” in *Handbook of Markov Chain Monte Carlo*, Handbooks of Modern Statistical Methods, eds. S. Brooks, A. Gelman, G. Jones, and X.-L. Meng, pp. 113–162, Boca Raton, FL: Chapman & Hall/CRC. [79]
- Neusser, K. (2016), *Time Series Econometrics*, Cham: Springer. [74]
- Paci, L., and Consonni, G. (2020), “Structural Learning Contemporaneous Dependencies in Graphical VAR Models,” *Computational Statistics & Data Analysis*, 144, 106880. [74]
- Philippe, A. (2006), “Bayesian Analysis of Autoregressive Moving Average Processes with Unknown Orders,” *Computational Statistics & Data Analysis*, 51, 1904–1923. [82]
- Piccolo, D. (1982), “The Size of the Stationarity and Invertibility Region of an Autoregressive-Moving Average Process,” *Journal of Time Series Analysis*, 3, 245–247. [75]
- Piironen, J., and Vehtari, A. (2017), “Sparsity Information and Regularization in the Horseshoe and other Shrinkage Priors,” *Electronic Journal of Statistics*, 11, 5018–5051. [78]
- Pourzanjani, A. A., Jiang, R. M., Mitchell, B., Atzberger, P. J., and Petzold, L. R. (2021), “Bayesian Inference Over the Stiefel Manifold via the Givens Representation,” *Bayesian Analysis*, 16, 639–666. [79]
- Roy, A., McElroy, T. S., and Linton, P. (2019), “Constrained Estimation of Causal, Invertible VARMA,” *Statistica Sinica*, 29, 455–478. [75,78,80,82]
- Stan Development Team. (2020), *RStan: The R Interface to Stan*, R package version 2.19.31. [79]
- Stephens, M. (2000), “Bayesian Analysis of Mixture Models with an Unknown Number of Components—An Alternative to Reversible Jump Methods,” *The Annals of Statistics*, 28, 40–74. [82]
- Vermaak, J., Andrieu, C., Doucet, A., and Godsill, S. J. (2004), “Reversible Jump Markov chain Monte Carlo Strategies for Bayesian Model Selection in Autoregressive Processes,” *Journal of Time Series Analysis*, 25, 785–809. [82]