Bayesian Vector Autoregressions: Estimation

Silvia Miranda-Agrippino*
Bank of England
Centre for Macroeconomics (LSE)

Giovanni Ricco†
University of Warwick
OFCE - SciencesPo

Revised 16 August 2018

Summary

Vector Autoregressions (VARs) are linear multivariate time-series models able to capture the joint dynamics of multiple time series. Bayesian inference treats the VAR parameters as random variables, and it provides a framework to estimate 'posterior' probability distribution of the location of the model parameters by combining information provided by a sample of observed data and prior information derived from a variety of sources, such as other macro o micro datasets, theoretical models, other macroeconomic phenomena, or introspection.

In empirical work in Economics and Finance, informative prior probability distributions are often adopted. These are intended to summarise stylised representations of the data generating process. For example, 'Minnesota' priors, one of the most commonly adopted macroeconomic priors for the VAR coefficients, express the belief that an independent random-walk model for each variable in the

*Monetary Analysis, Bank of England, Threadneedle Street, London EC2R 8AH, UK. Email: silvia.miranda-agrippino@bankofengland.co.uk Web: www.silviamirandaagrippino.com

†Department of Economics, The University of Warwick, The Social Sciences Building, Coventry, West Midlands CV4 7AL, UK. Email: G.Ricco@warwick.ac.uk Web: www.giovanni-ricco.com

Acknowledgments. We thank the Editor of the Oxford Research Encyclopedia and an anonymous referee for useful comments and suggestions. We are also grateful to Fabio Canova, Andrea Carriero, Matteo Ciccarelli, Domenico Giannone, Marek Jarociński, Marco del Negro, Massimiliano Marcellino, Giorgio Primiceri, Lucrezia Reichlin, and Frank Shorfheide for helpful comments and discussions. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank of England or any of its Committees.

system is a reasonable 'centre' for the beliefs about their time series behaviour. Other commonly adopted priors, the 'single-unit-root' and the 'sum-of-coefficients' priors are used to enforce beliefs about relations among the VAR coefficients, such as for example the existence of co-integrating relationships among variables, or of independent unit-roots.

Priors for macroeconomic variables are often adopted as 'conjugate prior distributions' – i.e. distributions that yields a posterior distribution in the same family as the prior p.d.f. –, in the form of Normal-Inverse-Wishart distributions that are conjugate prior for the likelihood of a VAR with normally distributed disturbances. Conjugate priors allow direct sampling from the posterior distribution and fast estimation. When this is not possible, numerical techniques such as Gibbs and Metropolis-Hastings sampling algorithms are adopted.

Bayesian techniques allow for the estimation of an ever-expanding class of sophisticated autoregressive models that includes conventional fixed-parameters VAR models; Large VARs incorporating hundreds of variables; Panel VARs, that permit analysing the joint dynamics of multiple time series of heterogeneous and interacting units; and VAR models that relax the assumption of fixed coefficients, such as Time-Varying Parameters, Threshold, and Markov Switching VARs.

Keywords: Bayesian inference, Vector Autoregression Models, BVAR, SVAR, forecasting
JEL Classification: C30, C32

Contents

1 Introduction

Vector Autoregressions (VARs) are linear multivariate time-series models able to capture the joint dynamics of multiple time series. The pioneering work of Sims (1980) proposed to replace the large-scale macroeconomic models popular in the 1960s with VARs, and suggested that Bayesian methods could have improved upon frequentist ones in estimating the model coefficients. Bayesian VARs (BVARs) with macroeconomic variables were first employed in forecasting by Litterman (1979) and Doan et al. (1984). Since then, VARs and BVARs have been a standard macroeconometric tool routinely used by scholars and policy makers for structural analysis, forecasting and scenario analysis in an ever growing number of applications.

The aim of this article is to review key ideas and contributions in the BVAR literature. A companion paper provides a brief survey of applications of BVARs in Economics and Finance, such as forecasting, scenario analysis and structural identification (Miranda-Agrippino and Ricco, 2018). An exhaustive survey of the literature is beyond the scope of this article due to space limitations. Readers are referred to a number of monographs and more detailed surveys available on different topics in the BVARs literature.[1,2]

Differently from frequentist statistics, Bayesian inference treats the VAR parameters as random variables, and provides a framework to update probability distributions about

[1]Several books provide excellent in-depth treatments of Bayesian inference. Among others, Zellner (1971), Gelman et al. (2003), Koop (2003) and Geweke (2005). Canova (2007) provides a book treatment of VARs and BVARs in the context of the methods for applied macroeconomic research. Several recent articles survey the literature on BVARs. Del Negro and Schorfheide (2011) have a deep and insightful discussion of BVAR with a broader focus on Bayesian macroeconometrics and DSGE models. Koop and Korobilis (2010) propose a discussion of Bayesian multivariate time series models with an in-depth discussion of time-varying parameters and stochastic volatility models. Geweke and Whiteman (2006a) and Karlsson (2013b) provide a detailed survey with a focus on forecasting with Bayesian Vector Autoregression. Ciccarelli and Rebucci (2003) survey BVARs in forecasting analysis with Euro Area data. Canova and Ciccarelli (2009, 2013) discuss panel Bayesian VARs. Finally, the reader is referred to Timmermann (2006) for an in-depth discussion on model averaging and forecast combination, a natural extension of the Bayesian framework.

[2]Dieppe et al. (2016) have developed the ready-to-use BEAR toolbox that implements many of the methods described in this article (`https://www.ecb.europa.eu/pub/research/working-papers/html/bear-toolbox.en.html`). Other useful code sources are those related to the books of Kroese and Chan (2014) (`http://joshuachan.org/code.html`) and Koop and Korobilis (2010) (`https://sites.google.com/site/dimitriskorobilis/matlab`).

the unobserved parameters conditional on the observed data. By providing such a framework, the Bayesian approach allows to incorporate prior information about the model parameters into post-sample probability statements. The 'prior' distributions about the location of the model parameters summarise pre-sample information available from a variety of sources, such as other macro o micro datasets, theoretical models, other macroeconomic phenomena, or introspection.

In the absence of pre-sample information, Bayesian VAR inference can be thought of as adopting 'non-informative' (or 'diffuse' or 'flat') priors, that express complete ignorance about the model parameters, in the light of the sample evidence summarised by the likelihood function (i.e. the probability density function of the data as a function of the parameters). Often, in such a case, Bayesian probability statements about the unknown parameters (conditional on the data) are very similar to classical confidence statements about the probability of random intervals around the true parameters value. For example, for a VAR with Gaussian errors and a flat prior on the model coefficients, the posterior distribution is centred at the maximum likelihood estimator (MLE), with variance given by the variance-covariance matrix of the residuals. Section 2 discusses inference in BVARs and 'non-informative' priors.

While non-informative priors can provide a useful benchmark, in empirical work with macroeconomic and financial variables informative priors are often adopted. In scientific data analysis, priors on the model coefficients do not incorporate the investigator's 'subjective' beliefs, instead, they summarise stylised representations of the data generating process. Conditional on a model, these widely held standardised priors aim at making the likelihood-based description of the data useful to investigators with potentially diverse prior beliefs (Sims, 2010b).[3]

The most commonly adopted macroeconomic priors for VARs are the the so-called 'Minnesota' priors (Litterman, 1980). They express the belief that an independent random-walk model for each variable in the system is a reasonable 'centre' for the be-

[3]Bayesian priors can often be interpreted as frequentist penalised regressions (see, for example, De Mol et al., 2008). A Gaussian prior for the regression coefficients, for example, can be thought of as a Ridge penalised regression. Having a double exponential (Laplace) prior on the coefficients is instead equivalent to a Lasso regularisation problem.

liefs about their time series behaviour. While not motivated by economic theory, they are computationally convenient priors, meant to capture commonly held beliefs about how economic time series behave. Minnesota priors can be cast in the form of a Normal-Inverse-Wishart (NIW) prior, which is the conjugate prior for the likelihood of a VAR with normally distributed disturbances (see Kadiyala and Karlsson, 1997). Conjugate priors are such that the posterior distribution belongs to the same family as the prior probability distribution. Hence, they allow for analytical tractability of the posterior, and computational speed. Because the data is incorporated into the posterior distribution only through the sufficient statistics, formulas for updating the prior into the posterior are in this case conveniently simple. It is often useful to think of the parameters of a prior distribution – known as 'hyperparameters' – as corresponding to having observed a certain number of 'dummy' or 'pseudo-' observations with properties specified by the prior beliefs on the VAR parameters. Minnesota priors can be formulated in terms of artificial data featuring pseudo observations for each of the regression coefficients, and that directly assert the prior on them.

Dummy observations can also implement prior beliefs about relations among the VAR coefficients, such as e.g. co-integration among variables. In this case, commonly used priors are formulated directly as linear joint stochastic restrictions among the coefficients.[4] This is, for example, the case of the 'single-unit root' prior, that is centred on a region of the VAR parameter space where either there is no intercept and the system contains at least one unit root, or the system is stationary and close to its steady state at the beginning of the sample (Sims, 1993).[5] Another instance in which dummy observations are used to establish relations among several coefficients is the 'sum-of-coefficients' prior, that incorporates the widely shared prior beliefs that economic variables can be represented by a process with unit roots and weak cross-sectional linkages (Litterman,

[4]In principle, dummy observations can also implement prior beliefs about nonlinear functions of the parameters (a short discussion on this is in Sims, 2005b).

[5]Such a prior is adopted to capture the belief that it is not plausible to assume that initial transients can explain a large part of observed long-run variation in economic time series. Since in a sample of given size there is no information on the behaviour of time series at frequencies longer than the sample size, the prior assumptions implicitly or explicitly elicited in the analysis will inform results. This is a clear example, in the inference in VARs, of an issue for which Bayesian inference provides a framework to make prior information explicit and available to scientific discussion on the inference in VAR models.

1979).[6] Section 3 discusses some of the priors commonly adopted in the economic literature.

The hyperparameters can be either fixed using prior information (and sometimes 'unorthodoxly' using sample information), or associated to hyperprior distributions that express beliefs about their values. A Bayesian model with more than one level of priors is called a hierarchical Bayes model. In empirical macroeconomic modelling, the hyperparameters associated with the informativeness of the prior beliefs (i.e. the tightness of the prior distribution) are usually left to the investigator's judgement. In order to select a value for these hyperparameters, the VAR literature has adopted mostly heuristic methodologies that minimise pre-specified loss functions over a pre-sample (e.g. the out-of-sample mean squared forecast error in Litterman, 1979, or the in-sample fit in Bańbura et al., 2010). Conversely, Giannone et al. (2015) specify hyperprior distributions and choose the hyperparameters that maximise their posterior probability distribution conditional on the data. Section 4 discusses hierarchical modelling and common approaches to choose hyperparameters not specified by prior information.

In Section 5 we discuss Bayesian inference in VAR models that relax the assumption of fixed coefficients in order to capture changes in the time series dynamics of macroeconomic and financial variables, such as VARs with autoregressive coefficients, Threshold and Markov Switching VARs.

Finally, in Section 6 we discuss Panel Bayesian VARs that generalise VAR models by describing the joint dynamics of multiple time series of potentially heterogenous and interacting units – as for examples, the economies of several countries, regions, or sectors.

2 Inference in BVARs

Vector Autoregressions (VARs) are linear stochastic models that describe the joint dynamics of multiple time series. Let y_t be an $n \times 1$ random vector that takes values in

[6]Several sets of pseudo-observations can be adopted at the same time. In fact, successive dummy observations modify the prior distribution as if they reflected successive observations of functions of the VAR parameters, affected by stochastic disturbances.

$\mathbb{R}_n$. The evolution of y_t – the endogenous variables – is described by a system of p-th order difference equations – the VAR(p):

$$y_t = A_1 y_{t-1} + \ldots + A_p y_{t-p} + c + u_t \ . \tag{1}$$

In Eq. (1), A_j, $j = 1, \ldots, p$ are $n \times n$ matrices of autoregressive coefficients, c is a vector of n intercepts, and u_t is an n-dimensional vector of one-step-ahead forecast errors, or reduced-form innovations. The vector of stochastic innovations, u_t, is an independent and identically distributed random variable for each t. The distribution from which u_t is drawn determines the distribution of y_t, conditional on its past $y_{1-p:t-1} \equiv \{y_{1-p}, \ldots, y_0, \ldots, y_{t-2}, y_{t-1}\}$. The standard assumption in the macroeconometric literature is that errors are Gaussian

$$u_t \sim i.i.d.\ \mathcal{N}(0, \Sigma) \ . \tag{2}$$

This implies that also the conditional distribution of y_t is Normal.[7,8]

Bayesian inference on the model in Eq. (1) amounts to updating prior beliefs about the VAR parameters, that are seen as stochastic variables, after having observed a sample $y_{1-p:t} \equiv \{y_{1-p}, \ldots, y_0, \ldots, y_{t-2}, y_t\}$. Prior beliefs about the VAR coefficients are summarised by a probability density function (p.d.f.), and updated using Bayes' Law

$$p(A, \Sigma | y_{1-p:t}) = \frac{p(A, \Sigma) p(y_{1-p:t} | A, \Sigma)}{p(y_{1-p:t})} \propto p(A, \Sigma) p(y_{1-p:t} | A, \Sigma) \ , \tag{3}$$

where we define $A \equiv [A_1, \ldots, A_p, c]'$ as a $k \times n$ matrix, with $k = np + 1$. The joint posterior distribution of the VAR(p) coefficients $p(A, \Sigma | y_{1-p:t})$ incorporates the inform-

[7] While the assumption of normally distributed errors makes the posterior p.d.f. tractable, modern computational methods permit straightforward characterisation of posterior distributions obtained under different assumptions. Among others, Chiu et al. (2017) and Panagiotelis and Smith (2008) depart from the normality assumption and allow for t-distributed errors.

[8] It is interesting to observe that in large samples, and under certain regularity conditions, the likelihood function converges to a Gaussian distribution, with mean at the maximum likelihood estimator (MLE) and covariance matrix given by the usual MLE estimator for the covariance matrix. This implies that conditioning on the MLE and using its asymptotic Gaussian distribution is, approximately in large samples, as good as conditioning on all the data (see discussion in Sims, 2010b).

ation contained in the prior distribution $p(A, \Sigma)$ – summarising the initial information about the model parameters –, and the sample information summarised by $p(y_{1-p:t}|A, \Sigma)$. Viewed as a function of the parameters, the sample information is the likelihood function.[9] The posterior distribution summarises the entire information available, and is used to conduct inference on the VAR parameters.

Given the autoregressive structure of the model, and the i.i.d. innovations, the (conditional) likelihood function of the sample observations $y_{1:T}$ – conditional on A, Σ and on the first p observations $y_{1-p:0}$ –, can be written as the product of the conditional distribution of each observation

$$p\left(y_{1:T}|A, \Sigma, y_{1-p:0}\right) = \prod_{t=1}^{T} p\left(y_t|A, \Sigma, y_{t-p:t-1}\right). \tag{4}$$

Under the assumption of Gaussian errors, the conditional likelihood of the VAR in Eq. (1) is

$$p\left(y_{1:T}|A, \Sigma, y_{1-p:0}\right) = \prod_{t=1}^{T} \frac{1}{(2\pi)^{n/2}} |\Sigma|^{-1} \exp\left\{-\frac{1}{2}\left(y_t - A'x_t\right)' \Sigma^{-1} \left(y_t - A'x_t\right)\right\}, \tag{5}$$

where $x_t' \equiv \begin{bmatrix} y_{t-1}' & \dots & y_{t-p}' & 1 \end{bmatrix}$.

The likelihood in Eq. (5) can be written in compact form, by using the seemingly unrelated regression (SUR) representation of the VAR

$$y = xA + u, \tag{6}$$

where the $T \times n$ matrices y and u and the $T \times k$ matrix x are defined as

$$y = \begin{bmatrix} y_1' \\ \vdots \\ y_T' \end{bmatrix}, \quad x = \begin{bmatrix} x_1' \\ \vdots \\ x_T' \end{bmatrix}, \quad u = \begin{bmatrix} u_1' \\ \vdots \\ u_T' \end{bmatrix}. \tag{7}$$

[9] The marginal p.d.f. for the observations, denoted as $p(y_{1-p:t})$, is a normalising constant and as such can be dropped when making inference about the model parameters.

Using this notation and standard properties of the trace operator, the conditional likelihood function can be equivalently expressed as

$$p\left(y_{1:T}|A,\Sigma,y_{1-p:0}\right)=\frac{1}{(2\pi)^{Tn/2}}|\Sigma|^{-T/2}\exp\left\{-\frac{1}{2}tr\left[\Sigma^{-1}\widehat{S}\right]\right\} \times\exp\left\{-\frac{1}{2}tr\left[\Sigma^{-1}(A-\widehat{A})'x'x(A-\widehat{A})\right]\right\}\ , \quad (8)$$

where $\widehat{A}$ is the maximum-likelihood estimator (MLE) of A, and $\widehat{S}$ the matrix of sums of squared residuals, i.e.

$$\widehat{A}=(x'x)^{-1}x'y, \qquad \widehat{S}=(y-x\widehat{A})'(y-x\widehat{A}). \quad (9)$$

The likelihood can also be written in terms of the vectorised representation of the VAR

$$\mathbf{y}=(\mathbb{I}_n\otimes x)\alpha+\mathbf{u}, \qquad \mathbf{u}\sim(0,\Sigma\otimes\mathbb{I}_T)\ , \quad (10)$$

where $\mathbf{y}\equiv vec(y)$ and $\mathbf{u}\equiv vec(u)$ are $Tn\times 1$ vectors, and $\alpha\equiv vec(A)$ is $nk\times 1$. In this vectorised notation the likelihood function is written as

$$p\left(y_{1:T}|A,\Sigma,y_{1-p:0}\right)=\frac{1}{(2\pi)^{Tn/2}}|\Sigma|^{-T/2}\exp\left\{-\frac{1}{2}tr\left[\Sigma^{-1}\widehat{S}\right]\right\} \times\exp\left\{-\frac{1}{2}(\alpha-\hat{\alpha})'[\Sigma^{-1}\otimes(x'x)](\alpha-\hat{\alpha})\right\}\ , \quad (11)$$

where, consistently, $\hat{\alpha}\equiv vec(\widehat{A})$ is $nk\times 1$. Detailed derivations for the multivariate Gaussian linear regression model can be found in Zellner (1971).

Given the likelihood function, Eq. (3) is used to update the prior information regarding the VAR parameters. An interesting case arises when we assume the absence of any information on the location of the model parameters. This setting can be formalised by assuming that α and Σ are independently distributed, i.e.,

$$p(\alpha,\Sigma)=p(\alpha)p(\Sigma), \quad (12)$$

with prior p.d.f.

$$
\begin{aligned}
p(\alpha) &\propto const., \\
p(\Sigma) &\propto |\Sigma|^{-(n+1)/2} \ .
\end{aligned} \tag{13}
$$

These priors are known as diffuse or Jeffreys' prior (Geisser, 1965; Tiao and Zellner, 1964). Jeffreys priors are proportional to the square root of the determinant of the Fisher information matrix, and are derived from the Jeffreys' 'invariance principle', meaning that the prior is invariant to re-parameterization (see Zellner, 1971).[10]

Given this set of priors, it is straightforward to derive the posterior distribution of the VAR parameters as

$$
\begin{aligned}
p\left(A, \Sigma | y_{1:T}\right) \propto & \\
& |\Sigma|^{-(T+n+1)/2} \exp\left\{-\frac{1}{2} tr\left(\Sigma^{-1} \otimes \mathbb{I}_T\right)\left[\mathbf{y} - \left(\mathbb{I}_n \otimes x\right)\hat{\alpha}\right]'\left[\mathbf{y} - \left(\mathbb{I}_n \otimes x\right)\hat{\alpha}\right]\right\} \\
& \qquad \times \exp\left\{-\frac{1}{2}\left(\alpha - \hat{\alpha}\right)'\left(\Sigma^{-1} \otimes x'x\right)\left(\alpha - \hat{\alpha}\right)\right\} \ ,
\end{aligned} \tag{14}
$$

where the proportionality factor has been dropped for convenience.

From the joint posterior in Eq. (14) one can readily deduce the form of the posterior for α, conditional on Σ and the observed sample. Also, the posterior can be integrated over α to obtain the marginal posterior for Σ. Therefore, it is possible to conveniently write the posterior distribution of the parameters as

$$
p(\alpha, \Sigma | y_{1:T}) = p(\alpha | \Sigma, y_{1:T}) p(\Sigma | y_{1:T}) \tag{15}
$$

[10] 'Non-informative' or 'flat' priors are designed to extract the maximum amount of expected information from the data. They maximise the difference (measured by Kullback-Leibler distance) between the posterior and the prior when the number of samples drawn goes to infinity. Jeffreys priors for VARs are 'improper', in the sense that they do not integrate to one over the parameter space. Hence, they cannot be thought of as well specified p.d.f. distributions. However, they can be obtained as degenerate limit of the Normal-Inverse-Wishart conjugate distribution, and their posterior is proper. For an in-depth discussion on non-informative priors in multi-parameter settings see Zellner (1971) and Bernardo and Smith (2009).

where

$$\Sigma|y \sim \mathcal{IW}\left((y - x\widehat{A})'(y - x\widehat{A}), T - k\right) \tag{16}$$

$$\alpha|\Sigma, y \sim \mathcal{N}\left(\hat{\alpha}, \Sigma \otimes (x'x)^{-1}\right) \ . \tag{17}$$

Hence, given the diffuse priors on α and Σ, the posterior for the autoregressive coefficients is centred at the MLE, with posterior variance $\Sigma \otimes (x'x)^{-1}$.[11] Interestingly, in this standard normal multivariate linear regression model, Bayesian probability statements about the parameters (given the data) have the same form as the frequentist pre-sample probability statements about the parameters' estimator (see also Sims, 2010b). This is a more general property, in fact, Kwan (1998) has shown that, under widely applicable regularity conditions, an estimator $\hat{\alpha}_T$ for which

$$\sqrt{T}(\hat{\alpha}_T - \alpha)|\alpha \xrightarrow[T\to\infty]{\mathcal{D}} \mathcal{N}(0, \Sigma)$$

allows, with high accuracy, to approximate the distribution of $\sqrt{T}(\alpha - \hat{\alpha}_T)|\hat{\alpha}$ as $\mathcal{N}(0, \Sigma)$ in large samples. Hence, it is often possible to interpret $(1 - \rho)$ approximate confidence sets generated from the frequentist asymptotic approximate distribution as if they were sets in the parameter space with posterior probability $(1 - \rho)$.

In potentially misspecified models for which linear regression coefficients are the object of interest, Müller (2013) proposes to adopt an artificial Gaussian posterior centred at the MLE but with a sandwich estimator for the covariance matrix. In fact, in the case of a misspecified model, the shape of the likelihood (the posterior) is asymptotically Gaussian and centred at the MLE, but of a different variance than the asymptotically normal sampling distribution of the MLE. This argument can be seen as a 'flipping' of the frequentist asymptotic statement that supports the use of a sandwich estimator for the covariance matrix in misspecified models, in line with the results in Kwan (1998).[12]

[11] The marginal posterior distribution of the $k \times n$ matrix A is matricvariate t (see Kadiyala and Karlsson, 1997)

$$A|\mathbf{y} \propto |(y - x\widehat{A})'(y - x\widehat{A}) + (A - \widehat{A})'x'x(A - \widehat{A})|^{-T/2} \ . \tag{18}$$

[12] Müller (2013) shows that a Bayesian decision-maker can justify using OLS with a sandwich co-

An important case in which frequentist pre-sample probability statements and Bayesian post-sample probability statements about parameters diverge, is the case of time-series regression models with unit roots. In such cases, while the frequentist distribution of the estimator is skewed asymptotically, the likelihood, and hence the posterior p.d.f., remain unaffected (see Sims and Uhlig, 1991; Kim, 1994).

3 Informative Priors for Reduced-Form VARs

Informative prior probability distributions incorporate information about the VAR parameters that is available before some sample is observed. Such prior information can be contained in samples of past data – from the same or a related system –, or can be elicited from introspection, casual observation, and theoretical models. The first case is sometimes referred to as a 'data-based' prior, while the second as a 'nondata-based' prior.

An important case arises when the prior probability distribution yields a posterior distribution for the parameters in the same family as the prior p.d.f. In this case the prior is called a natural conjugate prior for the likelihood function (Raiffa and Schlaifer, 1961). In general, it has been shown that exponential distributions are the only class of distributions that admit a natural conjugate prior, due to these having a fixed number of sufficient statistics that does not increase as the sample size T increases (see e.g. Gelman et al., 2013). Because the data is incorporated into the posterior distribution only through the sufficient statistics, formulas for updating the prior into the posterior are in these cases conveniently simple.

Prior distributions can be expressed in terms of coefficients, known as hyperparameters, whose functions are sufficient statistics for the model parameters. It is often useful to think of the hyperparameters of a conjugate prior distribution as corresponding to

variance matrix when the probability limit of the OLS estimator is the object of interest, despite the fact that the linear regression model is known not to be the true model (see discussion in Sims, 2010b). Miranda-Agrippino and Ricco (2017) use this intuition to construct coverage bands for impulse responses estimated with Bayesian Local Projections (BLP). This method can be thought of as a generalisation of BVARs that estimates a different model for different forecast horizons – as in direct forecasts – and hence induces autocorrelation in the reduced-form residuals that violate the the i.i.d. assumption in Eq. (61).

having observed a certain number of pseudo-observations with properties specified by the priors on the parameters. In general, for nearly all conjugate prior distributions, the hyperparameters can be interpreted in terms of 'dummy' or pseudo-observations. The basic idea is to add to the observed sample extra 'data' that express prior beliefs about the hyperparameters. The prior then takes the form of the likelihood function of these dummy observations. Hyperparameters can be either fixed using prior information, or associated to hyperprior distributions that express beliefs about their values. A Bayesian model with more than one level of priors is called a hierarchical Bayes model. In this section we review some of the most commonly used priors for VARs with macroeconomic and financial variables, while we discuss the choice of the hyperpriors and hierarchical modelling in Section 4.

3.1 Natural Conjugate Normal-Inverse Wishart Priors

The Normal-Inverse Wishart (NIW) conjugate priors, part of the exponential family, are commonly used prior distributions for (A, Σ) in VARs with Gaussian errors. These assume a multivariate normal distribution for the regression coefficients, and an Inverse Wishart specification for the covariance matrix of the error term, and can be written as

$$\Sigma \sim \mathcal{IW}(\underline{S},\ \underline{d}) \tag{19}$$

$$\alpha|\Sigma \sim \mathcal{N}(\underline{\alpha},\ \Sigma \otimes \underline{\Omega})\ , \tag{20}$$

where $(\underline{S}, \underline{d}, \underline{\alpha}, \underline{\Omega})$ are the priors' hyperparameters. $\underline{d}$ and $\underline{S}$ denote, respectively, the degrees of freedom and the scale of the prior Inverse-Wishart distribution for the variance-covariance matrix of the residuals. $\underline{\alpha}$ is the prior mean of the VAR coefficients, and $\underline{\Omega}$ acts as a prior on the variance-covariance matrix of the dummy regressors.[13] The

[13]The prior mean of the VAR coefficients is $\mathbb{E}[\alpha] = \underline{\alpha}$, for $\underline{d} > n$, while the variance is $\mathbb{V}ar[\alpha] = (\underline{d} - n - 1)^{-1}\underline{S} \otimes \underline{\Omega}$, for $\underline{d} > n + 1$. Setting $\underline{d} = \max\{n + 2, n + 2h - T\}$ ensures that both the prior variances of A and the posterior variances of the forecasts at $T + h$ are defined.

posterior distribution can be analytically derived and is given by

$$\Sigma|\mathbf{y} \sim \mathcal{IW}(\overline{S},\ \overline{d}) \tag{21}$$

$$\alpha|\Sigma, \mathbf{y} \sim \mathcal{N}(\overline{\alpha},\ \Sigma \otimes \overline{\Omega}), \tag{22}$$

where

$$\overline{\Omega} = (\underline{\Omega} + x'x)^{-1}, \tag{23}$$

$$\overline{\alpha} \equiv vec(\overline{A}) = vec\left(\overline{\Omega}\ \left(\underline{\Omega}^{-1}\underline{A} + x'x\widehat{A}\right)\right), \tag{24}$$

$$\overline{S} = \widehat{A}'x'x\widehat{A} + \underline{A}'\underline{\Omega}^{-1}\underline{A} + \underline{S} + (y - x\widehat{A})'(y - x\widehat{A}) - \overline{A}'(\underline{\Omega}^{-1} + x'x)\overline{A}\ . \tag{25}$$

Comparing Eqs. (16) - (17) to Eqs. (19) - (20), it is evident that informative priors can be thought of as equivalent to having observed dummy observations (y_d, x_d) of size T_d, such that

$$\underline{S} = (y_d - x_d\underline{A})'(y_d - x_d\underline{A}), \tag{26}$$

$$\underline{d} = T_d - k, \tag{27}$$

$$\underline{\alpha} = vec(\underline{A}) = vec\left((x_d'x_d)^{-1}x_d'y_d\right), \tag{28}$$

$$\underline{\Omega} = (x_d'x_d)^{-1}\ . \tag{29}$$

This idea was first proposed for a classical estimator for stochastically restricted coefficients by Theil (1963). Once a set of pseudo-observations able to match the wished hyperparameters is found, the posterior can be equivalently estimated using the extended samples $y_* = [y', y_d']'$, $x_* = [x', x_d']'$ of size $T_* = T + T_d$ obtaining

$$\Sigma|\mathbf{y} \sim \mathcal{IW}\left(S_*,\ T_* + \underline{d}\right) \tag{30}$$

$$\alpha|\Sigma, \mathbf{y} \sim \mathcal{N}\left(\alpha_*,\ \Sigma \otimes (x_*'x_*)^{-1}\right). \tag{31}$$

Indeed, it is easy to verify that the posterior moments obtained with the starred variables coincide with those in Eqs. (21) - (22). The posterior estimator efficiently combines

sample and prior information using their precisions as weights in the spirit of the mixed estimation of Theil and Goldberger (1961). Posterior inference can be conducted via direct sampling.

Algorithm 1: Direct Monte Carlo Sampling from Posterior of VAR Parameters.

For $s = 1, \ldots, n_{sim}$:

1. Draw $\Sigma^{(s)}$ from the Inverse-Wishart distribution $\Sigma|\mathbf{y} \sim \mathcal{IW}\left(S_*,\ T_* + \underline{d}\right)$.
2. Draw $A^{(s)}$ from the Normal distribution of $A^{(s)}|\Sigma^{(s)}, \mathbf{y} \sim \mathcal{N}\left(\alpha_*,\ \Sigma^{(s)} \otimes (x_*'x_*)^{-1}\right)$.

When it is not possible to sample directly from the posterior distribution, as in this case, Markov chain Monte Carlo (MCMC) algorithms are usually adopted (see e.g. Chib, 2001).[14]

An important feature of the NIW priors in Eqs. (19) - (20) is the Kronecker factorisation that appears in the Gaussian prior for α. As discussed in the previous section, because the same set of regressors appears in each equation, homoskedastic VARs can be written as SUR models. This symmetry across equations means that homoskedastic VAR models have a Kronecker factorisation in the likelihood, which in turn implies that estimation can be broken into n separate least-squares calculations, each only of dimension $np + 1$. The symmetry in the likelihood can be inherited by the posterior, if the prior adopted also features a Kronecker structure as in Eq. (20). This is a desirable property that guarantees tractability of the posterior p.d.f. and computational speed. However, such a specification can result in unappealing restrictions and may not fit the

[14]The key idea of MCMC algorithms is to construct a Markov chain for $\theta \equiv (A, \Sigma)$ which has the posterior as its (unique) limiting stationary distribution, and such that random draws can be sampled from the transition kernel $p(\theta^{(s+1)}|\theta^{(s)})$. Tierney (1994) and Geweke (2005) discuss the conditions for the convergence of the chain to the posterior distribution when starting from an arbitrary point in the parameter space. Typically, a large number of initial draws (known as burn-in sample) is discarded to avoid including portions of the chain which have not yet converged to the posterior. Also, even if convergent, the chain may move very slowly in the parameter space due to e.g. autocorrelation between the draws, and a very large number of draws may be needed. See also Karlsson (2013a) for a discussion on this point and on empirical diagnostic tests to assess the chain convergence. References include Geweke (1999); Chib and Greenberg (1995); Geweke and Whiteman (2006b).

actual prior beliefs one has – see discussions in Kadiyala and Karlsson (1997), and Sims and Zha (1998). In fact, it forces symmetry across equations, because the coefficients of each equation have the same prior variance matrix (up to a scale factor given by the elements of Σ). There may be situations in which theory suggests 'asymmetric restrictions' may be desirable instead, e.g. money neutrality implies that the money supply does not Granger-cause real output.[15] Also, the Kronecker structure implies that prior beliefs must be correlated across the equations of the reduced form representation of the VAR, with a correlation structure that is proportional to that of the disturbances.

3.2 Minnesota Prior

In macroeconomic and financial applications, the parameters of the NIW prior in Eqs. (19) - (20) are often chosen so that prior expectations and variances of A coincide with the so-called 'Minnesota' prior, that was originally proposed in Litterman (1980, 1986).[16] The basic intuition behind this prior is that the behaviour of most macroeconomic variables is well approximated by a random walk with drift. Hence, it 'centres' the distribution of the coefficients in A at a value that implies a random-walk behaviour for all the elements in y_t

$$y_t = c + y_{t-1} + u_t. \tag{32}$$

While not motivated by economic theory, these are computationally convenient priors, meant to capture commonly held beliefs about how economic time series behave.

The Minnesota prior assumes the coefficients $A_1, \ldots, A_p$ to be a priori independent

[15] Such restrictions can be accommodated by replacing Eq. (19) with a truncated Normal distribution. In this case, however, posterior moments are not available analytically and must be evaluated numerically, with consequential complications and loss of efficiency with respect to the MCMC algorithm discussed above (see Hajivassiliou and Ruud, 1994; Kadiyala and Karlsson, 1997, for further details).

[16] The original formulation of Litterman (1980)'s prior was of the form

$$\alpha \sim \mathcal{N}(\underline{\alpha}, \underline{\Gamma}),$$

where $\underline{\Gamma} \equiv diag([\underline{\gamma}_1^2, \ldots, \underline{\gamma}_n^2])$ is assumed to be fixed, known, and diagonal. Highfield (1992) and Kadiyala and Karlsson (1997) observed that by modifying Litterman's prior to make it symmetric across equations in the form of a NIW prior, the posterior p.d.f. was tractable.

and normally distributed, with the following moments

$$\mathbb{E}\left[(A_\ell)_{ij}|\Sigma\right] = \begin{cases} \delta_i & i=j,\ \ell=1 \\ 0 & \text{otherwise} \end{cases} \qquad \mathbb{V}ar\left[(A_\ell)_{ij}|\Sigma\right] = \begin{cases} \frac{\lambda_1^2}{f(\ell)} & \text{for } i=j, \forall \ell \\ \frac{\lambda_1^2}{f(\ell)}\frac{\Sigma_{ij}}{\omega_j^2} & \text{for } i\neq j, \forall \ell. \end{cases} \tag{33}$$

In Eq. (33), $(A_\ell)_{ij}$ denotes the coefficient of variable j in equation i at lag ℓ. In the original formulation of the prior $\delta_i = 1$, in accordance with Eq. (32). The random-walk assumption, however, may not be appropriate if the variables in y_t were characterised by substantial mean-reversion. For stationary series, or series that have been transformed to achieve stationarity, Bańbura et al. (2010) centre the distribution around zero (i.e. $\delta_i = 0$). The prior also assumes that lags of other variables are less informative than own lags, and that most recent lags of a variable tend to be more informative than more distant lags. This intuition is formalised with $f(\ell)$. A common choice for this function is a harmonic lag decay – i.e. $f(\ell) = \ell^{\lambda_2}$, a special case of which is $f(\ell) = \ell$ –, where the severity of the lag decay is regulated by the hyperparameter λ_2. The factor Σ_{ij}/ω_j^2 accounts for the different scales of variables i and j. The hyperparameters ω_j^2 are often fixed using sample information, for example from univariate regressions of each variable onto its own lags.

Importantly, λ_1 is a hyperparameter that controls the overall tightness of the random walk prior. If $\lambda_1 = 0$ the prior information dominates, and the VAR reduces to a vector of univariate models. Conversely, as $\lambda_1 \to \infty$ the prior becomes less informative, and the posterior mostly mirrors sample information. We discuss the choice of the free hyperparameters in Section 4.

The Minnesota prior can be implemented using dummy observations. Priors on the A coefficients are implemented via the following pseudo-observations

$$\begin{aligned} y_d^{(1)} &= \begin{bmatrix} diag([\delta_1\omega_1, \ldots, \delta_n\omega_n])/\lambda_1 \\ 0_{n(p-1)\times n} \end{bmatrix}, \\ x_d^{(1)} &= \begin{bmatrix} J_p \otimes diag([\omega_1, \ldots, \omega_n])/\lambda_1 & 0_{np\times 1} \end{bmatrix}, \end{aligned} \tag{34}$$

where $J_p = diag([1^{\lambda_2}, 2^{\lambda_2}, \ldots, p^{\lambda_2}])$ with geometric lag decay.[17] To provide intuition on how the prior is implemented using artificial observations, we consider the simplified case of a $n = 2,\ p = 2$ VAR for the pseudo-observations . The first n rows of Eq. (34) impose priors on A_1; that is, on the coefficients of the first lag. In the $n = 2,\ p = 2$ case one obtains,

$$\begin{pmatrix} \frac{\delta_1\omega_1}{\lambda_1} & 0 \\ 0 & \frac{\delta_2\omega_2}{\lambda_1} \end{pmatrix} = \begin{pmatrix} \frac{\omega_1}{\lambda_1} & 0 & 0 & 0 & 0 \\ 0 & \frac{\omega_2}{\lambda_1} & 0 & 0 & 0 \end{pmatrix} A + \begin{pmatrix} (u_d^{(1)})_{1,1} & (u_d^{(1)})_{2,1} \\ (u_d^{(1)})_{1,2} & (u_d^{(1)})_{2,2} \end{pmatrix} \tag{35}$$

that implies, for example, the following equations for the elements $(1,1)$ and $(1,2)$ of A_1

$$\frac{\delta_1\omega_1}{\lambda_1} = \frac{\omega_1}{\lambda_1}(A_1)_{1,1} + (u_d^{(1)})_{1,1} \implies (A_1)_{1,1} \sim \mathcal{N}\left(\delta_1, \frac{\Sigma_{1,1}\lambda_1^2}{\omega_1^2}\right),$$

$$0 = \frac{\omega_1}{\lambda_1}(A_1)_{2,1} + (u_d^{(1)})_{2,1} \implies (A_1)_{2,1} \sim \mathcal{N}\left(0, \frac{\Sigma_{2,1}\lambda_1^2}{\omega_1^2}\right).$$

Similar restrictions are obtained for the elements the elements $(2,1)$ and $(2,2)$ of A_1. The following $(n-1)p$ rows in Eq. (34) implement priors on the coefficients of the other lags. In fact, we readily obtain

$$\begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix} = \begin{pmatrix} 0 & 0 & \frac{2^{\lambda_2}\omega_1}{\lambda_1} & 0 & 0 \\ 0 & 0 & 0 & \frac{2^{\lambda_2}\omega_2}{\lambda_1} & 0 \end{pmatrix} A + \begin{pmatrix} (u_d^{(1)})_{1,1} & (u_d^{(1)})_{2,1} \\ (u_d^{(1)})_{1,2} & (u_d^{(1)})_{2,2} \end{pmatrix} \tag{36}$$

which for example implies the following restriction for the element $(1,1)$ of A_2

$$0 = \frac{2^{\lambda_2}\omega_1}{\lambda_1}(A_2)_{1,1} + (u_d^{(1)})_{1,1} \implies (A_2)_{1,1} \sim \mathcal{N}\left(0, \frac{\Sigma_{1,1}\lambda_1^2}{2^{2\lambda_2}\omega_1^2}\right).$$

Similar restrictions obtain for the other elements of A_2. Priors beliefs on the residual

[17] Given the dummy observations in Eq. (34), the matrix Ω in Eq. (19) is diagonal and of the form

$$\underset{[k\times k]}{\Omega} = ((x_d^{(1)})' x_d^{(1)})^{-1} = diag\left(\left[\frac{\lambda_1^2}{\omega_1^2}, \ldots, \frac{\lambda_1^2}{\omega_n^2}, \frac{\lambda_1^2}{2^{2\lambda_2}\omega_1^2}, \ldots, \frac{\lambda_1^2}{2^{2\lambda_2}\omega_n^2}, \ldots, \frac{\lambda_1^2}{p^{2\lambda_2}\omega_1^2}, \ldots, \frac{\lambda_1^2}{p^{2\lambda_2}\omega_n^2}, 0\right]\right).$$

covariance matrix Σ can instead be implemented by the following block of dummies

$$y_d^{(2)} = \left[\begin{array}{c} 1_{\lambda_3 \times 1} \otimes diag([\omega_1, \ldots, \omega_n]) \end{array} \right] \tag{37}$$

$$x_d^{(2)} = \left[\begin{array}{cc} 0_{\lambda_3 n \times np} & 0_{\lambda_3 n \times 1} \end{array} \right] . \tag{38}$$

In the $n = 2,\ p = 2$ case, they correspond to appending to the VAR equations λ_3 replications of

$$\begin{pmatrix} \omega_1 & 0 \\ 0 & \omega_2 \end{pmatrix} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix} A + \begin{pmatrix} (u_d^{(2)})_{1,1} & (u_d^{(2)})_{2,1} \\ (u_d^{(2)})_{1,2} & (u_d^{(2)})_{2,2} \end{pmatrix} . \tag{39}$$

λ_3 is the hyperparameter that determines the tightness of the prior on Σ. To understand how this works, it is sufficient to consider that given λ_3 artificial observations z_i, with $z_i \sim \mathcal{N}(0, \sigma_z^2)$, an estimator for the covariance is given by $\lambda_3^{-1} \sum_{i=1}^{\lambda_3} z_i^2$.

Finally, uninformative priors for the intercept are often implemented with the following set of pseudo-observations

$$y_d^{(3)} = \left[\begin{array}{c} 0_{1 \times n} \end{array} \right], \qquad x_d^{(3)} = \left[\begin{array}{cc} 0_{1 \times np} & \epsilon \end{array} \right],$$

where ϵ is a hyperparameter usually set to a very small number.[18]

3.3 Priors for VAR with Unit Roots and Trends

Sims (1996, 2000) observed that flat-prior VARs, or more generally estimation methods that condition on initial values, tend to attribute an implausibly large share of the variation in observed time series to deterministic – and hence entirely predictable – components. The issue stems from the fact that ML and OLS estimators that condition on the initial observations and treat them as non-stochastic do not apply any penalisation to parameters values that imply that these observations are very distant from the variables' steady state (or their trend if non-stationary). As a consequence,

[18] Canova (1992, 1993) propose a set of artificial observations to account for seasonal patterns and potentially other peaks in the spectral densities.

complex transient dynamics from the initial conditions to the steady state are treated as plausible, and can explain an 'implausibly' large share of the low-frequency variation of the data. This typically translates into poor out-of-sample forecasts. To understand the intuition, consider the univariate model

$$y_t = c + ay_{t-1} + u_t. \tag{40}$$

Iterating Eq. (40) backward yields

$$y_t = \left[a^t y_0 + \sum_{j=0}^{t-1} a^j c\right] + \left[\sum_{j=0}^{t-1} a^j u_{t-j}\right], \tag{41}$$

which, if $|a| < 1$, reduces to

$$y_t = \left[a^t \left(y_0 - \frac{c}{1-a}\right) + \frac{c}{1-a}\right] + \left[\sum_{j=0}^{t-1} a^j u_{t-j}\right]. \tag{42}$$

The first term in square brackets in Eq. (41) is the deterministic component: the evolution of y_t from the initial conditions y_0, absent any shocks. The second term instead captures the stochastic evolution of y_t due to the shocks realised between $[0, t-1]$. $c/(1-a)$ in Eq. (42) is the unconditional mean of y_t. If y_t is close to non-stationary – i.e. $a \simeq 1$ –, the MLE estimator of the unconditional mean of y_t may be very far from y_0, and the 'reversion to the mean' from y_0 is then used to fit the data (see Eq. 42).

One way to deal with this issue is to use the unconditional likelihood, by explicitly incorporating the density of the initial observations in the inference. However, because most macroeconomic time series are effectively nonstationary, it is not obvious how the density of the initial observations should be specified.[19] Another approach, following Sims and Zha (1998); Sims (2000), is to instead specify priors that downplay the im-

[19]This approach requires the use of iterative nonlinear optimisation methods. The main issue with this approach is that nonstationary models have no unconditional – viz. ergodic – distribution of the initial conditions. Also, while near-nonstationary models may have an ergodic distribution, the time required to arrive at the ergodic distribution from arbitrary initial conditions may be very long. For this reason, using such a method requires strong beliefs about the stationarity of the model, which is rarely the case in macroeconomics, and imposing the ergodic distribution on the first p observations may be unreasonable (see Sims, 2005a).

portance of the initial observations, and hence reduce the explanatory power of the deterministic component.

These types of priors, implemented through artificial observations, aim to reduce the importance that the deterministic component has in explaining a large share of the in-sample variation of the data, eventually improving forecasting performances out-of-sample (see Sims, 1996; Sims and Zha, 1998, for a richer discussion on this point).[20]

The 'co-persistence' (or 'single-unit-root' or 'dummy initial observation') prior (Sims, 1993) reflects the belief that when all lagged y_t's are at some level $\bar{y}_0$, y_t tends to persist at that level. It is implemented using the following artificial observation

$$\underset{[1\times n]}{y_d^{(4)}} = \left[\frac{\bar{y}_{0,1}}{\lambda_4}, \ldots, \frac{\bar{y}_{0,n}}{\lambda_4}\right] \qquad \underset{[1\times k]}{x_d^{(4)}} = [y_d^{(4)}, \ldots, y_d^{(4)}, 1/\lambda_4], \tag{43}$$

where $\bar{y}_{0,i},\ i = 1, \ldots, n$ are the average of the initial values of each variable, and usually set to be equal to the average of the first p observations in the sample, and $k = np + 1$. Writing down the implied system of equations $y_d^{(4)} = Ax_d^{(4)} + u_d^{(4)}$ one obtains the following stochastic restriction on the VAR coefficients

$$[\mathbb{I}_n - A(1)]\,\bar{y}_0 - c = \lambda_4 u_d^{(4)}, \tag{44}$$

where $\mathbb{I}_n - A(1) = (\mathbb{I}_n - A_1 - \ldots - A_p)$. The hyperparameter λ_4 controls the tightness of this stochastic constraint. The prior is uninformative for $\lambda_4 \to \infty$. Conversely, as $\lambda_4 \to 0$ the model tends to a form where either there is at least one explosive common unit root and the constant c is equal to zero ($\bar{y}_0$ is the eigenvector of the unit root), or the VAR is stationary, c is different from zero, and the initial conditions are close to the implied unconditional mean ($\bar{y}_0 = [\mathbb{I}_n - A(1)]^{-1} c$). In the stationary form, this prior does not rule out cointegrated models. This prior induces prior correlation among all the VAR coefficients in each equation, including the constant.[21]

[20] The treatment of unit root in Bayesian and frequentist inference has been hotly debated. Among others, important contributions are Sims (1988, 1991), Sims and Uhlig (1991), Koop and Steel (1991), Phillips (1991a,b), Uhlig (1994a,b), Müller and Elliott (2003); Jarociński and Marcet (2011, 2014). The Journal of Applied Econometrics October/December 1991 Volume 6, Issue 4 has been entirely dedicated to this debate.

[21] To put a heavier weight on the presence of a unit root, one could add to the observation in Eq.

The 'sums-of-coefficients' (or 'no-cointegration') prior (Doan et al., 1984), captures the belief that when the average lagged values of a variable $y_{j,t}$ is at some level $\bar{y}_{0,j}$, then $\bar{y}_{0,j}$ is likely to be a good forecast of $y_{j,t}$. It also implies that knowing the average of lagged values of variable j does not help in predicting a variable $i \neq j$. This prior is implemented using n artificial observations, one for each variable in y_t

$$\underset{[n\times n]}{y_d^{(5)}} = diag\left(\left[\frac{\bar{y}_{0,1}}{\lambda_5},\ldots,\frac{\bar{y}_{0,n}}{\lambda_5}\right]\right) \qquad \underset{[n\times k]}{x_d^{(5)}} = [y_d^{(5)},\ldots,y_d^{(5)},0]. \tag{45}$$

The prior implied by these dummy observations is centred at 1 for the sum of coefficients on own lags for each variable, and at 0 for the sum of coefficients on other variables' lags. It also introduces correlation among the coefficients of each variable in each equation. In fact, it is easy to show that equation by equation this priors implies the stochastic constraint

$$[1-(A_1)_{jj}-\ldots-(A_p)_{jj}]\,\bar{y}_{0,j} = \lambda_5(u_d^{(5)})_j \qquad \forall j\ , \tag{46}$$

where $(A_\ell)_{jj}$ denotes the coefficient of variable j in equation j at lag ℓ. The hyperparameter λ_5 controls the variance of these prior beliefs. As $\lambda_5 \to \infty$ the prior becomes uninformative, while $\lambda_5 \to 0$ implies that each variable is an independent unit-root process, and there are no co-integration relationships.[22]

The Bayesian analysis of cointegrated VARs is an active area of research, (a detailed survey is in Koop et al. 2006).[23] Giannone et al. (2016) elicit theory-based priors for the long run of persistent variables which shrink towards a random walk those linear

(43) an additional artificial observation that enforces the belief that $c = 0$. Alternatively, one could modify Eq. (43) to have a zero in place of λ_4^{-1} as the observation corresponding to the intercept. In this case, the prior gives no plausibility to stationary models and, if used in isolation, allows for at least a single unit root without any restriction on c. Hence, despite the presence of a unit root, it may not necessarily reduce the importance of the deterministic component (see Sims, 2005a).

[22]The sums-of-coefficients observations of Eq. (45) do not imply any restriction on the vector of intercepts c, since the artificial observations loading on the constant are set to zero. Therefore, this prior allows for a non-zero constant, and hence for a linearly trending drift. To assign smaller probability to versions of the model in which deterministic transient components are much more important than the error term in explaining the series variance, one has to add to Eq. (45) artificial observations that favour $c = 0$ (see Sims, 2005a).

[23]Among many others, contributions to the treatment of cointegration in Bayesian VARs are in Kleibergen and van Dijk (1994), Geweke (1996), Villani (2001), Kleibergen and Paap (2002), Strachan and Inder (2004), Koop et al. (2011), Jochmann and Koop (2015).

combination of variables that are likely to have a unit root. Conversely, combinations which are likely to be stationary (i.e. cointegrating relationships among variables) are shrunk towards stationary processes. Operationally, this is achieved by rewriting the VAR in Eq. (1) as

$$\begin{aligned}\Delta y_t &= \Pi y_{t-1} + P_1 \Delta y_{t-1} + \ldots + P_p \Delta y_{t-p+1} + c + \xi_t \\ &= \Pi F^{-1} F y_{t-1} + P_1 \Delta y_{t-1} + \ldots + P_p \Delta y_{t-p+1} + c + \xi_t, \end{aligned} \tag{47}$$

where $\Pi = A_1 + \ldots + A_p - \mathbb{I}_n$, $P_j = -(A_{j+1} + \ldots + A_p)$, and F is any invertible n-dimensional matrix. The problem is then specified as setting a prior for $\widetilde{\Pi} \equiv \Pi F^{-1}$, conditional on a specific choice of F. F defines the relevant linear combinations of the variables in y_t which macroeconomic theory suggest to be a priori stationary or otherwise.

Another alternative is in Villani (2009). Here the VAR is written as

$$y_t = \rho_0 + \rho_1 t + \tilde{y}_t, \qquad \tilde{y}_t = A_1 \tilde{y}_{t-1} + \ldots + A_p \tilde{y}_{t-p} + u_t, \qquad u_t \sim i.i.d. \mathcal{N}(0, \Sigma) \tag{48}$$

where ρ_0 and ρ_1 are $n \times 1$ vectors. The first term, $\rho_0 + \rho_1 t$, captures a linear deterministic trend of y_t, whereas the law of motion of $\tilde{y}_t$ captures stochastic fluctuations around the deterministic trend, which can be either stationary or non-stationary. This alternative specification allows to separate beliefs about the deterministic trend component from beliefs about the persistence of fluctuations around this trend. Let $A = [A_1, \ldots, A_p]'$ and $\rho = [\rho_0', \rho_1']'$. It can be shown that if the prior distribution of ρ conditional on A and Σ is Normal, the (conditional) posterior distribution of ρ is also Normal (see also Del Negro and Schorfheide, 2011, for details). Hence, posterior inference can be implemented via Gibbs sampling.

3.4 Priors from Structural Models

DeJong et al. (1993), Ingram and Whiteman (1994), Del Negro and Schorfheide (2004) have proposed the use of priors for VARs that are derived from Dynamic Stochastic

General Equilibrium (DSGE) models. This approach bridges VARs and DSGEs by constructing families of prior distributions informed by the restrictions that a DSGE-model implies on the VAR coefficients. This modelling approach is sometimes referred to as DSGE-VAR. Ingram and Whiteman (1994) derive prior information from the basic stochastic growth model of King et al. (1988) and report that a BVAR based on the Real Business Cycle model prior outperforms a BVAR with a Litterman prior in forecasting real economic activity. Del Negro and Schorfheide (2004) extend and generalise this approach, and show how to conduct policy simulations within this framework.

Schematically, the exercises can be thought of as follows. First, time-series are simulated from a DSGE model. Second, a VAR is estimated from these simulated data. Population moments of the simulated data computed from the DSGE model solution are considered in place of sample moments. Since the DSGE model depends on unknown structural parameters, hierarchical prior modelling is adopted by specifying a distribution on the DSGE model parameters. A tightness parameter controls the weight of the DSGE model prior relative to the weight of the actual sample. Finally, Markov Chain Monte Carlo methods are used to generate draws from the joint posterior distribution of the VAR and DSGE model parameters.

3.5 Priors for Model Selection

It is standard practice in VAR models to pre-select the relevant variables to be included in the system (and with how many lags). This procedure may be thought of as having dogmatic priors about which variables have non-zero coefficients in the system. The challenge is in selecting among an expansive set of potential models. Indeed, for a VAR with n endogenous variables, q additional potentially exogenous variables including a constant, and p lags, there are $2^{(q+pn)n+n(n-1)/2}$ possible models.

Jarociński and Maćkowiak (2017) propose to select the variables to be included in the system by systematically assessing the posterior probability of 'Granger causal priority' (Sims, 2010a) in a BVAR with conjugate priors. Granger causal priority answers questions of the form "Is variable z relevant for variable x, after controlling for other

variables in the system?" The authors provide a closed form expression for the posterior probability of Granger causal priority, and suggest that variables associated with high Granger causal priority probabilities can be omitted from a VAR with the variables of interest.

Alternatively, one can adopt priors that support model selection and enforce sparsity. A variety of techniques, including double exponential (Laplace) prior, spike-and-slab prior, etc., have been adopted to handle this issue. Some recent theoretical and empirical contributions on this topic are in Mitchell and Beauchamp (1988), George et al. (2008), Koop (2013), Korobilis (2013), Bhattacharya et al. (2015), Griffin and Brown (2010, 2017), Giannone et al. (2017), Huber and Feldkircher (2017).

4 Hyperpriors and Hierarchical Modelling

As seen in the previous section, the informativeness of prior beliefs on the VAR parameters often depends on a set of free hyperparameters. Let $\lambda \equiv [\lambda_1, \lambda_2, \ldots]$ denote the vector collecting all the hyperparameters not fixed using (pre)sample information, and θ denote all the VAR parameters, i.e. A and Σ. The prior distribution of θ is thus effectively $p_\lambda(\theta)$. Choosing a value for λ alters the tightness of the prior distribution, and hence determines how strictly the prior is enforced on the data.

In order to set the informativeness of the prior distribution of the VAR coefficients, the literature has initially used mostly heuristic methodologies. Litterman (1980) and Doan et al. (1984), for example, choose a value for the hyperparameters that maximises the out-of-sample forecasting performance over a pre-sample. Conversely, Bańbura et al. (2010) propose to choose the shrinkage parameters that yield a desired in-sample fit, in order to control for overfitting. Subsequent studies have then either used these as 'default' values, or adopted either one of these criteria. Robertson and Tallman (1999); Wright (2009); Giannone et al. (2014) opt for the first, while e.g. Giannone et al. (2008); Bloor and Matheson (2011); Carriero et al. (2009); Koop (2013) follow Bańbura et al. (2010).

In VARs, Giannone et al. (2015) observe that, from a purely Bayesian perspective,

choosing λ is conceptually identical to conducting inference on any other unknown parameter of the model. Specifically, the model is interpreted as a hierarchical one (Berger, 1985; Koop, 2003) and λ can be chosen as the maximiser of

$$\begin{aligned} p(\lambda|\mathbf{y}) &\propto \int p(\mathbf{y}|\theta, \lambda, y_{1-p:0}) p(\theta|\lambda) d\theta \cdot p(\lambda) \\ &= p(\mathbf{y}|\lambda, y_{1-p:0}) \cdot p(\lambda) \ . \end{aligned} \tag{49}$$

This method is also known in the literature as the Maximum Likelihood Type II (ML-II) approach to prior selection (Berger, 1985; Canova, 2007). In Eq. (49), $p(\lambda|\mathbf{y})$ is the posterior distribution of λ conditional on the data, and $p(\lambda)$ denotes a prior probability density specified on the hyperparameters themselves, and also known as the hyperprior distribution. In such hierarchical model, the prior distribution for the VAR coefficients is treated as a conditional prior, that is $p_\lambda(\theta)$ is replaced by $p(\theta|\lambda)$. In the case of a NIW family of distributions, the prior structure becomes $p(\alpha|\Sigma, \lambda)p(\Sigma|\lambda)p(\lambda)$. $p(\mathbf{y}|\lambda, y_{1-p:0})$ is the marginal likelihood (ML), and is obtained as the density of the data as a function of λ, after integrating out all the VAR parameters. Conveniently, with conjugate priors the ML is available in closed form.

Conversely, the joint posterior of α, Σ and λ is not available in closed form. However, with NIW priors for θ, Giannone et al. (2015) set up the following Metropolis-Hasting sampler for the joint distribution

Algorithm 2: MCMC Sampler for a VAR with Hierarchical Prior.
For $s = 1, \ldots, n_{sim}$:

1. Draw a candidate vector λ^* from the random walk distribution $\lambda^* \sim \mathcal{N}(\lambda^{s-1}, \kappa H^{-1})$, where H is the Hessian of the negative of the log-posterior at the peak for λ, and κ is a tuning constant. Choose

$$\lambda^{(s)} = \begin{cases} \lambda^* & \text{with probability } = \min\left\{1, \frac{p(\mathbf{y}|\lambda^*)}{p(\mathbf{y}|\lambda^{(s-1)})}\right\} \\ \lambda^{(s-1)} & \text{otherwise.} \end{cases}$$

2. Draw $\Sigma^{(s)}$ form the full conditional posterior $\Sigma|\mathbf{y}, \lambda^{(s)}$ in Eq. (21).

3. Draw $A^{(s)}$ from the full conditional posterior $A^{(s)}|\mathbf{y}, \Sigma^{(s)}, \lambda^{(s)}$ in Eq. (22).

In a similar fashion, Belmonte et al. (2014) apply a hierarchical structure to time-varying parameters (TVP) models and specify priors for Bayesian Lasso shrinkage parameters to determine whether coefficients in a forecasting model for inflation are zero, constant, or time-varying in a data driven way.

Carriero et al. (2015) evaluate the forecasting performance of BVARs where tightness hyperparameters are chosen as the maximisers of Eq. (49) or rather set to default values and find that the former route yields modest but statistically significant gains in forecasting accuracy particularly at short horizons.

5 Time-Varying Parameter, State-Dependent, Stochastic Volatility VARs

Models that allow parameters to change over time are increasingly popular in empirical research, in recognition of the fact that they can capture structural changes in the economy. In fact, it seems to be a common belief that the properties of many (if not most) macroeconomic time series have changed over time, and can change across regimes or phases of the business cycle. Model parameters either change frequently and gradually over time according to a multivariate autoregressive process – as in e.g. in Time-Varying Parameters VARs (TVP-VARs) –, or they change abruptly and infrequently as in e.g. Markov-switching or structural-break models.

5.1 Time-varying parameters VAR (TVP-VAR)

Time-varying parameters VARs differ from fixed-coefficient VARs in that they allow the parameters of the model to vary over time, according to a specified law of motion.[24] TVP-VARs often include also stochastic volatility (SV), which allows for time variation

[24]Review articles are in Del Negro and Schorfheide (2011); Koop and Korobilis (2010); Lubik and Matthes (2015).

in the variance of the stochastic disturbances.[25] Doan et al. (1984) were first to show how estimation of a TVP-VAR with Litterman priors could be conducted by casting the VAR in state space form and using Kalman filtering techniques. This same specification is in Sims (1993). Bayesian time varying parameter VARs have become popular in empirical macroeconomics following the work of Cogley and Sargent (2002, 2005) and Primiceri (2005) who provided the foundations for Bayesian inference in these models, and used then innovations in MCMC algorithms to improve on their computational feasibility.

The basic TVP-VAR is of the form

$$y_t = A_{1,t}y_{t-1} + \ldots + A_{p,t}y_{t-p} + c_t + u_t \ , \tag{50}$$

where the constant coefficients of Eq. (1) are replaced by the time-varying $A_{j,t}$. Eq. (50) can be rewritten in compact form as

$$y_t = x_t A_t + u_t \ , \tag{51}$$

where x_t is defined as in Eq. (5), and $A_t = [A_{1,t}, \ldots, A_{p,t}, c_t]'$. It is common to assume that the coefficients follow a random-walk process

$$\alpha_t = \alpha_{t-1} + \varsigma_t \qquad \varsigma_t \sim i.i.d.\ \mathcal{N}(0, \Upsilon) \ , \tag{52}$$

where $\alpha_t \equiv vec(A_t)$. The covariance matrix Υ is usually restricted to be diagonal, and the innovations ς_t to be uncorrelated with u_t, with u_t distributed as in Eq. (61). The law of motion for α_t in Eq. (52) – i.e. the state equation –, implies that $\alpha_{t+1}|\alpha_t, \Upsilon \sim \mathcal{N}(\alpha_t, \Upsilon)$, which can be used as a prior distribution for α_{t+1}. Hence, the prior for all the states (i.e. $\alpha_t \ \forall t$) is a product of normal distributions. For the initial vector of the VAR coefficients Cogley and Sargent (2002, 2005) use a prior of the form $\alpha_1 \sim \mathcal{N}(\underline{\alpha}_{1|0}, \underline{\Upsilon}_{1|0})$, where $\underline{\alpha}_{1|0}$ and $\underline{\Upsilon}_{1|0}$ are set by estimating a fixed-coefficient VAR with a flat prior on a pre-sample.[26] If the Gaussian prior for the states is complemented with IW priors for

[25] Stochastic volatility in Bayesian VARs was initially introduced in Uhlig (1997).

[26] See also the discussion in Karlsson (2013a) for additional details on the specification of the prior for α_t.

both Σ and Υ, then sampling from the joint posterior is possible with a Gibbs sampling algorithm

Algorithm 3: Gibbs Sampling from Posterior of TVP-VAR Parameters. Select starting values for $\Sigma^{(0)}$ and $\Upsilon^{(0)}$. For $s = 1, \ldots, n_{sim}$:

1. Draw $\alpha_T^{(s)}$ from the full conditional posterior

$$\alpha_T^{(s)} | y_{1:T}, \Sigma^{(s-1)}, \Upsilon^{(s-1)} \sim \mathcal{N}(\alpha_{T|T}, \Upsilon_{T|T})$$

   obtained from the Kalman filter. For $t = T-1, \ldots, 1$ draw $\alpha_t^{(s)}$ from the full conditional posterior

$$\alpha_t^{(s)} | y_{1:T}, \Sigma^{(s-1)}, \Upsilon^{(s-1)} \sim \mathcal{N}(\alpha_{t|T}, \Upsilon_{t|T})$$

   obtained from a simulation smoother.

2. Draw $\Upsilon^{(s)}$ from

$$\Upsilon^{(s)} | \alpha_{1:T}^{(s)} \sim \mathcal{IW}\left(\underline{S}_\Upsilon + \sum_{t=1}^{T} \left[\alpha_{t+1}^{(s)} - \alpha_t^{(s)}\right] \left[\alpha_{t+1}^{(s)} - \alpha_t^{(s)}\right]', \underline{d}_\Upsilon + T\right).$$

3. Draw $\Sigma^{(s)}$ from

$$\Sigma^{(s)} | \mathbf{y}, \alpha_{1:T}^{(s)} \sim \mathcal{IW}\left(\underline{S} + \sum_{t=1}^{T} \left[\mathbf{y} - (\mathbb{I}_n \otimes x)\alpha_t^{(s)}\right] \left[\mathbf{y} - (\mathbb{I}_n \otimes x)\alpha_t^{(s)}\right]', \underline{d} + T\right).$$

When stochastic volatility is added to the framework, the VAR innovations are assumed to be still normally distributed, but with variance that evolves over time (see Cogley and Sargent, 2002, 2005; Primiceri, 2005)

$$u_t \sim \mathcal{N}(0, \Sigma_t) \ , \qquad \Sigma_t = K^{-1} \Xi_t (K^{-1})' \ , \tag{53}$$

where K is a lower-triangular matrix with ones on the main diagonal, and Ξ_t a diagonal

matrix with elements evolving following a geometric random-walk process

$$ln(\Xi_t)_j = ln(\Xi_{t-1})_j + \eta_{j,t} \qquad \eta_{j,t} \sim i.i.d.\ \mathcal{N}(0, \sigma^2_{\eta,j}) \ . \tag{54}$$

The prior distributions for Υ and $\sigma^2_{\eta,j}$ $j = 1, \ldots, n$ can be used to express beliefs about the magnitude of the period-to-period drift in the VAR coefficients, and the changes in the volatility of the VAR innovations respectively. In practice, these priors are chosen to ensure that innovations to the parameters are small enough that the short- and medium-run dynamics of y_t are not swamped by the random-walk behaviour of A_t and Ξ_t. Primiceri (2005) extends the above TVP-VAR by also allowing the nonzero off-diagonal elements of the contemporaneous covariance matrix K to evolve as random-walk processes (i.e. K is replaced by K_t to allow for an arbitrary time-varying correlation structure). A Gibbs sampler to draw from the posterior distribution of the parameters is in Primiceri (2005).

5.2 Markov Switching, Threshold, and Smooth Transition VARs

Contrary to the drifting coefficients models discussed in the previous section, Markov switching (MS) VARs are designed to capture abrupt changes in the dynamics of y_t.[27] These can be viewed as models that allow for at least one structural break to occur within the sample, with the timing of the break being unknown. They are of the form

$$y_t = A(s_t)x_t + u_t, \qquad u_t \sim \mathcal{N}(0, \Sigma(s_t)), \tag{55}$$

where x_t is defined as in Eq. (5). The matrix of autoregressive coefficients $A(s_t)$ and the variance of the error term $\Sigma(s_t)$ are a function of a discrete m-state Markov process s_t with fixed transition probabilities

$$\pi_{ij} \equiv p(s_t = \mathcal{S}_i | s_t = \mathcal{S}_j) \text{ for } i, j \in [1, \ldots, m]. \tag{56}$$

[27] The book by Kim and Nelson (1999) is the standard reference for frequentist and Bayesian estimation of Markov switching models.

If $\pi_{ii} = 1$ for some $i \in [1, \ldots, m]$, then $\mathcal{S}_i$ is an absorbing state from which the system is not allowed to move away. Suppose $m = 2$, and that both $A(s_t)$ and $\Sigma(s_t)$ change simultaneously when switching from $\mathcal{S}_1$ to $\mathcal{S}_2$ and vice versa. If a NIW prior is specified for $A(s_t)$ and $\Sigma(s_t)$, and π_{11} and π_{22} have independent Beta prior distributions, a Gibbs sampler can be used to sample from the posterior (see e.g. Del Negro and Schorfheide, 2011).

A MS-VAR with non-recurrent states is called a 'change-point' model (see Chib, 1998; Bauwens and Rombouts, 2012). Generalising the specification to allow for more states, with the appropriate transition probabilities, allows to adapt the change-point model to the case of several structural breaks (see also Koop and Potter, 2007, 2009; Liu et al., 2017, for models where the number of change-points is unknown). Important extensions regard the transmission of structural shocks in the presence of structural breaks and in a time-varying coefficient environment discussed in e.g. Sims and Zha (2006) and Koop et al. (2011) who also allow for cointegration.

In threshold VARs (TVARs), the coefficients of the model change across regimes when an observable variable exceeds a given threshold value. Bayesian inference in TVAR models is discussed in detail in Geweke and Terui (1993) and Chen and Lee (1995). A TVAR with two regimes can be written as

$$y_t = Ax_t + \Theta(\tau_{t-d} - \tau)A^* x_t + u_t, \tag{57}$$

where A and A^* are $n \times k$ matrices that collect the autoregressive coefficients of the two regimes, $\Theta(\cdot)$ is a Heaviside step function, i.e. a discontinuous function whose value is zero for a negative argument, and one for a positive argument, τ_{t-d} is threshold variable at lag d, and τ is a potentially unobserved threshold value. The system in Eq. (57) can be easily generalised to allow for multiple regimes. TVARs have been applied to several problems in the economic literature (see, for example Koop and Potter, 1999; Ricco et al., 2016; Alessandri and Mumtaz, 2017).

If the coefficients gradually migrate to the new state(s), the model is called a smooth-

transition VAR (STVAR). A STVAR model with two regimes can be written as

$$y_t = (1 - G(w_t; \vartheta, w))Ax_t + G(w_t; \vartheta, w)A^* x_t + u_t, \tag{58}$$

where A^*, A, and x_t are defined as in Eq. (57). The function $G(w_t; \vartheta, w)$ governs the transition across states, and is a function of the observable variable w_t, and of the parameters ϑ and w. In an exponential smooth-transition (EST) VAR, typically

$$G(w_t; \vartheta, w) = \frac{1}{1 + \exp\{-\vartheta(w_t - w)/\sigma_w\}} \tag{59}$$

where $\vartheta > 0$ determines the speed of transition across regimes, w can be thought of as a threshold value, and σ_w is the sample standard deviation of w_t. The higher ϑ the more abrupt the transition, the more the model collapses into a fixed threshold VAR. Among others, Gefang and Strachan (2009) and Gefang (2012) apply Bayesian techniques to estimate Smooth-transition VAR models.

6 Bayesian Panel VARs

Panel VARs generalise VAR models by describing the joint dynamics of multiple time series of heterogenous and interacting units – as for examples, the economies of several countries, regions, or sectors. Thorough reviews are in Canova and Ciccarelli (2013) and in Dieppe et al. (2016).

A panel VAR describes the the evolution of $y_{t,i}$ – the vector of $n \times 1$ endogenous variables of each unit $i \in [1, \dots N]$ – by a system of p-th order VARs

$$y_{t,i} = \sum_{j=1}^{N} [A_{1,ij} y_{t-1,j} + \dots + A_{p,ij} y_{t-p,j} + c_j] + H_i w_t + u_{t,i} \ , \tag{60}$$

where w_t is a vector of m exogenous controls. The innovations are generally assumed to be *i.i.d.* and Gaussian

$$u_{t,i} \sim i.i.d.\ \mathcal{N}(0, \sigma_i^2) \ , \tag{61}$$

while being possibly correlated across units.

Stacking over the N units, the model assumes the form of a VAR(p) with exogenous controls

$$y_t = A_1 y_{t-1} + \ldots + A_p y_{t-p} + c + H w_t + u_t \ , \qquad u_t \sim i.i.d.\ \mathcal{N}(0, \Sigma) \ . \tag{62}$$

In Eq. (63)

$$\underset{[Nn\times 1]}{y_t} = \begin{pmatrix} y_{t,1} \\ \vdots \\ y_{t,N} \end{pmatrix}, \quad \underset{[Nn\times 1]}{u_t} = \begin{pmatrix} u_{t,1} \\ \vdots \\ u_{t,N} \end{pmatrix} \ , \tag{63}$$

moreover

$$\underset{[Nn\times Nn]}{A_\ell} = \begin{pmatrix} A_{\ell,11} & \ldots & A_{\ell,1N} \\ \vdots & \ddots & \vdots \\ A_{\ell,1N} & \ldots & A_{\ell,NN} \end{pmatrix} \ \ell = 1, \ldots, p \ ,$$

and

$$\underset{[Nn\times Nn]}{\Sigma} = \begin{pmatrix} \Sigma_{11} & \ldots & \Sigma_{1N} \\ \vdots & \ddots & \vdots \\ \Sigma_{1N} & \ldots & \Sigma_{NN} \end{pmatrix} \qquad \underset{[Nn\times m]}{H} = \begin{pmatrix} H_1 \\ \vdots \\ H_N \end{pmatrix} \ .$$

While in Eq. (63) the system appears as a standard VAR, its panel structure is captured by three properties: (i) Dynamic interdependencies – the dynamics of the variables in each unit depend on the lagged values of the other endogenous variables in the unit and possibly all other units, i.e. $A_{\ell,jk} \neq 0$ for $j \neq k$; (ii) Static interdependencies – the innovations $u_{t,i}$ can be correlated across units, i.e. $\Sigma_{ij} \neq 0$ for $i \neq j$; (iii) Cross-unit (sub-sectional) heterogeneity – the VAR coefficients and residual variances can be unit-specific, i.e. $A_{\ell,ik} \neq A_{\ell,jk}$, $H_i \neq H_j$ and $\Sigma_{ii} \neq \Sigma_{jj}$ for $i \neq j$.[28]

If all of these properties are present in the data, and relationships do not exist among the coefficients, the system is a VAR with a large cross-section (i.e. a Large VAR) and can be estimated with standard macroeconomic priors such as e.g. the Minnesota priors

[28] An additional potential property of the Panel VAR is the time-variation in the VAR coefficients. For ease of notation we abstract from this in the following exposition.

of Section 3.[29] However, it is often possible to assume that some of these properties are relevant to system of interest.

If the units do not have dynamic or static interdependencies (i.e. $A_{\ell,jk} = 0$ and $\Sigma_{jk} = 0$ for $j \neq k$) and the dynamic coefficients are homogenous across units (i.e. $A_{\ell,jj} = \bar{A}_\ell$, $H_j = \bar{H}$, and $\Sigma_{jj} = \bar{\Sigma}\ \forall j$), then

$$A_\ell = \mathbb{I}_N \otimes \bar{A}_\ell, \qquad H = 1_{N\times 1} \otimes \bar{H}, \qquad \Sigma = \mathbb{I}_N \otimes \bar{\Sigma} \tag{64}$$

and the system simplifies into a single pooled VAR, with only $n \times (np + m)$ coefficients to be estimated.[30] By stacking the observations first over different units and then over different times, the system can be cast in the standard SUR representation (Eq. 6) and estimated with standard priors (e.g. Normal-Inverse Wishart prirors) and techniques.

If the dynamic coefficients are heterogenous across units but no dynamic or static interdependencies are present, then the system breaks up into N independent VARs, with the following SUR representation

$$y_i = x_i A_i + u_i\ . \tag{65}$$

A random coefficients model for Eq. (65) assumes that the coefficients from each unit can be thought of as random draws from a common distribution. For example,

$$\alpha_i \sim \mathcal{N}(a, \Sigma_a)\ , \tag{66}$$

where $\alpha_i \equiv vec(A_i)$. Eq. (66) can be thought of as an exchangeable Bayesian prior on the

[29] A Panel VAR can always be estimated as a Large VAR using standard macroeconomic priors (Bańbura et al., 2010). However, this implies (comes to the cost of) treating all the variables symmetrically thus disregarding the unit structure, and the fact that different variables may measure the same quantities in different units. Also, for very large systems the need to adopt too tight priors to overcome the issue of dimensionality may distort the posterior distribution.

[30] Alternatively, one could use standard priors to estimate a VAR for each of the N units separately, and then average the results across units. Such a mean group estimator is inefficient relative to the pooled estimator under dynamic homogeneity, but gives consistent estimates of the average system dynamic effects if dynamic heterogeneity is present. Conversely, the pooled estimator is inconsistent under dynamic heterogeneity due to the presence of correlation between the regressors and the error term.

units' coefficients – viz. the unit indices i are uninformative, in the sense that they can be exchanged without any loss of information. This approach was proposed by Zellner and Hong (1989) who used a 'Minnesota' type prior with fixed and known residual covariance matrix, a diagonal Σ_a with overall tightness hyperparameter λ_a, and a plug-in pooled estimator for $\boldsymbol{a}$. Jarociñski (2010) proposes instead a fully Bayesian model in which all the parameters are treated as random variables, and a sophisticated hierarchical prior approach is adopted. Estimation can be archived using a Gibbs sampler.

If dynamic interdependencies are allowed, the estimation problem becomes more complex. Canova and Ciccarelli (2004, 2009) have suggested solutions based on different Bayesian and cross-sectional shrinkage techniques that can deal with the issue of parameters proliferation that arises in these cases. The approach works by assuming a factor structure for the the matrix of coefficients and can be estimated with standard Bayesian priors and a Gibbs sampler. This structural factor approach is very flexible, and can also be used to estimate Panel VARs with dynamic coefficients that evolve over time, as done in e.g. Ciccarelli et al. (2012) and Canova and Ciccarelli, 2013.

References

Alessandri, Piergiorgio and Haroon Mumtaz (2017) "Financial conditions and density forecasts for US output and inflation," *Review of Economic Dynamics*, Vol. 24, pp. 66–78, March.

Bańbura, Marta, Domenico Giannone, and Lucrezia Reichlin (2010) "Large Bayesian vector auto regressions," *Journal of Applied Econometrics*, Vol. 25, No. 1, pp. 71–92.

Bauwens, Luc and Jeroen V.K. Rombouts (2012) "On marginal likelihood computation in change-point models," *Computational Statistics & Data Analysis*, Vol. 56, No. 11, pp. 3415–3429.

Belmonte, Miguel A.G., Gary Koop, and Dimitris Korobilis (2014) "Hierarchical Shrinkage in Time-Varying Parameter Models," *Journal of Forecasting*, Vol. 33, No. 1, pp. 80–94, January.

Berger, James O. (1985) *Statistical Decision Theory and Bayesian Analysis*: Springer.

Bernardo, J.M. and A.F.M. Smith (2009) *Bayesian Theory*, Wiley Series in Probability and Statistics: Wiley.

Bhattacharya, Anirban, Debdeep Pati, Natesh S. Pillai, and David B. Dunson (2015) "Dirichlet-Laplace Priors for Optimal Shrinkage," *Journal of the American Statistical Association*, Vol. 110, No. 512, pp. 1479–1490. PMID: 27019543.

Bloor, Chris and Troy Matheson (2011) "Real-time conditional forecasts with Bayesian VARs: An application to New Zealand," *The North American Journal of Economics and Finance*, Vol. 22, No. 1, pp. 26 – 42. Symposium on Nowcasting and Model Combination.

Canova, Fabio (1992) "An Alternative Approach to Modeling and Forecasting Seasonal Time Series," *Journal of Business & Economic Statistics*, Vol. 10, No. 1, pp. 97–108.

——— (1993) "Forecasting time series with common seasonal patterns," *Journal of Econometrics*, Vol. 55, No. 1-2, pp. 173–200.

——— (2007) *Methods for Applied Macroeconomic Research*: Princeton University Press.

Canova, Fabio and Matteo Ciccarelli (2004) "Forecasting and turning point predictions in a Bayesian panel VAR model," *Journal of Econometrics*, Vol. 120, No. 2, pp. 327–359, June.

——— (2009) "Estimating Multicountry VAR Models," *International Economic Review*, Vol. 50, No. 3, pp. 929–959.

——— (2013) "Panel Vector Autoregressive Models: A Survey," in Thomas B. Fomby, Lutz Kilian, and Anthony Murphy eds. *VAR Models in Macroeconomics – New Developments and Applications: Essays in Honor of Christopher A. Sims (Advances in Econometrics)*, Vol. 32: Emerald Group Publishing Limited, pp. 205 – 246.

Carriero, Andrea, Todd E. Clark, and Massimiliano Marcellino (2015) "Bayesian VARs: Specification Choices and Forecast Accuracy," *Journal of Applied Econometrics*, Vol. 30, No. 1, pp. 46–73.

Carriero, Andrea, George Kapetanios, and Massimiliano Marcellino (2009) "Forecasting exchange rates with a large Bayesian VAR," *International Journal of Forecasting*, Vol. 25, No. 2, pp. 400–417.

Chen, Cathy W. S. and Jack C. Lee (1995) "Bayesian Inference of Threshold Autoregressive Models," *Journal of Time Series Analysis*, Vol. 16, No. 5, pp. 483–492.

Chib, Siddhartha (1998) "Estimation and comparison of multiple change-point models," *Journal of Econometrics*, Vol. 86, No. 2, pp. 221 – 241.

——— (2001) "Markov Chain Monte Carlo Methods: Computation and Inference," in James J. Heckman and Edward Leamer eds. *Handbook of Econometrics*, Vol. 5: Elsevier, pp. 3569 – 3649.

Chib, Siddhartha and Edward Greenberg (1995) "Understanding the Metropolis-Hastings Algorithm," *The American Statistician*, Vol. 49, No. 4, pp. 327–335.

Chiu, Ching-Wai (Jeremy), Haroon Mumtaz, and Gábor Pintér (2017) "Forecasting with VAR models: Fat tails and stochastic volatility," *International Journal of Forecasting*, Vol. 33, No. 4, pp. 1124 – 1143.

Ciccarelli, Matteo and Alessandro Rebucci (2003) "Bayesian Vars; A Survey of the Recent Literature with An Application to the European Monetary System," IMF Working Papers 03/102, International Monetary Fund.

Ciccarelli, Matteo, Eva Ortega, and Maria Teresa Valderrama (2012) "Heterogeneity and cross-country spillovers in macroeconomic-financial linkages," Working Paper Series 1498, European Central Bank.

Cogley, Timothy and Thomas J. Sargent (2002) "Evolving Post-World War II U.S. Inflation Dynamics," in *NBER Macroeconomics Annual 2001, Volume 16*: National Bureau of Economic Research, Inc, pp. 331–388.

——— (2005) "Drift and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S," *Review of Economic Dynamics*, Vol. 8, No. 2, pp. 262–302, April.

De Mol, Christine, Domenico Giannone, and Lucrezia Reichlin (2008) "Forecasting using a large number of predictors: Is Bayesian shrinkage a valid alternative to principal components?" *Journal of Econometrics*, Vol. 146, No. 2, pp. 318 – 328. Honoring the research contributions of Charles R. Nelson.

DeJong, David N., Beth Ingram, and Charles H. Whiteman (1993) "Analyzing VARs with monetary business cycle model priors," *Proceedings of the American Statistical Association, Bayesian Statistics Section*, pp. 160–169.

Del Negro, Marco and Frank Schorfheide (2004) "Priors from General Equilibrium Models for VARS," *International Economic Review*, Vol. 45, No. 2, pp. 643–673, May.

——— (2011) "Bayesian Macroeconometrics," in John Geweke, Gary Koop, and Herman Van Dijk eds. *The Oxford Handbook of Bayesian Econometrics*: Oxford University Press, pp. 293–389.

Dieppe, Alistair, Björn van Roye, and Romain Legrand (2016) "The BEAR toolbox," Working Paper Series 1934, European Central Bank.

Doan, Thomas, Robert Litterman, and Christopher Sims (1984) "Forecasting and conditional projection using realistic prior distributions," *Econometric Reviews*, Vol. 3, No. 1, pp. 1–100.

Gefang, Deborah (2012) "Money-output Causality Revisited – A Bayesian Logistic Smooth Transition VECM Perspective," *Oxford Bulletin of Economics and Statistics*, Vol. 74, No. 1, pp. 131–151, February.

Gefang, Deborah and Rodney Strachan (2009) "Nonlinear Impacts of International Business Cycles on the U.K. – A Bayesian Smooth Transition VAR Approach," *Studies in Nonlinear Dynamics & Econometrics*, Vol. 14, No. 1, pp. 1–33, December.

Geisser, Seymour (1965) "Bayesian Estimation in Multivariate Analysis," *The Annals of Mathematical Statistics*, Vol. 36, No. 1, pp. 150–159, 02.

Gelman, A., J.B. Carlin, H.S. Stern, and D.B. Rubin (2003) *Bayesian Data Analysis, Second Edition*, Chapman & Hall/CRC Texts in Statistical Science: Taylor & Francis.

Gelman, A., J.B. Carlin, H.S. Stern, D.B. Dunson, A. Vehtari, and D.B. Rubin (2013) *Bayesian Data Analysis, Third Edition*, Chapman & Hall/CRC Texts in Statistical Science: Taylor & Francis.

George, Edward I., Dongchu Sun, and Shawn Ni (2008) "Bayesian stochastic search for VAR model restrictions," *Journal of Econometrics*, Vol. 142, No. 1, pp. 553 – 580.

Geweke, John (1996) "Bayesian reduced rank regression in econometrics," *Journal of Econometrics*, Vol. 75, No. 1, pp. 121–146, November.

——— (1999) "Using simulation methods for bayesian econometric models: inference, development,and communication," *Econometric Reviews*, Vol. 18, No. 1, pp. 1–73.

——— (2005) *Contemporary Bayesian Econometrics and Statistics*, Wiley Series in Probability and Statistics: Wiley.

Geweke, John and Nobuhiko Terui (1993) "Bayesian Threshold Autoregressive Models for Nonlinear Time Series," *Journal of Time Series Analysis*, Vol. 14, No. 5, pp. 441–454.

Geweke, John and Charles Whiteman (2006a) "Bayesian Forecasting," in G. Elliott, C. Granger, and A. Timmermann eds. *Handbook of Economic Forecasting*, Vol. 1: Elsevier, 1st edition, Chap. 01, pp. 3–80.

——— (2006b) "Bayesian Forecasting," in G. Elliott, C.W.J. Granger, and A. Timmermann eds. *Handbook of Economic Forecasting*, Vol. 1: Elsevier, pp. 3 – 80.

Giannone, Domenico, Michele Lenza, and Giorgio E. Primiceri (2015) "Prior Selection for Vector Autoregressions," *The Review of Economics and Statistics*, Vol. 2, No. 97, pp. 436–451, May.

——— (2016) "Priors for the Long Run," CEPR Discussion Papers 11261, C.E.P.R. Discussion Papers.

——— (2017) "Economic Predictions with Big Data: The Illusion Of Sparsity," CEPR Discussion Papers 12256, C.E.P.R. Discussion Papers.

Giannone, Domenico, Michele Lenza, and Lucrezia Reichlin (2008) "Explaining The Great Moderation: It Is Not The Shocks," *Journal of the European Economic Association*, Vol. 6, No. 2-3, pp. 621–633, 04-05.

Giannone, Domenico, Michele Lenza, Daphne Momferatou, and Luca Onorante (2014) "Short-term inflation projections: A Bayesian vector autoregressive approach," *International Journal of Forecasting*, Vol. 30, No. 3, pp. 635–644.

Griffin, Jim E. and Philip J. Brown (2010) "Inference with normal-gamma prior distributions in regression problems," *Bayesian Anal.*, Vol. 5, No. 1, pp. 171–188, 03.

Griffin, Jim and Phil Brown (2017) "Hierarchical Shrinkage Priors for Regression Models," *Bayesian Analysis*, Vol. 12, No. 1, pp. 135–159, 03.

Hajivassiliou, Vassilis A. and Paul A. Ruud (1994) "Classical estimation methods for LDV models using simulation," in *Handbook of Econometrics*, Vol. 4: Elsevier, pp. 2383 – 2441.

Highfield, Richard A. (1992) "Forecasting Similar Time Series with Bayesian Pooling Methods: Application to Forecasting European Output Growth," in Prem K. Goel and N. Sreenivas Iyengar eds. *Bayesian Analysis in Statistics and Econometrics*, pp. 303–326, New York, NY: Springer New York.

Huber, Florian and Martin Feldkircher (2017) "Adaptive Shrinkage in Bayesian Vector Autoregressive Models," *Journal of Business & Economic Statistics*, Vol. 0, No. 0, pp. 1–13.

Ingram, Beth F. and Charles H. Whiteman (1994) "Supplanting the 'Minnesota' prior: Forecasting macroeconomic time series using real business cycle model priors," *Journal of Monetary Economics*, Vol. 34, No. 3, pp. 497–510, December.

Jarociñski, Marek (2010) "Responses to monetary policy shocks in the east and the west of Europe: a comparison," *Journal of Applied Econometrics*, Vol. 25, No. 5, pp. 833–868.

Jarociński, Marek and Bartosz Maćkowiak (2017) "Granger Causal Priority and Choice of Variables in Vector Autoregressions," *The Review of Economics and Statistics*, Vol. 99, No. 2, pp. 319–329, May.

Jarociński, Marek and Albert Marcet (2011) "Autoregressions in Small Samples, Priors about Observables and Initial Conditions," CEP Discussion Papers dp1061, Centre for Economic Performance, LSE.

——— (2014) "Contrasting Bayesian and Frequentist Approaches to Autoregressions: the Role of the Initial Condition," Working Papers 776, Barcelona Graduate School of Economics.

Jochmann, Markus and Gary Koop (2015) "Regime-switching cointegration," *Studies in Nonlinear Dynamics & Econometrics*, Vol. 19, No. 1, pp. 35–48, February.

Kadiyala, Rao K. and Sune Karlsson (1997) "Numerical Methods for Estimation and Inference in Bayesian VAR-Models," *Journal of Applied Econometrics*, Vol. 12, No. 2, pp. 99–132, March-Apr.

Karlsson, Sune (2013a) "Forecasting with Bayesian Vector Autoregression," in Graham Elliott and Allan Timmermann eds. *Handbook of Economic Forecasting*, Vol. 2 of Handbook of Economic Forecasting: Elsevier, pp. 791 – 897.

——— (2013b) *Forecasting with Bayesian Vector Autoregression*, Vol. 2 of Handbook of Economic Forecasting, Chap. 0, pp. 791–897: Elsevier.

Kim, Jae Young (1994) "Bayesian Asymptotic Theory in a Time Series Model with a Possible Nonstationary Process," *Econometric Theory*, Vol. 10, No. 3/4, pp. 764–773.

Kim, Chang Jin and Charles Nelson (1999) *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications*, Vol. 1: The MIT Press, 1st edition.

King, Robert G., Charles I. Plosser, and Sergio T. Rebelo (1988) "Production, growth and business cycles: I. The basic neoclassical model," *Journal of Monetary Economics*, Vol. 21, No. 2, pp. 195 – 232.

Kleibergen, Frank and Richard Paap (2002) "Priors, posteriors and bayes factors for a Bayesian analysis of cointegration," *Journal of Econometrics*, Vol. 111, No. 2, pp. 223–249, December.

Kleibergen, Frank and Herman K. van Dijk (1994) "On the Shape of the Likelihood/Posterior in Cointegration Models," *Econometric Theory*, Vol. 10, No. 3-4, pp. 514–551, August.

Koop, Gary (2003) *Bayesian Econometrics*: Wiley.

——— (2013) "Forecasting with Medium and Large Bayesian VARS," *Journal of Applied Econometrics*, Vol. 28, No. 2, pp. 177–203, March.

Koop, Gary and Dimitris Korobilis (2010) "Bayesian Multivariate Time Series Methods for Empirical Macroeconomics," *Foundations and Trends(R) in Econometrics*, Vol. 3, No. 4, pp. 267–358, July.

Koop, Gary and Simon M. Potter (1999) "Dynamic Asymmetries in U.S. Unemployment," *Journal of Business & Economic Statistics*, Vol. 17, No. 3, pp. 298–312.

——— (2007) "Estimation and Forecasting in Models with Multiple Breaks," *The Review of Economic Studies*, Vol. 74, No. 3, pp. 763–789.

——— (2009) "Prior Elicitation In Multiple Change-Point Models," *International Economic Review*, Vol. 50, No. 3, pp. 751–772, August.

Koop, Gary and Mark F. J. Steel (1991) "A comment on: "To criticize the critics: An objective bayesian analysis of stochastic trends", By Peter C. B. Phillips," *Journal of Applied Econometrics*, Vol. 6, No. 4, pp. 365–370.

Koop, Gary, Roberto Leon-Gonzalez, and Rodney W. Strachan (2011) "Bayesian inference in a time varying cointegration model," *Journal of Econometrics*, Vol. 165, No. 2, pp. 210–220.

Koop, G., R.W. Strachan, H.K. van Dijk, and M. Villani (2006) "Monetary policy shocks: What have we learned and to what end?" in T. C.Mills and K. P. Patterson eds. *Palgrave Handbook of Econometrics*, Vol. 1: Basingstoke: Palgrave Macmillan, pp. 871–98.

Korobilis, Dimitris (2013) "VAR Forecasting Using Bayesian Variable Selection," *Journal of Applied Econometrics*, Vol. 28, No. 2, pp. 204–230.

Kroese, Dirk P. and Joshua Chan (2014) *Statistical Modeling and Computation*: Springer.

Kwan, Yum K. (1998) "Asymptotic Bayesian analysis based on a limited information estimator," *Journal of Econometrics*, Vol. 88, No. 1, pp. 99–121, November.

Litterman, Robert B. (1979) "Techniques of forecasting using vector autoregressions," Working Papers 115, Federal Reserve Bank of Minneapolis.

——— (1980) "A Bayesian Procedure for Forecasting with Vector Autoregression," working papers, MIT Department of Economics.

——— (1986) "Forecasting with Bayesian Vector Autoregressions-Five Years of Experience," *Journal of Business & Economic Statistics*, Vol. 4, No. 1, pp. 25–38, January.

Liu, Philip, Haroon Mumtaz, Konstantinos Theodoridis, and Francesco Zanetti (2017) "Changing Macroeconomic Dynamics at the Zero Lower Bound," *Journal of Business & Economic Statistics*, Vol. 0, No. ja, pp. 0–0.

Lubik, Thomas A. and Christian Matthes (2015) "Time-Varying Parameter Vector Autoregressions: Specification, Estimation, and an Application," *Economic Quarterly*, Vol. 4Q, pp. 323–352.

Miranda-Agrippino, Silvia and Giovanni Ricco (2017) "The transmission of monetary policy shocks," Bank of England working papers 657, Bank of England.

——— (2018) "Bayesian Vector Autoregressions: Applications," in Jonathan Hamilton ed. *Oxford Research Encyclopedia in Economics and Finance*: Oxford University Press.

Mitchell, T. J. and J. J. Beauchamp (1988) "Bayesian Variable Selection in Linear Regression," *Journal of the American Statistical Association*, Vol. 83, No. 404, pp. 1023–1032.

Müller, Ulrich K. (2013) "Risk of Bayesian Inference in Misspecified Models, and the Sandwich Covariance Matrix," *Econometrica*, Vol. 81, No. 5, pp. 1805–1849.

Müller, Ulrich K. and Graham Elliott (2003) "Tests for Unit Roots and the Initial Condition," *Econometrica*, Vol. 71, No. 4, pp. 1269–1286.

Panagiotelis, Anastasios and Michael Smith (2008) "Bayesian density forecasting of intraday electricity prices using multivariate skew t distributions," *International Journal of Forecasting*, Vol. 24, No. 4, pp. 710–727.

Phillips, P. C. B. (1991a) "Bayesian routes and unit roots: De rebus prioribus semper est disputandum," *Journal of Applied Econometrics*, Vol. 6, No. 4, pp. 435–473.

——— (1991b) "To criticize the critics: An objective bayesian analysis of stochastic trends," *Journal of Applied Econometrics*, Vol. 6, No. 4, pp. 333–364.

Primiceri, Giorgio E. (2005) "Time Varying Structural Vector Autoregressions and Monetary Policy," *Review of Economic Studies*, Vol. 72, No. 3, pp. 821–852.

Raiffa, H. and R. Schlaifer (1961) *Applied statistical decision theory*, Studies in managerial economics: Division of Research, Graduate School of Business Adminitration, Harvard University.

Ricco, Giovanni, Giovanni Callegari, and Jacopo Cimadomo (2016) "Signals from the government: Policy disagreement and the transmission of fiscal shocks," *Journal of Monetary Economics*, Vol. 82, No. C, pp. 107–118.

Robertson, John C. and Ellis W. Tallman (1999) "Vector autoregressions: forecasting and reality," *Economic Review*, Vol. Issue Q1, pp. 4–18.

Sims, Christopher A. (1980) "Macroeconomics and Reality," *Econometrica*, Vol. 48, No. 1, pp. 1–48, January.

——— (1988) "Bayesian skepticism on unit root econometrics," *Journal of Economic Dynamics and Control*, Vol. 12, No. 2, pp. 463 – 474.

——— (1991) "Comment by Christopher A. Sims on 'To criticize the critics', by Peter C. B. Phillips," *Journal of Applied Econometrics*, Vol. 6, No. 4, pp. 423–434.

——— (1993) "A Nine-Variable Probabilistic Macroeconomic Forecasting Model," in *Business Cycles, Indicators and Forecasting*: National Bureau of Economic Research, Inc, pp. 179–212.

——— (1996) "Inference For Multivariate Time Series Models With Trend,"Technical report, Princeton University, mimeo.

——— (2000) "Using a likelihood perspective to sharpen econometric discourse: Three examples," *Journal of Econometrics*, Vol. 95, No. 2, pp. 443–462.

——— (2005a) "Conjugate Dummy Observation Priors for VARs,"Technical report, Princeton University, mimeo.

——— (2005b) "Dummy Observation Priors Revisited,"Technical report, Princeton University, mimeo.

——— (2010a) "Causal Ordering and Exogeneity,"Technical report, Princeton University, mimeo.

——— (2010b) "Understanding Non-Bayesians,"Technical report, Princeton University, mimeo.

Sims, Christopher A and Harald Uhlig (1991) "Understanding Unit Rooters: A Helicopter Tour," *Econometrica*, Vol. 59, No. 6, pp. 1591–1599, November.

Sims, Christopher A. and Tao Zha (1998) "Bayesian Methods for Dynamic Multivariate Models," *International Economic Review*, Vol. 39, No. 4, pp. 949–68, November.

——— (2006) "Were There Regime Switches in U.S. Monetary Policy?" *American Economic Review*, Vol. 96, No. 1, pp. 54–81, March.

Strachan, Rodney W. and Brett Inder (2004) "Bayesian analysis of the error correction model," *Journal of Econometrics*, Vol. 123, No. 2, pp. 307–325, December.

Theil, H. (1963) "On the Use of Incomplete Prior Information in Regression Analysis," *Journal of the American Statistical Association*, Vol. 58, No. 302, pp. 401–414.

Theil, H. and A. S. Goldberger (1961) "On Pure and Mixed Statistical Estimation in Economics," *International Economic Review*, Vol. 2, No. 1, pp. 65–78.

Tiao, George C. and Arnold Zellner (1964) "On the Bayesian Estimation of Multivariate Regression," *Journal of the Royal Statistical Society. Series B (Methodological)*, Vol. 26, No. 2, pp. 277–285.

Tierney, Luke (1994) "Markov Chains for Exploring Posterior Distributions," *Ann. Statist.*, Vol. 22, No. 4, pp. 1701–1728, 12.

Timmermann, Allan (2006) "Forecast Combinations," in G. Elliott, C. Granger, and A. Timmermann eds. *Handbook of Economic Forecasting*, Vol. 1: Elsevier, Chap. 4, pp. 135–196.

Uhlig, Harald (1994a) "On Jeffreys Prior when Using the Exact Likelihood Function," *Econometric Theory*, Vol. 10, No. 3-4, pp. 633–644.

——— (1994b) "What Macroeconomists Should Know about Unit Roots: A Bayesian Perspective," *Econometric Theory*, Vol. 10, No. 3-4, pp. 645–671.

——— (1997) "Bayesian Vector Autoregressions with Stochastic Volatility," *Econometrica*, Vol. 65, No. 1, pp. 59–74, January.

Villani, Mattias (2001) "Bayesian prediction with cointegrated vector autoregressions," *International Journal of Forecasting*, Vol. 17, No. 4, pp. 585 – 605.

——— (2009) "Steady-state priors for vector autoregressions," *Journal of Applied Econometrics*, Vol. 24, No. 4, pp. 630–650.

Wright, Jonathan H. (2009) "Forecasting US inflation by Bayesian model averaging," *Journal of Forecasting*, Vol. 28, No. 2, pp. 131–144.

Zellner, Arnold (1971) *An introduction to Bayesian inference in econometrics*, Wiley Classics Library: Wiley-Interscience.

Zellner, Arnold and Chansik Hong (1989) "Forecasting international growth rates using Bayesian shrinkage and other procedures," *Journal of Econometrics*, Vol. 40, No. 1, pp. 183–202, January.