

Bayesian Estimates for Vector Autoregressive Models

Shawn Ni

Department of Economics, University of Missouri, Columbia, MO 65211

Dongchu Sun

Department of Statistics, University of Missouri, Columbia, MO 65211

This article examines frequentist risks of Bayesian estimates of vector autoregressive (VAR) regression coefficient and error covariance matrices under competing loss functions, under various noninformative priors, and in the normal and Student- t models. Simulation results show that for the regression coefficient matrix, an asymmetric LINEX estimator does better overall than the posterior mean. No dominating estimator emerges for the error covariance matrix. We find that the choice of prior has a more significant effect on the estimates than the form of estimator. For the VAR regression coefficients, a shrinkage prior dominates a constant prior. For the error covariance matrix, Yang and Berger's reference prior dominates the Jeffreys prior. Estimation of a VAR using U.S. macroeconomic data yields significantly different estimates under competing priors.

KEY WORDS: Bayesian VAR; LINEX loss; Noninformative priors; Pseudoentropy loss; Quadratic loss; Student- t distribution.

1. INTRODUCTION

The purpose of this article is to explore properties of Bayesian estimates of vector autoregression (VAR) models under several possible choices of sampling distributions, loss functions, and priors. Over the past two decades, VAR models have become popular tools for analyzing time series data. Applications of VAR to such policy questions as the macroeconomic effects of monetary supply shocks are too numerous to list (see, e.g., Sims 1980, 1992; Christiano, Eichenbaum, and Evans 1999). The popularity of VAR models justifies a close look at the technical issues pertaining to VAR estimation and finite-sample inference. A typical VAR has hundreds of parameters, which are often estimated through least squares or maximum likelihood by VAR users. There are several problems with the frequentist approach to VAR estimation, however. First, for some distributions of data (such as Student- t), MLE does not have an analytical form or simply does not exist. Second, drawing finite-sample inference for the VAR parameters is a challenge. Frequentist finite-sample distributions cannot be derived in closed form even for normal errors, whereas asymptotic theory may not be applicable to a VAR with a large number of parameters and limited observations of macroeconomic data. Third, in some applications of VAR models, nonlinear functions of VAR parameters are the focus of research. For example, impulse responses are often considered easier to interpret than the regression coefficients. For another example, eigenvalues of the regression coefficient matrix are of interest, because they determine the long-term dynamics of the VAR. Finite-sample frequentist analysis of these nonlinear functions is more difficult than analysis of the VAR parameters themselves.

The difficulties encountered in the frequentist approach of VAR inference can be circumvented by the Bayesian approach, which combines information from data with the researcher's prior. Many applied econometricians find Bayesian methods appealing not only for philosophical reasons, but also, and more importantly, for their effectiveness for finite-sample inference.

Bayesian inference is influenced by the choice of prior. Elicitation of priors is an important step of Bayesian analysis. When

researchers have information on the nature of parameters of interest, they may use informative priors to reflect their beliefs. For example, Doan, Litterman, and Sims (1984) and Litterman (1986) observed that many macroeconomic time series approximately follow random-walk processes and developed an informative prior known as the "Minnesota prior" that reflects this pattern. In recent Bayesian studies, Pastor (2000) and Pastor and Stambaugh (2000) used finance theory for elicitation of informative priors. Hollifield, Koop, and Li (2001) developed a method for eliciting an informative prior for VAR variance decompositions.

In many applications of VAR models, because of the large number of parameters involved, researchers often find it desirable to use vague or noninformative priors. However, examination of the effects of using different noninformative priors on VAR posterior distributions is relatively rare. (For some studies along this line, see Kadiyala and Karlsson 1997; Ni and Sun 2003; Sun and Ni 2004.) The default prior in the literature for the VAR regression coefficients is the constant prior, and the default prior for the error covariance matrix is the Jeffreys prior or a modified version of it used in the RATS software package (called the RATS prior hereinafter). These combinations of priors allow for easy simulation of posterior distributions and are widely used in macroeconomic studies. The Jeffreys prior for the covariance matrix has also been widely used for multifactor asset pricing models in the finance literature (some recent examples include Kandel, McCulloch, and Stambaugh 1995; Kandel and Stambaugh 1996; Lamoureux and Zhou 1996; Pastor and Stambaugh 2000; Barberis 2000). However, the Jeffreys prior is known to be undesirable in high-dimensional settings (see Berger and Bernardo 1992). Bernardo (1979) and Berger and Bernardo (1992) proposed an alternative approach for deriving a reference prior. Following this approach, the reference prior for the covariance matrix in an iid

normal model was derived by Yang and Berger (1994). The constant prior on the regression coefficient matrix is known to be inadmissible under quadratic loss for estimation of an unknown mean of a vector with iid normal observations (see Berger and Strawderman 1996). An alternative to the constant prior is a “shrinkage” prior (e.g., Baranchik 1964; Berger and Strawderman 1996), which dominates the constant prior under squared error loss for estimating an unknown normal mean in iid cases but has scant applications to the VAR model.

For Bayesian VAR analysis, the subject of interest is the posterior of the VAR parameters. However, reporting the posterior distribution of each parameter is infeasible even for a small VAR. Some summaries of posterior quantities, such as the posterior risks of Bayesian estimates, would be desirable. Bayesian estimates are derived from minimization of expected posterior loss in the parameter space; however, the choice of loss function determines the form of the Bayesian estimator. In applications of Bayesian procedures, the posterior means of VAR regression coefficients and the error covariance matrix are usually reported as the Bayesian estimates. The posterior mean is optimal for certain loss functions. Bayesian estimates derived from minimizing other commonly used loss functions are rarely studied for VAR models. These loss functions include Yang and Berger’s (1994) quadratic and pseudoinverse entropy losses for the error covariance matrix and Zellner’s (1986) LINEX loss for the regression coefficient matrix. The fact that Bayesian estimates derived from these loss functions differ from the posterior mean may be of interest for macroeconomists. For instance, the posterior mean of the regression coefficient matrix is often biased. An asymmetric LINEX estimator may be helpful in correcting the bias.

The main objective of this article is to examine the effect of using competing priors and estimates for VAR models under various assumptions on the distribution of data. We explore Bayesian estimates under three competing loss functions on the covariance matrix (pseudoentropy loss, quadratic loss, and inverse pseudoentropy loss). We also compare the quadratic and asymmetric LINEX loss functions with respect to the VAR regression coefficient matrix. We consider each estimator under the following prior combinations: Jeffreys, RATS, or Yang and Berger’s reference prior on the covariance matrix, and a constant or shrinkage prior on the VAR coefficient matrix. The noninformative priors for the covariance matrix are derived under the normality assumption. We evaluate their performance for both normal and Student- t VAR models. Our study on the Student- t model extends Geweke’s (1993) analysis on univariate time series models to the vector case. We examine the competing VAR estimates in terms of frequentist average losses. We also investigate the posterior of the eigenvalues of the VAR regression coefficient matrix and the posterior of the impulse response functions.

We find that the asymmetric LINEX estimator for the VAR regression coefficient matrix does better overall than the posterior mean. The performance of competing estimates for the error covariance matrix is mixed. The choice of prior has more significant effects on the estimates than the form of the estimates. The shrinkage prior on the VAR regression coefficient matrix dominates the constant prior, whereas Yang and Berger’s reference prior on the covariance matrix dominates the Jeffreys

prior. These conclusions are drawn from a numerical example with the normal and Student- t errors. Estimation of a VAR using U.S. macroeconomic data produces significantly different estimates under the shrinkage and constant priors.

The rest of the article is organized as follows. In Section 2 we define notation for the VAR models with the normal or Student- t errors. In Section 3 we introduce several loss functions for VAR estimation and derive corresponding Bayesian estimates. In Section 4 we present priors, especially noninformative priors. In Section 5 we derive full conditional distributions under both normal and t -errors. Several algorithms under t -errors are given in the Appendix. In Section 6 we give simulation results for illustration, and also compare Bayesian estimates of a VAR using quarterly data of the U.S. economy. In Section 7 we offer concluding remarks.

2. VECTOR AUTOREGRESSIVE MODELS

A VAR of a p -dimensional column variable, \mathbf{y}_t , typically has the form

$$\mathbf{y}'_t = \mathbf{c} + \sum_{i=1}^L \mathbf{y}'_{t-i} \mathbf{B}_i + \boldsymbol{\epsilon}'_t \quad (1)$$

for $t = 1, \dots, T$, where L is a known positive integer, \mathbf{c} is a $1 \times p$ unknown vector, and \mathbf{B}_i is an unknown $p \times p$ matrix. In the normal VAR, the errors $\boldsymbol{\epsilon}_1, \dots, \boldsymbol{\epsilon}_T$ are iid $N_p(\mathbf{0}, \boldsymbol{\Sigma})$, and $\boldsymbol{\Sigma}$ is a $p \times p$ positive definite error covariance matrix. Define $\mathbf{x}'_t = (1, \mathbf{y}'_{t-1}, \dots, \mathbf{y}'_{t-L})$,

$$\mathbf{Y} = \begin{pmatrix} \mathbf{y}'_1 \\ \vdots \\ \mathbf{y}'_T \end{pmatrix}, \quad \mathbf{X} = \begin{pmatrix} \mathbf{x}'_1 \\ \vdots \\ \mathbf{x}'_T \end{pmatrix},$$

$$\boldsymbol{\Phi} = \begin{pmatrix} \mathbf{c} \\ \mathbf{B}_1 \\ \vdots \\ \mathbf{B}_L \end{pmatrix}, \quad \boldsymbol{\epsilon} = \begin{pmatrix} \boldsymbol{\epsilon}'_1 \\ \vdots \\ \boldsymbol{\epsilon}'_T \end{pmatrix},$$

where \mathbf{x}'_t is a $1 \times (1 + Lp)$ row vector and \mathbf{Y} and $\boldsymbol{\epsilon}$ are $T \times p$ matrices. Here \mathbf{X} is a $T \times (1 + Lp)$ matrix of observations, and the regression coefficient matrix $\boldsymbol{\Phi}$ is a $(1 + Lp) \times p$ matrix of unknown parameters. We can rewrite (1) as

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\Phi} + \boldsymbol{\epsilon}. \quad (2)$$

For the case of normal errors, the likelihood function of $(\boldsymbol{\Phi}, \boldsymbol{\Sigma})$ is

$$l_N(\boldsymbol{\Phi}, \boldsymbol{\Sigma}) = \frac{1}{|\boldsymbol{\Sigma}|^{T/2}} \exp \left\{ -\frac{1}{2} \sum_{i=1}^T (\mathbf{y}_i - \mathbf{x}_i \boldsymbol{\Phi})' \boldsymbol{\Sigma}^{-1} (\mathbf{y}_i - \mathbf{x}_i \boldsymbol{\Phi}) \right\}$$

$$= \frac{1}{|\boldsymbol{\Sigma}|^{T/2}} \text{etr} \left\{ -\frac{1}{2} (\mathbf{Y} - \mathbf{X}\boldsymbol{\Phi}) \boldsymbol{\Sigma}^{-1} (\mathbf{Y} - \mathbf{X}\boldsymbol{\Phi})' \right\}. \quad (3)$$

Here and in the following $\text{etr}(\mathbf{A})$ represents $\exp(\text{trace}(\mathbf{A}))$ for a matrix \mathbf{A} . Under the normality assumption, the maximum likelihood estimators (MLEs) are

$$\hat{\boldsymbol{\Phi}}_{MLE} = (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbf{Y} \quad \text{and} \quad \hat{\boldsymbol{\Sigma}}_{MLE} = \mathbf{S}(\hat{\boldsymbol{\Phi}}_{MLE})/T, \quad (4)$$

where

$$\mathbf{S}(\Phi) = (\mathbf{Y} - \mathbf{X}\Phi)'(\mathbf{Y} - \mathbf{X}\Phi). \quad (5)$$

It is well known that normal distributions have light tails. Alternative heavier-tailed distributions for the errors ϵ_t include an independent multivariate- t with parameters $(\mathbf{0}, \Sigma)$ and degrees of freedom ν , denoted by $t_\nu(\mathbf{0}, \Sigma)$. The density of $t_\nu(\mathbf{0}, \Sigma)$ is given by

$$\begin{aligned} p(\mathbf{s}|\Sigma, \nu) &= \frac{\Gamma(\frac{1}{2}(\nu + p))}{(\pi\nu)^{p/2}\Gamma(\frac{\nu}{2})} \\ &\times |\Sigma|^{-1/2} \left(1 + \frac{1}{\nu} \mathbf{s}'\Sigma^{-1}\mathbf{s}\right)^{-(\nu+p)/2}, \quad \mathbf{s} \in \mathbb{R}^p \quad (6) \end{aligned}$$

(see Johnson and Kotz 1972, p. 134).

Regression models with Student- t errors have been studied by a number of researchers in different contexts. In Zellner's (1976) model, the error terms (of different periods) in a univariate regression model form a joint multivariate- t distribution. In his model, given degrees of freedom, the MLEs of the variance and regression coefficients are the same as the MLEs in the normal model. For independent and identically t -distributed vector models, Fernandez, Osiewalski, and Steel (1997) proved propositions on the robustness of both frequentist and Bayesian inferences developed under the normality assumption. Osiewalski and Steel (1993) showed that in a regression model with exogenous regressors and an improper prior, the posteriors of the parameters of interest are the same under the normal and under t -errors. In Geweke's (1993) univariate time series regression model, the error terms are independently t -distributed and have unknown degrees of freedom. In our VAR model, the vector of error terms in each period is assumed to follow a multivariate- t distribution with unknown degrees of freedom. The case with degree of freedom $\nu = 1$ corresponds to a multivariate Cauchy distribution (cf. Johnson and Kotz 1972, p. 134). When $\nu \rightarrow \infty$, $t_\nu(\mathbf{0}, \Sigma)$ goes to $N_p(\mathbf{0}, \Sigma)$.

The property that the MLE based on the t -distributed errors is the same as the MLE under normal errors no longer holds for the VAR model. For given degrees of freedom parameter ν , the MLE of (Φ, Σ) under the t -distributed errors is not available in closed form. If ν is unknown, then the MLE for (Φ, Σ, ν) may not even exist. Fernandez and Steel (1999) discussed the difficulties of conducting likelihood-based inference for t -distributed models when identical data are recorded from independent sampling. But we do not encounter such a problem in our VAR analysis on a given set of time series observations. To simulate Bayesian estimates, in this article we develop MCMC methods based on a hierarchical structure of the multivariate t -distribution.

3. LOSS FUNCTIONS AND BAYESIAN ESTIMATES

A Bayesian estimator of (Φ, Σ) depends on the distribution model, the prior, and the loss function. We consider loss functions that contain a part measuring the loss associated with the error covariance matrix and a part measuring the loss of the

VAR regression coefficients only. The overall loss with respect to estimator $(\hat{\Phi}, \hat{\Sigma})$ for (Φ, Σ) has the form

$$L(\hat{\Sigma}, \hat{\Phi}; \Sigma, \Phi) = L_\Sigma(\hat{\Sigma}; \Sigma) + L_\Phi(\hat{\Phi}; \Phi). \quad (7)$$

The question that we seek to answer is whether alternative loss functions result in Bayesian estimates with significantly different properties.

3.1 Loss Functions for Σ

First, we consider the loss function for Σ ,

$$L_{\Sigma 1}(\hat{\Sigma}; \Sigma) = \text{tr}(\hat{\Sigma}^{-1}\Sigma) - \log|\hat{\Sigma}^{-1}\Sigma| - p, \quad (8)$$

where p is the number of variables in the VAR. We refer this function as a pseudoentropy loss because it is an entropy loss with respect to Σ only, whereas the true entropy loss pertains to both Σ and Φ .

The second loss function on Σ is a quadratic loss,

$$L_{\Sigma 2}(\hat{\Sigma}; \Sigma) = \text{tr}(\hat{\Sigma}\Sigma^{-1} - \mathbf{I})^2. \quad (9)$$

The third loss function is the pseudoentropy function on Σ^{-1} considered by Yang and Berger (1994),

$$L_{\Sigma 3}(\hat{\Sigma}; \Sigma) = \text{tr}(\hat{\Sigma}\Sigma^{-1}) - \log|\hat{\Sigma}\Sigma^{-1}| - p. \quad (10)$$

Note that the loss functions with respect to Σ are separated from the loss functions with respect to Φ . As a result, Bayesian estimates can be derived separately from minimizing expected posterior loss functions regarding (Φ, Σ) if the minimum is finite. A generalized Bayesian estimator can be defined analogously regardless of whether the Bayes risk is finite.

Lemma 1. (a) Under the loss $L_{\Sigma 1}$, the generalized Bayesian estimator of Σ is $\hat{\Sigma}_1 = E(\Sigma|\mathbf{Y})$.

(b) Under the loss $L_{\Sigma 2}$, the generalized Bayesian estimator of Σ is $\hat{\Sigma}_2$, given by $\text{vec}(\hat{\Sigma}_2) = [E\{(\Sigma^{-1} \otimes \Sigma^{-1})|\mathbf{Y}\}]^{-1} \times \text{vec}\{E(\Sigma^{-1}|\mathbf{Y})\}$, where \otimes represents the Kronecker product.

(c) Under the loss $L_{\Sigma 3}$, the generalized Bayesian estimator of Σ is $\hat{\Sigma}_3 = \{E(\Sigma^{-1}|\mathbf{Y})\}^{-1}$.

Proof. For part (a), note that

$$\frac{\partial \log(|\hat{\Sigma}|)}{\partial \hat{\Sigma}} = \hat{\Sigma}^{-1} \quad \text{and} \quad \frac{\partial \text{tr}(\hat{\Sigma}^{-1}\Sigma)}{\partial \hat{\Sigma}} = -\hat{\Sigma}^{-1}\Sigma\hat{\Sigma}^{-1}.$$

The foregoing facts can be used to calculate the derivative of $E\{L_{\Sigma 1}(\hat{\Sigma}, \Sigma)|\mathbf{Y}\}$ with respect to $\hat{\Sigma}$. The desired result follows by setting the derivative to 0. The proof of part (b) is more tedious (see Yang and Berger 1994). The proof of part (c) follows from the fact that the derivative of $E\{L_{\Sigma 3}(\hat{\Sigma}, \Sigma)|\mathbf{Y}\}$ with respect to $\hat{\Sigma}$ is $E(\Sigma^{-1}|\mathbf{Y}) - \hat{\Sigma}^{-1}$.

3.2 Loss Functions for Φ

The most common loss for Φ is the quadratic loss

$$L_{\Phi 1}(\widehat{\Phi}, \Phi) = \text{trace}\{(\widehat{\Phi} - \Phi)' \mathbf{W}(\widehat{\Phi} - \Phi)\}, \quad (11)$$

where \mathbf{W} is a constant weighting matrix. If the weighting matrix \mathbf{W} is the identity matrix, then the loss of $L_{\Phi 1}$ is simply the sum of squared errors of all elements of $\widehat{\Phi}$, $\sum_{i=1}^{1+Lp} \sum_{j=1}^p (\widehat{\phi}_{ij} - \phi_{ij})^2$.

The quadratic function is symmetric. LINEX loss, an asymmetric loss function, was explored by Zellner (1986) for estimating an iid normal mean under a conjugate prior. LINEX loss has the form

$$L_{\Phi 2}(\widehat{\Phi}, \Phi) = \sum_{i=1}^{1+Lp} \sum_{j=1}^p [\exp\{a_{ij}(\widehat{\phi}_{ij} - \phi_{ij})\} - a_{ij}(\widehat{\phi}_{ij} - \phi_{ij}) - 1], \quad (12)$$

where a_{ij} is a given constant. When a_{ij} is close to 0, the LINEX loss function is nearly symmetric and not much different from the quadratic loss. When a_{ij} is a large negative number, the LINEX loss is close to exponential when $\widehat{\phi}_{ij} < \phi_{ij}$ and close to linear otherwise. Hence if we suspected that the posterior mean had a downward bias, using the LINEX loss with a negative a_{ij} parameter should help correct the bias. The Bayesian estimates of Φ under these loss functions are well known and are given here for convenience.

Lemma 2. (a) Under the loss $L_{\Phi 1}$, the generalized Bayesian estimator of Φ is

$$\widehat{\Phi}_1 = E(\Phi | \mathbf{Y}). \quad (13)$$

(b) Under the LINEX loss function, each element of the Bayesian estimator $\widehat{\Phi}_2$ satisfies the condition

$$\widehat{\phi}_{ij} = -\frac{1}{a_{ij}} \log[E\{\exp(-a_{ij}\phi_{ij}) | \mathbf{Y}\}] \quad (14)$$

for $i = 1, \dots, Lp + 1$ and $j = 1, \dots, p$.

Proving the lemma is straightforward. The lemma shows that the Bayesian estimate of Φ under the quadratic loss function is the posterior mean, whereas the estimate under the LINEX loss may be larger or smaller than the posterior mean, depending on the sign of the constant a_{ij} .

3.3 Impulse Response Functions

A covariance stationary VAR has the moving average representation $\mathbf{y}'_t = E_0 \mathbf{y}'_t + \sum_{j=0}^{t-1} \epsilon'_{t-j} \mathbf{H}_j$, where \mathbf{H}_0 is the $p \times p$ identity matrix and the impulse responses of \mathbf{y}_t to a shock ϵ_{t-j} occurring j periods earlier is $\mathbf{H}_j = \sum_{i=1}^j \mathbf{B}_i \mathbf{H}_{j-i}$, with $\mathbf{B}_i = \mathbf{0}$ for i larger than L . Note that the components of the vector of errors ϵ_t are correlated because the covariance matrix is unrestricted. For example, the forecasting error of short-term interest rates may be correlated with that of inflation. Suppose that a short-term interest rate is the monetary policy indicator. A monetary policy shock should be represented by a shock in the short-term interest rate uncorrelated with other shocks. Thus more economic meaningful impulse responses are those to orthogonalized (structural) errors. Orthogonalization

of the errors can be achieved through Cholesky decomposition of the covariance matrix, $\Sigma = \Psi' \Psi$, where Ψ is an upper-triangular positive definite matrix. VAR errors are mapped to structural shocks through $\mathbf{u}'_t = \epsilon'_t \Psi^{-1}$. The covariance matrix of the structural error vector \mathbf{u}_t is the identity matrix. The impulse response to structural shocks occurring j periods earlier is given by $\mathbf{Z}_j = \Psi \mathbf{H}_j$. By definition, impulse responses are nonlinear functions of (Φ, Σ) . The nonlinearity makes it difficult to conduct frequentist inference but does not pose difficulty for Bayesian simulations as long as the posteriors of (Φ, Σ) can be simulated.

The loss function for estimation of the impulse response \mathbf{Z}_j (with forecasting horizon j) is assumed to be $L(\mathbf{Z}, \widehat{\mathbf{Z}}_j) = \text{trace}\{(\mathbf{Z}_j - \widehat{\mathbf{Z}}_j)' \Omega (\mathbf{Z}_j - \widehat{\mathbf{Z}}_j)\}$. The matrix Ω weighs the estimation error of each element of the impulse responses. The weighting may be dictated by the economic significance of the element. As long as Ω is constant, the Bayesian estimate $\widehat{\mathbf{Z}}_j$ is the posterior mean $E(\mathbf{Z}_j | \mathbf{Y})$. In the numerical example in this article, we simply set Ω as the identity matrix.

4. PRIORS

Bayesian analysis requires explicit specification of a prior on the parameters. Noninformative priors are commonly used by VAR users, because it is difficult to find a universally justifiable subjective prior. Note that noninformative priors for (Σ, Φ) are not unique because they can be derived according to different principles (see Kass and Wasserman 1996).

The most popular noninformative prior for Σ is the Jeffreys prior (see Geisser 1965; Tiao and Zellner 1964). The Jeffreys prior is derived from the "invariance principle," according to which the prior is invariant to reparameterization (see Jeffreys 1961; Zellner 1971). Specifically, for the VAR covariance matrix, the Jeffreys prior is $\pi_J(\Sigma) \propto |\Sigma|^{-(p+1)/2}$. The prior for Σ in RATS is a modified version of the Jeffreys prior, $\pi_A(\Sigma) \propto |\Sigma|^{-(L+1)p/2-1}$.

It has been noted, however, that the Jeffreys prior often gives unsatisfactory results for multiparameter problems. Bernardo (1979) proposed an information-based approach for deriving a reference prior by breaking down a single multiparameter problem into consecutive problems with fewer numbers of parameters. The form of the reference prior depends on the inferential problem at hand and on researchers' ordering of parameters in terms of perceived importance. (For examples where the reference priors produce more desirable estimates than the Jeffreys priors, see Berger and Bernardo 1992 and Sun and Berger 1998, among others.) In estimating the variance-covariance matrix Σ based on an iid random sample from a normal population with known mean, Yang and Berger (1994) reparameterized the matrix Σ as $\mathbf{O}' \mathbf{D} \mathbf{O}$, where \mathbf{D} is a diagonal matrix whose elements are the eigenvalues of Σ (in increasing or decreasing order) and \mathbf{O} is an orthogonal matrix. The following reference prior is derived by giving priority to vectorized \mathbf{D} over vectorized \mathbf{O} : $\pi_R(\Sigma) \propto |\Sigma|^{-1} \prod_{1 \leq i < j \leq p} (\lambda_i - \lambda_j)^{-1}$, where $\lambda_1 > \lambda_2 > \dots > \lambda_p$ are eigenvalues of Σ .

The prior for (Φ, Σ) can be obtained by putting together priors for Φ and Σ . In practice, it is often convenient to consider the vectorized Φ , $\text{vec}(\Phi)$, which we denote by ϕ . A common

expression of ignorance about $\boldsymbol{\phi}$ is a (flat) constant prior. A popular noninformative prior for multivariate regression models, called the diffuse prior, consists of a constant prior for $\boldsymbol{\phi}$ and the Jeffreys prior for $\boldsymbol{\Sigma}$. The joint density of the constant Jeffreys prior for $(\boldsymbol{\Phi}, \boldsymbol{\Sigma})$ [or $(\boldsymbol{\phi}, \boldsymbol{\Sigma})$] is in the form $\pi_{\text{CJ}}(\boldsymbol{\phi}, \boldsymbol{\Sigma}) \propto \pi_{\text{J}}(\boldsymbol{\Sigma})$. The constant RATS prior $\pi_{\text{CA}}(\boldsymbol{\Phi}, \boldsymbol{\Sigma}) \propto \pi_{\text{A}}(\boldsymbol{\Sigma})$ is the default choice in RATS and has been used in hundreds of published works in empirical macroeconomics. (For an argument touting the constant RATS instead of the constant Jeffreys prior, see Sims and Zha 1999.) Finally, the constant reference prior, which has not been commonly used for VAR models, takes the form $\pi_{\text{CR}}(\boldsymbol{\phi}, \boldsymbol{\Sigma}) \propto \pi_{\text{R}}(\boldsymbol{\Sigma})$.

For estimating an unknown multivariate normal mean, motivated by Stein's (1956) result on inadmissibility of the MLE, some authors (e.g., Baranchik 1964; Berger and Strawderman 1996) have advocated the following "shrinkage" prior as an alternative to the constant prior for $\boldsymbol{\phi}$:

$$\pi_{\text{S}}(\boldsymbol{\phi}) \propto \|\boldsymbol{\phi}\|^{-(J-2)}, \quad \boldsymbol{\phi} \in \mathbb{R}^J, \quad (15)$$

where the dimension of $\boldsymbol{\phi}$, J , equals $p(Lp + 1)$. Note that the prior (15) has a two-stage hierarchical structure. Let $\pi_{\text{S}}(\boldsymbol{\phi}|\delta)$ be the normal density $N_{\text{J}}(\mathbf{0}, \delta \mathbf{I}_J)$. We can write

$$(\boldsymbol{\phi}|\delta) \sim N_{\text{J}}(\mathbf{0}, \delta \mathbf{I}_J) \quad \text{and} \quad \pi(\delta) \propto 1, \quad (16)$$

because

$$\begin{aligned} \int_0^\infty \pi_{\text{S}}(\boldsymbol{\phi}|\delta) \pi(\delta) d\delta &\propto \int_0^\infty \frac{1}{(2\pi\delta)^{J/2}} \exp\left\{-\frac{1}{2\delta} \boldsymbol{\phi}'\boldsymbol{\phi}\right\} d\delta \\ &= \frac{\Gamma(J/2)}{(2\pi)^{J/2} (\boldsymbol{\phi}'\boldsymbol{\phi})^{J/2-1}}, \end{aligned}$$

which is proportional to (15).

We consider three noninformative priors for $(\boldsymbol{\Phi}, \boldsymbol{\Sigma})$ with the shrinkage prior on $\boldsymbol{\Phi}$: the shrinkage Jeffreys prior $\pi_{\text{SJ}}(\boldsymbol{\phi}, \boldsymbol{\Sigma}) = \pi_{\text{S}}(\boldsymbol{\phi})\pi_{\text{J}}(\boldsymbol{\Sigma})$, the shrinkage RATS prior $\pi_{\text{SA}}(\boldsymbol{\phi}, \boldsymbol{\Sigma}) = \pi_{\text{S}}(\boldsymbol{\phi}) \times \pi_{\text{A}}(\boldsymbol{\Sigma})$, and the shrinkage reference prior $\pi_{\text{SR}}(\boldsymbol{\phi}, \boldsymbol{\Sigma}) = \pi_{\text{S}}(\boldsymbol{\phi}) \times \pi_{\text{R}}(\boldsymbol{\Sigma})$. The noninformative priors for $(\boldsymbol{\Phi}, \boldsymbol{\Sigma})$ are improper (i.e., the integrals of these in the parameter space are infinite). Ni and Sun (2003) and Sun and Ni (2004) provided conditions under which the posteriors of $(\boldsymbol{\Phi}, \boldsymbol{\Sigma})$ are proper with the prior combinations considered in this article.

We now turn to the prior for the degrees-of-freedom parameter, ν , in the Student- t errors. The parameter ν in the Student- t likelihood function is commonly considered a positive integer but can be any positive number. We apply a gamma (a, b) prior for $w = \nu/2$, the density of which is given by

$$[w] = \frac{b^a}{\Gamma(a)} w^{a-1} \exp(-bw), \quad w > 0, \quad (17)$$

where (a, b) are positive known constants. When $a = 1$, it is the exponential prior given by Geweke (1993).

5. BAYESIAN COMPUTATION AND CONDITIONAL POSTERiors

The posterior quantities of $(\boldsymbol{\phi}, \boldsymbol{\Sigma})$ are not available in closed form for most of the priors that we consider. In this study we use Gibbs sampling Markov chain Monte Carlo (MCMC) methods to sample from the posteriors (see Gelfand and Smith 1990). The first step of the MCMC computation is to find the full conditional distributions of $(\boldsymbol{\phi}, \boldsymbol{\Sigma})$.

5.1 Conditional Posteriors Under the Normal Errors

We make use of the following results for the normal model.

Fact 1. Consider the constant Jeffreys prior for $(\boldsymbol{\phi}, \boldsymbol{\Sigma})$.

- The conditional posterior of $\boldsymbol{\phi}$ given $(\boldsymbol{\Sigma}, \mathbf{Y})$ is $N_{\text{J}}(\widehat{\boldsymbol{\phi}}_{\text{MLE}}, \boldsymbol{\Sigma} \otimes (\mathbf{X}'\mathbf{X})^{-1})$
- The marginal posterior of $\boldsymbol{\Sigma}$ given \mathbf{Y} is inverse Wishart $(\mathbf{S}(\widehat{\boldsymbol{\phi}}_{\text{MLE}}), T - Lp - 1)$. Here we use Anderson's (1984) definition of the inverse Wishart distribution.

Fact 2. Consider the constant RATS prior π_{CA} .

- The conditional posterior of $\boldsymbol{\phi}$ given $(\boldsymbol{\Sigma}; \mathbf{Y})$ is $N_{\text{J}}(\widehat{\boldsymbol{\phi}}_{\text{MLE}}, \boldsymbol{\Sigma} \otimes (\mathbf{X}'\mathbf{X})^{-1})$.
- The marginal posterior of $\boldsymbol{\Sigma}$ given \mathbf{Y} is inverse Wishart $(\mathbf{S}(\widehat{\boldsymbol{\phi}}_{\text{MLE}}), T)$.

Fact 3. Consider the constant reference prior π_{CR} .

- The conditional distribution of $\boldsymbol{\phi}$ given $(\boldsymbol{\Sigma}; \mathbf{Y})$ is $N_{\text{J}}(\widehat{\boldsymbol{\phi}}_{\text{MLE}}, \boldsymbol{\Sigma} \otimes (\mathbf{X}'\mathbf{X})^{-1})$.
- The conditional density of $\boldsymbol{\Sigma}$ given $(\boldsymbol{\phi}; \mathbf{Y})$ is

$$\pi(\boldsymbol{\Sigma}|\boldsymbol{\phi}; \mathbf{Y}) \propto \frac{\text{etr}\{-1/2\boldsymbol{\Sigma}^{-1}\mathbf{S}(\boldsymbol{\Phi})\}}{|\boldsymbol{\Sigma}|^{T/2+1} \prod_{1 \leq i < j \leq p} (\lambda_i - \lambda_j)}, \quad (18)$$

where $\mathbf{S}(\boldsymbol{\Phi})$ is defined by (5).

For the shrinkage prior, the hierarchical structure (16) suggests a nice computational formula. Instead of simulating from the conditional distribution of $\boldsymbol{\Phi}$ and $\boldsymbol{\Sigma}$ within each Gibbs cycle, we use δ as a latent variable and simulate from $\boldsymbol{\Phi}$, $\boldsymbol{\Sigma}$, and δ based on the following fact.

Fact 4. Consider the shrinkage reference prior π_{SR} . The full conditional posteriors of $(\boldsymbol{\Sigma}, \boldsymbol{\phi}, \delta)$ given data \mathbf{Y} are as follows:

- The conditional density of $\boldsymbol{\Sigma}$ given $(\boldsymbol{\phi}, \delta; \mathbf{Y})$ is given in (18).
- The conditional distribution of $\boldsymbol{\phi}$ given $(\delta, \boldsymbol{\Sigma}; \mathbf{Y})$ is $N_{\text{J}}(\boldsymbol{\mu}_{\text{S}}, \mathbf{V}_{\text{S}})$, where

$$\boldsymbol{\mu}_{\text{S}} = \delta(\boldsymbol{\Sigma} \otimes (\mathbf{X}'\mathbf{X})^{-1} + \delta \mathbf{I}_J)^{-1} \widehat{\boldsymbol{\phi}}_{\text{MLE}} \quad (19)$$

and

$$\mathbf{V}_{\text{S}} = \left(\boldsymbol{\Sigma}^{-1} \otimes (\mathbf{X}'\mathbf{X}) + \frac{1}{\delta} \mathbf{I}_J \right)^{-1}. \quad (20)$$

- The conditional distribution of δ , given $(\boldsymbol{\phi}, \boldsymbol{\Sigma}; \mathbf{Y})$, is inverse gamma $(\frac{J}{2} - 1, \frac{1}{2} \boldsymbol{\phi}'\boldsymbol{\phi})$.

These conditional posteriors are used for MCMC simulations of the posteriors of $(\boldsymbol{\Phi}, \boldsymbol{\Sigma})$. Ni and Sun (2003) offered proofs for some of the foregoing facts as well as detailed algorithms of MCMC procedures. The Bayesian estimates of $(\boldsymbol{\Phi}, \boldsymbol{\Sigma})$ under alternative loss functions and priors can be computed when the posterior distributions are simulated.

5.2 Conditional Posteriors Under Student- t Errors

The conditional posterior distributions of Φ and Σ under the t -errors are quite complicated. We can, however, use the hierarchical structure of the t -distribution to introduce latent variables. In particular, let $(s|q) \sim N_p(\mathbf{0}, q^{-1}\Sigma)$, where $q \sim \text{gamma}(v/2, v/2)$ has density given by

$$p(q|v) = \frac{(v/2)^{v/2} q^{v/2-1}}{\Gamma(v/2)} \exp\left(-\frac{vq}{2}\right) \quad \text{for } q > 0. \quad (21)$$

It follows that the marginal density of s is

$$\begin{aligned} p(s|\Sigma, v) &= \int_0^\infty p(s|q, \Sigma, v) p(q|v) dq \\ &\propto |\Sigma|^{-1/2} (v + s' \Sigma^{-1} s)^{-(v+p)/2}, \end{aligned}$$

which is the same as the density of $t_v(\mathbf{0}, \Sigma)$ given by (6).

Now we introduce independent latent variables q_1, \dots, q_T that are iid $\sim \text{gamma}(v/2, v/2)$. For given q_1, \dots, q_T and Σ , VAR error vectors $\epsilon_1, \dots, \epsilon_T$ are independent, with $\epsilon_t \sim N_p(\mathbf{0}, q_t^{-1}\Sigma)$. Let $\mathbf{Q} = \text{diag}(q_1, \dots, q_T)$. The joint likelihood of $(\Phi, \Sigma, \mathbf{Q}, v)$ is then

$$\begin{aligned} l^*(\mathbf{Q}, \Phi, \Sigma, v) &\propto \frac{(v/2)^{vT/2} \prod_{t=1}^T q_t^{(p+v)/2-1} \exp(-vq_t/2)}{[\Gamma(v/2)]^T |\Sigma|^{T/2}} \\ &\quad \times \exp\left\{-\sum_{t=1}^T \frac{q_t (\mathbf{y}_t - \mathbf{x}_t \Phi)' \Sigma^{-1} (\mathbf{y}_t - \mathbf{x}_t \Phi)}{2}\right\} \\ &= \frac{(v/2)^{vT/2} \prod_{t=1}^T q_t^{(p+v)/2-1} \exp(-vq_t/2)}{[\Gamma(v/2)]^T |\Sigma|^{T/2}} \\ &\quad \times \text{etr}\left\{-\frac{1}{2} \mathbf{S}_{\mathbf{Q}}(\Phi) \Sigma^{-1}\right\}, \quad (22) \end{aligned}$$

where

$$\mathbf{S}_{\mathbf{Q}}(\Phi) = (\mathbf{Y} - \mathbf{X}\Phi)' \mathbf{Q} (\mathbf{Y} - \mathbf{X}\Phi). \quad (23)$$

The following facts are useful for Bayesian simulation in the Student- t model.

Fact 5. Consider the constant Jeffreys prior for (ϕ, Σ) and gamma(a, b) prior for $w = v/2$.

- (a) The conditional posterior of ϕ given $(\Sigma, \mathbf{Q}, v; \mathbf{Y})$ is $N(\hat{\phi}_{\mathbf{Q}}, \Sigma \otimes (\mathbf{X}'\mathbf{Q}\mathbf{X})^{-1})$. Here $\hat{\phi}_{\mathbf{Q}}$ is defined as vectorized $\hat{\Phi}_{\mathbf{Q}}$, and

$$\hat{\Phi}_{\mathbf{Q}} = (\mathbf{X}'\mathbf{Q}\mathbf{X})^{-1} \mathbf{X}'\mathbf{Q}\mathbf{Y}. \quad (24)$$

- (b) The conditional posterior of Σ given $(\phi, \mathbf{Q}, v; \mathbf{Y})$ is inverse Wishart $(\mathbf{S}_{\mathbf{Q}}(\Phi), T)$.
(c) Given $(\phi, \Sigma, v; \mathbf{Y})$, q_1, \dots, q_T are independent and

$$\begin{aligned} (q_t | \phi, \Sigma, v; \mathbf{Y}) \\ \sim \text{gamma}\left(\frac{1}{2}(v+p), \frac{1}{2}\{v + (\mathbf{y}_t - \mathbf{x}_t \Phi)' \Sigma^{-1} (\mathbf{y}_t - \mathbf{x}_t \Phi)\}\right). \quad (25) \end{aligned}$$

- (d) Given $(\phi, \Sigma, \mathbf{Q}; \mathbf{Y})$, the conditional density of $w = v/2$ has the form

$$\begin{aligned} [w | \phi, \Sigma, \mathbf{Q}; \mathbf{Y}] \\ = \frac{w^{Tw+a-1}}{[\Gamma(w)]^T} \left\{ \prod_{t=1}^T q_t \right\}^w \exp\left\{-\left(b + \sum_{t=1}^T q_t\right)w\right\}, \end{aligned}$$

for $w > 0$. (26)

- (e) If $T/2 + a - 1 > 0$, then the conditional posterior of w in part (d) is log-concave.

Proof. The proof of parts (a)–(d) is easy. For part (e), denote $g(w) = \log[w | \phi, \Sigma, \mathbf{Q}; \mathbf{Y}]$. It is enough to show that $\frac{\partial^2}{\partial w^2} g(w) < 0$. Clearly,

$$\begin{aligned} g(w) &= (Tw + a - 1) \log(w) - T \log\{\Gamma(w)\} \\ &\quad + w \log\left(\prod_{t=1}^T q_t\right) - \left(b + \sum_{t=1}^T q_t\right)w. \end{aligned}$$

It follows from formula (1.46) of Bowman and Shenton (1988) that

$$\frac{d^2}{dw^2} \log[\Gamma(w)] = \frac{1}{w} + \frac{1}{2w^2} + \frac{2\pi}{w} \int_0^\infty \frac{y\sqrt{s}}{(w^2 + s)(y-1)^2} ds,$$

where $y = \exp(2\pi\sqrt{s})$. Consequently,

$$\frac{d^2}{dw^2} g(w) = -\frac{T/2 + a - 1}{w^2} - \frac{2\pi T}{w} \int_0^\infty \frac{y\sqrt{s}}{(w^2 + s)(y-1)^2} ds,$$

which is negative if $a + T/2 - 1 > 0$.

For MCMC simulations, one can sample successively from the full conditional distribution, given by Fact 5(a)–(d), treating $(\phi, \Sigma, \mathbf{Q}, v)$ in four blocks in updating. Alternatively, we can divide $(\phi, \Sigma, \mathbf{Q}, v)$ into three blocks, \mathbf{Q} , (ϕ, Σ) , and v . We first sample $(\mathbf{Q} | \phi, \Sigma, v; \mathbf{Y})$ according to Fact 5(c); within the block (ϕ, Σ) , we sample $(\phi | \Sigma, \mathbf{Q}, v; \mathbf{Y})$ according to Fact 5(a), then instead of using property Fact 5(b), we sample $(\Sigma | \mathbf{Q}, v; \mathbf{Y})$ from inverse Wishart $(\mathbf{S}_{\mathbf{Q}}(\hat{\Phi}_{\mathbf{Q}}), T - Lp - 1)$. Chib, Osiewalski, and Steel (1991) offered a theoretical discussion on simulation of the degrees of freedom parameter v . Gilks and Wild (1992) proposed an adaptive rejection sampling scheme for simulating from a log-concave density. We will adopt Gilks and Wild's method for simulating the conditional posterior of w based on Fact 5(e).

Fact 6. Consider π_{CA} , the constant RATS prior for (ϕ, Σ) , and gamma(a, b) for $w = v/2$. The conditional posteriors of ϕ , \mathbf{Q} , and v are the same as parts (a), (c), and (d) in Fact 5, whereas the conditional posterior of Σ given $(\mathbf{Q}, v; \mathbf{Y})$ is inverse Wishart $(\mathbf{S}_{\mathbf{Q}}(\hat{\Phi}_{\mathbf{Q}}), T)$.

Fact 7. Consider π_{CR} , the constant reference prior for (ϕ, Σ) , and gamma(a, b) prior for $w = v/2$. The conditional posteriors of ϕ , \mathbf{Q} , and v are the same as parts (a), (c), and (d) in Fact 5, whereas the conditional posterior density of Σ given $(\phi, \mathbf{Q}, v; \mathbf{Y})$ is

$$\pi(\Sigma | \phi, \mathbf{Q}, v; \mathbf{Y}) \propto \frac{\text{etr}\{-1/2 \Sigma^{-1} \mathbf{S}_{\mathbf{Q}}(\phi)\}}{|\Sigma|^{T/2+1} \prod_{1 \leq i < j \leq p} (\lambda_i - \lambda_j)}. \quad (27)$$

Fact 8. Consider π_{SR} , the shrinkage reference prior and gamma(a, b) prior for $w = v/2$.

- (a) The conditional distribution of ϕ given $(\delta, \Sigma, \mathbf{Q}, v; \mathbf{Y})$ is $N_J(\mu_{\mathbf{Q}}, \mathbf{V}_{\mathbf{Q}})$, where

$$\mu_{\mathbf{Q}} = \delta(\Sigma \otimes (\mathbf{X}'\mathbf{Q}\mathbf{X})^{-1} + \delta\mathbf{I}_J)^{-1}\hat{\phi}_{\mathbf{Q}} \quad (28)$$

and

$$\mathbf{V}_{\mathbf{Q}} = \left(\Sigma^{-1} \otimes \mathbf{X}'\mathbf{Q}\mathbf{X} + \frac{1}{\delta}\mathbf{I}_J \right)^{-1}. \quad (29)$$

- (b) The conditional distribution of δ given $(\phi, \Sigma, \mathbf{Q}, v; \mathbf{Y})$ is inverse gamma($\frac{J}{2} - 1, \frac{1}{2}\phi'\phi$).
- (c) The conditional density of Σ given $(\phi, \delta, \mathbf{Q}; \mathbf{Y})$ is given by (27).
- (d) For $t = 1, \dots, T$, $(q_t|\phi, \delta, \Sigma, v; \mathbf{Y})$ has the distribution (25).
- (e) Given $(\phi, \delta, \Sigma, \mathbf{Q}; \mathbf{Y})$, the conditional posterior of $w = v/2$ is the same as Fact 5(d).

6. NUMERICAL EXAMPLES

6.1 Simulation Results

In what follows we use a numerical example to evaluate the properties of competing Bayesian estimates. We first generate data samples from a VAR with known parameters, then compute the Bayesian estimates via MCMC simulations. Using the MCMC results, we evaluate the performance of the Bayesian estimates under competing priors in terms of the frequentist risks given the true parameters. The frequentist risks under a loss L [i.e., $E_{\mathbf{Y}|\Phi, \Sigma}L(\Sigma, \hat{\Sigma})$ and $E_{\mathbf{Y}|\Phi, \Sigma}L(\Phi, \hat{\Phi})$] of the MLE and Bayesian estimates using alternative priors on Σ and Φ are estimated as the average loss pertaining to estimates $\hat{\Sigma}$ and $\hat{\Phi}$ across generated data samples. For instance, $E_{\mathbf{Y}|\Phi, \Sigma}L(\Phi, \hat{\Phi})$ is estimated as $\frac{1}{N} \sum_{n=1}^N L(\Phi, \hat{\Phi}_n)$. Here $\hat{\Phi}_n$ is the Bayesian estimate from sample n ($n = 1, \dots, N$), and Φ is the true parameter matrix Φ .

We denote the estimates according to a loss function and a prior. For example, $\hat{\Sigma}_{1CA}$ represents the estimator of Σ under loss $L_{\Sigma 1}$ and the constant RATS prior combination, and $\hat{\Phi}_{2SR}$ represents the estimator of Φ under loss $L_{\Phi 2}$ and the shrinkage reference prior combination. Each row of Tables 1 and 2 reports the frequentist average and standard deviations of losses of the corresponding estimator under different loss functions.

As mentioned in Section 1, besides frequentist risks of estimates of parameters of the VAR, we are interested in estimates of certain nonlinear functions of these parameters, such as impulse responses. Losses pertaining to impulse responses are computed as follows. For the n th dataset generated in the experiment, we denote the impulse response matrix for the i th step after the shock by $\hat{\mathbf{Z}}_i^{(n)}$. The accuracy in estimation of the impulse responses (with forecasting horizon i) is measured by the frequentist average of sum of squared errors,

$$R_{Imp,i} = \frac{1}{N} \sum_{n=1}^N \text{trace}\{(\mathbf{Z}_i - \hat{\mathbf{Z}}_i^{(n)})'(\mathbf{Z}_i - \hat{\mathbf{Z}}_i^{(n)})\}. \quad (30)$$

Table 1. Frequentist Average Losses (with standard deviations in parentheses) of Competing Bayes Estimates of Σ in the Numerical Example With Normal Errors

	$L_{\Sigma 1}$	$L_{\Sigma 2}$	$L_{\Sigma 3}$
$\hat{\Sigma}_{MLE}$.861 (.403)	.681 (.189)	.516 (.178)
$\hat{\Sigma}_{1CA}$.608 (.308)	.646 (.222)	.415 (.153)
$\hat{\Sigma}_{2CA}$	1.187 (.501)	.803 (.192)	.660 (.205)
$\hat{\Sigma}_{3CA}$.861 (.403)	.681 (.189)	.516 (.178)
$\hat{\Sigma}_{1CJ}$.450 (.222)	.800 (.344)	.389 (.142)
$\hat{\Sigma}_{2CJ}$.862 (.403)	.681 (.189)	.516 (.178)
$\hat{\Sigma}_{3CJ}$.609 (.309)	.645 (.222)	.415 (.153)
$\hat{\Sigma}_{1CR}$.281 (.172)	.434 (.219)	.234 (.113)
$\hat{\Sigma}_{2CR}$.546 (.314)	.489 (.184)	.353 (.161)
$\hat{\Sigma}_{3CR}$.386 (.238)	.419 (.178)	.273 (.133)
$\hat{\Sigma}_{1SA}$.609 (.308)	.646 (.222)	.415 (.153)
$\hat{\Sigma}_{2SA}$	1.187 (.501)	.803 (.192)	.661 (.205)
$\hat{\Sigma}_{3SA}$.862 (.403)	.682 (.189)	.516 (.178)
$\hat{\Sigma}_{1SJ}$.449 (.221)	.801 (.345)	.389 (.143)
$\hat{\Sigma}_{2SJ}$.862 (.403)	.682 (.189)	.516 (.178)
$\hat{\Sigma}_{3SJ}$.609 (.308)	.646 (.222)	.415 (.153)
$\hat{\Sigma}_{1SR}$.261 (.163)	.418 (.208)	.221 (.107)
$\hat{\Sigma}_{2SR}$.505 (.300)	.464 (.177)	.332 (.154)
$\hat{\Sigma}_{3SR}$.356 (.226)	.399 (.169)	.256 (.126)

Besides to evaluating the overall performance of VAR estimates, we also investigate stationarity of the estimated VARs. VARs with one lag are stationary if the absolute values of the real eigenvalues (and amplitudes of the complex eigenvalues) of the regression coefficient matrix \mathbf{B}_1 are less than unity. VARs of more than one lag can be rewritten as a VAR with one lag through transformation of variables. For simplicity, simulations in this study were based on a VAR with one lag. For each prior, we compared the eigenvalues of the posterior mean of the \mathbf{B}_1 matrix with those used for generating the data.

We generated $N = 1,000$ samples from a five-variable VAR model with one lag (i.e., $p = 5$ and $L = 1$) with true parameters given in the following. The elements of Σ are $\sigma_{i,j} = 1.0$ if $i = j$ and .5 otherwise, which implies that the VAR errors have a pairwise correlation of .5. The elements of Φ are $\mathbf{c} = \mathbf{0}$ and $\mathbf{B}_1 = \mathbf{I}_5$. By assumption, the VAR regression coefficients contains unit roots. The number of sample observations, T , is 50. In computing the Bayesian estimates of (Φ, Σ) , we ran $M = 10,000$ MCMC cycles with 500 burn-in cycles for each of the 1,000 samples. We choose the weighting matrix \mathbf{W} in the loss function $L_{\Phi 1}$ to be the identity matrix. The parameter a in

Table 2. Frequentist Average Losses (with standard deviations in parentheses) of Competing Bayes Estimates of Φ in the Numerical Example With Normal Errors

	$L_{\Phi 1}$	$L_{\Phi 2}$
$\hat{\Phi}_{MLE}$	11.183 (14.070)	11.614 (6.898)
$\hat{\Phi}_{1CA}$	11.184 (14.068)	11.611 (6.893)
$\hat{\Phi}_{2CA}$	11.135 (14.055)	9.416 (5.154)
$\hat{\Phi}_{1CJ}$	11.184 (14.094)	11.613 (6.905)
$\hat{\Phi}_{2CJ}$	11.134 (14.079)	9.156 (4.947)
$\hat{\Phi}_{1CR}$	11.185 (14.070)	11.615 (6.894)
$\hat{\Phi}_{2CR}$	11.135 (14.057)	9.319 (5.066)
$\hat{\Phi}_{1SA}$	1.552 (.597)	7.254 (3.349)
$\hat{\Phi}_{2SA}$	1.523 (.596)	6.313 (2.786)
$\hat{\Phi}_{1SJ}$	1.419 (.501)	7.174 (3.250)
$\hat{\Phi}_{2SJ}$	1.387 (.499)	6.155 (2.653)
$\hat{\Phi}_{1SR}$	1.261 (.359)	7.176 (3.332)
$\hat{\Phi}_{2SR}$	1.231 (.355)	6.198 (2.739)

loss $L_{\Sigma 2}$ is set as follows. We let the parameter corresponding to the intercept term ϕ_{1j} be .001, implying a near-symmetric loss, and let that parameter corresponding to nonintercept term be -4.0 . (The value -4 was used in Zellner 1986 for estimating the normal means.) We first consider the normal model and then a Student- t model.

Simulations for the six prior combinations took less than 2 minutes for each of the 1,000 samples on a Pentium 4 1.7-GHz PC. The simulations converged quite quickly. There is little difference when the Markov chain is shortened to 5,000 from 10,000. The acceptance rates for the Metropolis step for simulation of Σ under the reference prior are about 30%.

Several conclusions can be drawn from the numerical example:

a. The estimates of Σ are largely unrelated to the priors on Φ , and estimates of Φ are not affected much by the priors on Σ . Tables 1 and 2 report the frequentist risks of the estimates of Σ and Φ under competing priors for the VAR with normal errors. In Table 1 the Σ -related losses of $\hat{\Sigma}_{1CA}$ (based on the constant RATS prior) are quite similar to those of $\hat{\Sigma}_{1SA}$ (based on the shrinkage RATS prior), and in Table 2 the Φ -related losses of $\hat{\Phi}_{1CA}$ (based on the constant RATS prior) are similar to those of $\hat{\Phi}_{1CJ}$ (based on the constant Jeffreys prior) and $\hat{\Phi}_{1CR}$ (based on the constant reference prior). This is due to the assumption that the loss functions are separable in Σ and Φ .

b. Under the constant prior, the estimator $\hat{\Phi}_2$ obtained from minimizing the LINEX loss function is slightly worse than estimator $\hat{\Phi}_1$ in terms of quadratic loss, but it is moderately better in terms of LINEX loss. Hence overall $\hat{\Phi}_2$ is better than the posterior mean. The Bayesian estimates of Φ under the shrinkage prior yield smaller frequentist average losses than their counterparts under the constant prior.

c. In terms of frequentist risk, the influence of prior exceeds that of the loss functions. Table 1 shows that the reference prior dominates other priors on Σ . Under the reference prior with any loss function, the average losses associated with Σ are reduced for all estimates. Table 2 shows that the Bayesian estimates of Φ based on the shrinkage prior yield smaller frequentist risks than their counterparts based on a constant prior. Even with strong asymmetry ($a = -4.0$) in the loss function $L_{\Phi 2}$, there is not much difference between the two estimates of Φ under the shrinkage prior. In contrast, prior choice is critically important. Under the constant prior, both Bayesian estimates perform poorly, even in comparison with the MLE. The large standard deviations of the frequentist losses under the constant prior indicate that these Bayesian estimates are not successful in dealing with outliers. In comparison, under the shrinkage prior, both estimates dominate the MLE in terms of the frequentist risk associated with Φ . The Bayesian estimator under the shrinkage prior improves over the MLE by reducing bias for some elements and at the same time substantially reducing variances. Table 3 reports the frequentist risks of impulse responses for the normal model. The frequentist risks of impulse responses are smaller under the shrinkage prior than under the constant prior.

d. A question of interest for applied econometricians is whether the conclusions on comparison of priors are robust when the VAR errors are not normal. We find that they are. We simulate 1,000 samples using the Student- t model with degrees of freedom $\nu = 8$ and the parameter values for Σ and Φ

Table 3. Frequentist Average Losses of Impulse Responses in the Numerical Example With Normal Errors

Horizon	MLE	CA	CJ	CR	SA	SJ	SR
1	.963	.911	.848	.854	.734	.662	.640
2	1.745	1.657	1.603	1.640	1.331	1.249	1.250
3	2.389	2.276	2.238	2.290	1.851	1.772	1.781
4	2.898	2.767	2.747	2.804	2.282	2.208	2.220
5	3.302	3.158	3.152	3.210	2.638	2.568	2.580
6	3.625	3.470	3.476	3.533	2.931	2.867	2.877
7	3.886	3.723	3.738	3.793	3.176	3.116	3.125
8	4.101	3.929	3.953	4.005	3.382	3.325	3.332
9	4.278	4.101	4.131	4.180	3.556	3.502	3.508
10	4.427	4.246	4.283	4.330	3.706	3.653	3.659
11	4.554	4.371	4.417	4.460	3.835	3.785	3.789
12	4.662	4.483	4.540	4.579	3.947	3.900	3.903

unchanged. We find similar performance of the estimates and priors as in the normal model. To save space, we do not report the counterparts of Tables 1 and 2. We use $\text{gamma}(1, .5)$ as the prior for the nuisance parameter ν . Experiments with alternative priors on ν show that our general conclusion is robust. Table 4 reports frequentist losses of impulse responses, with the normal VAR errors replaced by Student- t -distributed errors and degrees of freedom treated as an unknown parameter. The posteriors are simulated under the algorithm for t -errors given in the Appendix. Consistent with our finding on the estimates of VAR parameters under competing priors and data-generating models, Table 4 is not qualitatively different from Table 3.

e. Economists are often concerned with inference on stationarity of time series models. Much research has been done on the properties of MLE of autoregressive (AR) models. It is known that the MLE of the AR(1) coefficient with unit root has a downward bias. MacKinnon and Smith (1998) showed that the downward bias of ML estimates for an AR(1) coefficient is nonlinear in the true AR parameter. When the true parameter is near unity, the downward bias is particularly severe. There is a rich Bayesian literature on the unit root model. Phillips (1991) and Sims (1991) debated the merit of using a constant prior for the AR coefficient. Stambaugh (1999) conducted Bayesian analysis of a model with two equations, one a regression of asset return on lagged dividend yield and the other an AR(1) equation of the dividend yield. He showed that the posterior mean of the slope parameter in the asset return equation is smaller than the bias-corrected ordinary least squares (OLS), whereas that of the dividend yield equation is larger. Zellner and Hong

Table 4. Frequentist Average Losses of Impulse Responses in the Numerical Example With Student- t Errors (true degrees of freedom parameter $\nu = 8$)

Horizon	CA	CJ	CR	SA	SJ	SR
1	.930	.867	.875	.747	.673	.654
2	1.684	1.630	1.670	1.334	1.252	1.255
3	2.307	2.272	2.327	1.843	1.764	1.776
4	2.798	2.782	2.842	2.264	2.190	2.206
5	3.186	3.184	3.246	2.610	2.541	2.558
6	3.492	3.502	3.564	2.895	2.831	2.849
7	3.737	3.756	3.816	3.133	3.073	3.090
8	3.934	3.960	4.019	3.333	3.276	3.293
9	4.095	4.127	4.184	3.502	3.448	3.465
10	4.229	4.266	4.321	3.646	3.595	3.612
11	4.342	4.386	4.438	3.771	3.721	3.738
12	4.441	4.493	4.542	3.879	3.832	3.848

(1989) estimated an AR(3) model for gross domestic product (GDP) growth and found that 80% of the posterior eigenvalues had a complex pair.

We now examine the posterior mean of eigenvalues of the VAR regression coefficient matrix \mathbf{B}_1 . The true coefficient matrix of the VAR in the numerical example has five eigenvalues of unity. We found that all eigenvalues of the MLE and posterior mean of Bayesian estimates of the matrix were smaller in magnitude than the true value. About 40% of the eigenvalues in the 1,000 samples turned out to be complex numbers. The imaginary parts of these complex numbers were much smaller than the real parts. In what follows we focus on the eigenvalues with the largest and smallest amplitude. For the normal model, the average of the largest eigenvalues of the MLE over the 1,000 samples is .960, and the standard deviation is .036. The average of the smallest eigenvalues is .510, with a standard deviation of .145. There is clearly a downward bias even for the largest eigenvalues. For the 1,000 samples, the statistics of the eigenvalues for posterior means of \mathbf{B}_1 under constant prior were almost identical to those of the MLE. The well-known downward bias of the MLE for AR(1) coefficients exists for the eigenvalues of the posterior mean of \mathbf{B}_1 as well under the constant prior. Under the shrinkage prior, the downward bias in eigenvalues of \mathbf{B}_1 improves slightly. For example, the average of largest eigenvalues under the shrinkage reference prior is .965, and the standard deviation is .030. The average of the smallest eigenvalues is .534, and the standard deviation is .133. A closer examination of the impulse responses reveals from a different angle a pattern of the downward bias of the MLE and the Bayesian estimates under the constant prior. For instance, the true impulse responses of the fifth variable to a shock of the fifth equation at the first lag (denoted as $z_{1,(5,5)}$) is .775, and at the 12th lag, the element $z_{12,(5,5)}$ is .775 as well. The frequentist averages of the same elements corresponding to the MLE are $\hat{z}_{1,(5,5)} = .397$ and $\hat{z}_{12,(5,5)} = .004$. Similar numbers are obtained for Bayesian estimates under the constant priors. The same elements corresponding to the Bayesian estimates under the shrinkage reference prior are $\hat{z}_{1,(5,5)} = .605$ and $\hat{z}_{12,(5,5)} = .216$. The downward bias under the shrinkage prior is still quite sizeable but is much smaller than that of the MLE and Bayesian estimates under the constant prior.

We conducted more simulations with several parameter settings. To save space, we do not report the details of the simulations, but briefly describe the results instead. When the sample size T is increased, the role of the prior diminishes. When the size of the VAR, p , or the number of lags, L , increases, the opposite is true. When the degrees of freedom parameter of the Student- t distribution ν increases, the results become more similar to the case with normal errors. The shrinkage prior tends to shrink the intercept terms more than the nonintercept terms. When the true value of the intercept term is set at a positive value, the dominance of the shrinkage prior over the constant prior becomes less prominent with respect to the symmetric quadratic loss. Stronger correlations of the VAR variables amplify the variances of MLEs and the variances of Bayesian estimates under the constant prior. The shrinkage prior is effective in reducing the variances of the Bayesian estimates and results in smaller losses compared to the constant prior when the VAR errors have strong pairwise correlations. Finally, regarding the

loss function, when the parameter in the asymmetric loss function a is set at -8 instead of -4 , the Bayesian estimator $\hat{\Phi}_2$ becomes more different from the posterior mean $\hat{\Phi}_1$. The downward bias of the posterior mean is better corrected by the estimator based on the loss function with stronger asymmetry.

The foregoing exercises demonstrate that the frequentist properties of alternative estimates can be substantially different under competing priors. In the next section we compare Bayesian VAR estimates of the U.S. economy with a data sample of 172 quarters. We show that the estimates are quite different, and it is important to be aware of the effect of using alternative priors.

6.2 Estimating a Vector Autoregressive Model of the U.S. Economy

In the past two decades, Bayesian VAR models have been commonly used for analyzing multivariate time series macroeconomic data and addressing policy questions. However, not much attention has been given to the sensitivity of results with respect to researchers' choice of estimator and prior. In what follows we compare various Bayesian estimates of a six-variable VAR using quarterly data of the U.S. economy from 1959 Q1 to 2001 Q4. We set the lag length of the VAR at 2, in accordance with the Schwarz criterion. The variables are real GDP, GDP deflator, world commodity price, the Federal Funds rates, nonborrowed reserves, and M2 money stock. The commodity price data were obtained from the International Monetary Fund, and the rest of the data were obtained from the FRED database at the Federal Reserve Bank of St Louis. All variables except the Federal Funds rates are in logarithms. These variables frequently appear in macroeconomics-related VARs (e.g., Sims 1992; Gordon and Leeper 1994; Sims and Zha 1998; Christiano et al. 1999). We focus on the role of priors by limiting our attention to the normal model. The Gibbs sampling algorithms for the posteriors under various priors in this application are based on the conditional posteriors given in Section 5.1. The number of MCMC cycles is set at 10,000.

6.2.1 Bayesian Estimates and Posterior Risks. As in our numerical example, under the constant prior the posterior mean of Φ is very close to the MLE but quite different from that under the shrinkage prior. To be specific, the most prominent differences of the two estimates under alternative priors on Φ are in the third and the ninth rows of Φ , which correspond to the first and second lag parameters of the GDP deflator variable. Most elements of the third row of $\hat{\Phi}$ under the constant prior are similar in magnitude to the elements in the ninth row of the same column but have the opposite signs. The pattern of estimates also emerge in the GDP variable (of the second and eighth rows), the nonborrowed reserves variable (of the sixth and the twelfth rows), and the M2 money stock variable (of the seventh and thirteenth rows). The estimates of this VAR suggest some degree of collinearity, which is not uncommon for macroeconomic applications of VARs because macroeconomic time series data often exhibit strong serial and pairwise correlations. Moreover, the VAR models are "overparameterized," with no restrictions on the matrix Φ . In empirical applications such as the present one, the MLE estimates of the first and second lag coefficients are not only of similar magnitude and opposite signs, but also are often found to be very sensitive to model

specification and sample period. The fact that the estimates of Φ under the constant prior are similar to the MLEs suggests the possibility of improvement by using alternative priors in place of the constant prior.

As discussed earlier, the shrinkage estimator of James and Stein (1961), which motivates the shrinkage prior, reduces quadratic frequentist loss in estimating a multivariate normal mean. The James–Stein estimator is also known for improving efficiency in the presence of multicollinearity in regression models. Assessing the improvement of shrinkage prior-based estimates requires computation of losses. In real applications, frequentist risks cannot be calculated, because the true parameters are unknown. Theoretical results on the admissibility of estimates under alternative priors have not been established in the VAR framework, but most likely the use of a shrinkage prior for Φ improves upon the MLE and Bayesian estimates based on the constant prior.

Unlike the frequentist risk, the posterior risk for a given dataset can be computed using the MCMC simulation output. For some loss functions, this can be done at little additional cost. For example, with the posterior mean estimator $\hat{\Sigma} = E(\Sigma|Y)$, the posterior loss of $L_{\Sigma_1}(\hat{\Sigma}, \Sigma)$ is

$$E[\{\text{trace}(\hat{\Sigma}^{-1}\Sigma) - \log|\hat{\Sigma}^{-1}\Sigma| - p\}|Y] = \log|\hat{\Sigma}| - E\{(\log|\Sigma|)|Y\}. \quad (31)$$

To compute $E\{(\log|\Sigma|)|Y\}$, we decompose the Σ_k matrix in the k th MCMC cycle as $\Sigma_k = \mathbf{O}_k \mathbf{D}_k \mathbf{O}'_k$, where \mathbf{D}_k is the diagonal matrix comprising eigenvalues of Σ_k ,—that is, $\mathbf{D}_k = \text{diag}(d_{k1}, d_{k2}, \dots, d_{kp})$ —and \mathbf{O}_k is an orthogonal matrix with $\mathbf{O}_k \mathbf{O}'_k = \mathbf{I}$. It follows that $\hat{E}\{(\log|\Sigma|)|Y\} = \frac{1}{M} \sum_{k=1}^M \times$

$\sum_{i=1}^p \log|d_{ki}|$, which can be computed in MCMC runs. The posterior risk of the estimator $\hat{\Phi} = E(\Phi|Y)$ associated with the loss $L_{\Phi_1}(\hat{\Phi}, \Phi)$ is

$$E[\text{trace}\{(\hat{\Phi} - \Phi)'(\hat{\Phi} - \Phi)\}|Y] = \text{trace}[E(\Phi'\Phi|Y) - \{E(\Phi|Y)\}'\{E(\Phi|Y)\}]. \quad (32)$$

Both $E(\Phi|Y)$ and $E(\Phi'\Phi|Y)$ can be computed in the process of MCMC simulations without storage of the entire MCMC output.

The posterior risks of the posterior mean estimator under alternative priors turn out to be quite different. For instance, under the constant RATS prior the posterior risks (31) and (32) are .126 and 13.703. Under the shrinkage reference prior combination, the corresponding posterior losses are .123 and 6.470. The posterior losses under the shrinkage reference prior are smaller because the posterior distributions are tighter for the distribution of the VAR regression coefficient matrix Φ .

6.2.2 Impulse Responses. We now compare impulse responses under the constant RATS and shrinkage reference priors. To save space, we plot only responses of GDP to shocks in the GDP deflator. We compare the GDP response with an inflation shock under competing priors. As noted earlier, under the constant prior, the third and ninth rows of $\hat{\Phi}_{1CA}$ show similar magnitude but opposite signs. For example, the posterior means of elements $\phi_{(3,1)}$ and $\phi_{(9,1)}$ are 3.057 and -3.714 . In contrast, under the shrinkage prior, the posterior means of elements $\phi_{(3,1)}$ and $\phi_{(9,1)}$ are .358 and -1.134 . The difference in the Bayesian estimates $\hat{\Phi}_{1CA}$ and $\hat{\Phi}_{1SR}$ corresponds to quite different impulse responses. Figure 1 shows that under the con-

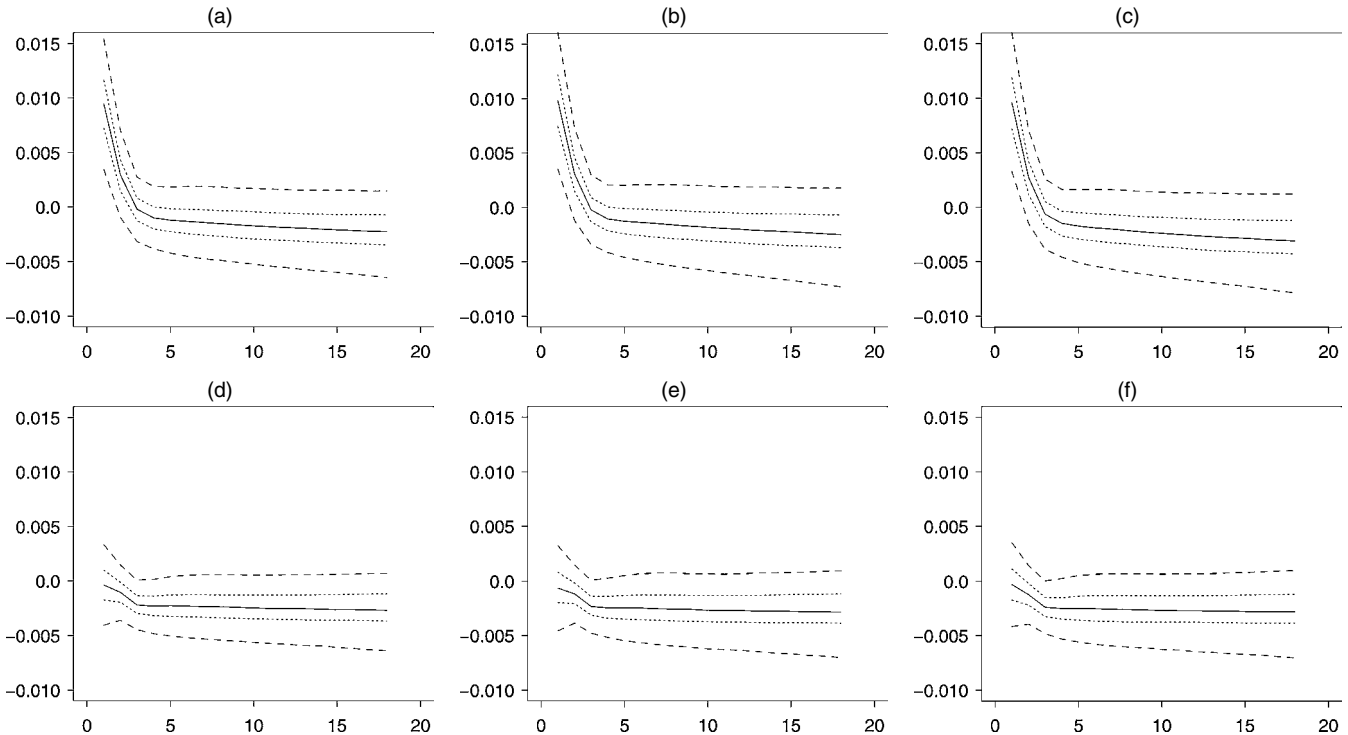


Figure 1. Responses of GDP to an Inflation Shock. The dashed line, represent the 10th or 90th percentile; dotted line, the 32th or 68th percentile; solid line, the posterior mean of impulse responses. The panels are GDP responses (a) under the constant RATS prior; (b) under the constant Jeffreys prior; (c) under the constant reference prior; (d) under the shrinkage RATS prior; (e) under the shrinkage Jeffreys prior; and (f) under the shrinkage reference prior.

stant prior, after a one-unit inflation shock, GDP initially surges by as much as 1% and then quickly drops. The long-run effect is about $-.2\%$. Given the fact that price level increases after an inflation shock, the shock may be interpreted as a positive shock in demand. Under the shrinkage prior, a distinctly different pattern emerges. GDP does not change by much immediately after the inflation shock and drifts slowly below 0. Under the shrinkage prior, an inflation shock plays the role of a negative supply shock (with the consequence of increases in price and decreases in output). Obviously, using different priors on the same dataset leads to quite different conclusions.

7. CONCLUDING REMARKS

In this article we compare frequentist risks of several Bayesian estimates of the VAR regression coefficients Φ and covariance matrix Σ under competing priors and data distributions. The asymmetric LINEX estimator Φ does better overall than the posterior mean. We do not find an estimator for Σ dominating in all cases. It is therefore well advised to compare all estimates of Σ in applications. The choice of prior has stronger effect on the Bayesian estimates than the choice of loss function. The shrinkage prior on Φ dominates the constant prior in the numerical example. Yang and Berger's reference prior on Σ dominates the Jeffreys prior and the RATS prior. These conclusions also hold for the Student- t model as well as the normal model. Estimation of a VAR using U.S. macroeconomic data reveals significant difference between estimates of VAR regression coefficients under the shrinkage and constant priors. Similarly, impulse responses of GDP to an inflation shock are shown to be distinctly different under the competing priors.

This study may be extended in several directions. First, the list of noninformative priors that we examine in this article is by no means exhaustive. Other noninformative priors applicable to the VAR framework need to be explored. Second, in this article we derive Bayesian estimates from loss functions separable in Φ and Σ . As a result, inferences for Φ and Σ are largely independent. Our future research concerns joint Bayesian inference of Φ and Σ based on an intrinsic joint loss function, such as the entropy loss. Third, we consider only linear VAR models. Recently, applications of nonlinear VAR models (e.g., Altissimo and Giovanni 2001) have become more common. Extending our analysis to nonlinear Bayesian VARs will be an interesting future research topic.

ACKNOWLEDGMENTS

This research was supported by a Research Board grant from the University of Missouri System. Sun's research is also supported by National Science Foundation grants DMS-99-72598 and SES-0095919 and a grant from the Missouri Department of Conservation. The authors thank two referees and an associate editor for their valuable comments. They also thank Professor Arnold Zellner and the referees for suggesting the simulation study on stationarity of the estimated VARs. They are especially grateful to Professor Paul L. Speckman for his careful proofreading and helpful comments.

APPENDIX: SAMPLING FROM THE POSTERIOR OF (Φ, Σ) UNDER STUDENT- t ERRORS

The algorithms for MCMC computations of posterior distributions of (Φ, Σ) depend on the priors. For brevity, we outline only the algorithms with the constant prior on Φ and the Jeffreys and reference priors on Σ , and we outline the algorithms only with Student- t errors, which have one more step than the algorithms with the normal errors. We take the degrees of freedom of the Student- t distribution, ν , as unknown. Following the discussion after Fact 5, we use an MCMC algorithm to sample from the joint posterior distribution $(\Phi, \Sigma, \mathbf{Q}, \nu)$. The algorithm used for simulating the posterior under the constant Jeffreys prior is as follows:

Algorithm CJT. Suppose that at cycle k , we have $(\Phi_{k-1}, \Sigma_{k-1}, \mathbf{Q}_{k-1}, \nu_{k-1})$ sampled from cycle $k-1$:

- Step 1. Simulate $q_{tk} \sim \text{gamma}(\frac{1}{2}(\nu_{k-1} + p), \frac{1}{2}\{\nu_{k-1} + (\mathbf{y}_t - \mathbf{x}_t \Phi_{k-1})' \Sigma_{k-1}^{-1} (\mathbf{y}_t - \mathbf{x}_t \Phi_{k-1})\})$.
- Step 2. Calculate $\hat{\Phi}_k = (\mathbf{X}' \mathbf{Q}_k \mathbf{X})^{-1} \mathbf{X}' \mathbf{Q}_k \mathbf{Y}$, where $\mathbf{Q}_k = \text{diag}(q_{1k}, \dots, q_{Tk})$.
- Step 3. Simulate $\Sigma_k \sim \text{IW}(\mathbf{S}_k(\hat{\Phi}_k), T - Lp - 1)$, where $\mathbf{S}_k(\Phi) = (\mathbf{Y} - \mathbf{X}\Phi)' \mathbf{Q}_k (\mathbf{Y} - \mathbf{X}\Phi)$.
- Step 4. Simulate $\phi_k \sim \text{N}(\hat{\Phi}_k, \Sigma_k \otimes (\mathbf{X}' \mathbf{Q}_k \mathbf{X})^{-1})$.
- Step 5. Simulate w conditional on ϕ_k, Σ_k , and \mathbf{Q}_k from distribution (26) using the adaptive rejection sampling scheme of Gilks and Wild (1992), and let $\nu_k = 2w$.

The algorithm using the constant RATS prior is similar to the foregoing, with the exception that in step 3 the distribution of the inverse Wishart has different degrees of freedom: $\Sigma_k \sim \text{IW}(\mathbf{S}_k(\hat{\Phi}_k), T)$.

It is much more difficult to simulate from the conditional distribution of Σ under the reference prior. We adopt the hit-and-run algorithm used by Yang and Berger (1994). In implementing the algorithm, we use a one-to-one transformation of Σ , namely $\Sigma^* = \log(\Sigma)$ or $\Sigma = \exp(\Sigma^*)$ in the sense that $\Sigma = \sum_{j=0}^{\infty} \Sigma^{*j}/j!$. Then the conditional posterior density of Σ^* given $(\mathbf{Q}, \phi, \mathbf{Y}, \nu)$ is

$$\begin{aligned} \pi(\Sigma^* | \mathbf{Q}, \phi, \nu, \mathbf{Y}) &= \pi(\Sigma^* | \mathbf{S}_k(\phi)) \\ &\propto \frac{\text{etr}\{-T/2\Lambda^* - 1/2(\exp \Sigma^*)^{-1} \mathbf{S}_k(\phi)\}}{\prod_{i < j} (\lambda_i^* - \lambda_j^*)}, \quad (\text{A.1}) \end{aligned}$$

where $\Sigma^* = \mathbf{O}' \Lambda^* \mathbf{O}$, \mathbf{O} is an orthogonal matrix, and $\Lambda^* = \text{diag}(\lambda_1^*, \dots, \lambda_p^*)$ with $\lambda_1^* > \dots > \lambda_p^*$. Note that $\exp(\Sigma^*) = \mathbf{O}' \exp(\Lambda^*) \mathbf{O}$. To simulate Σ^* from (A.1), we use the following algorithm:

Algorithm CRT. Suppose that at cycle k , we have $(\Phi_{k-1}, \Sigma_{k-1}, \mathbf{Q}_{k-1}, \nu_{k-1})$ sampled from cycle $k-1$:

- Step 1. Same as step 1 of Algorithm CJT.
- Step 2. Same as step 2 of Algorithm CJT.
- Step 3. Simulate $\phi_k \sim \text{N}(\hat{\Phi}_k, \Sigma_{k-1} \otimes (\mathbf{X}' \mathbf{Q}_k \mathbf{X})^{-1})$.
- Step 4. Decompose $\Sigma_{k-1} = \mathbf{O}' \Lambda \mathbf{O}$, where $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_p)$, $\lambda_1 > \lambda_2 > \dots > \lambda_p$, and $\mathbf{O}' \mathbf{O} = \mathbf{I}$. Let $\lambda_i^* = \log(\lambda_i)$, $\Lambda^* = \text{diag}(\lambda_1^*, \dots, \lambda_p^*)$, and $\Sigma_{k-1}^* = \mathbf{O} \Lambda^* \mathbf{O}'$.

Step 5. Select a random symmetric $p \times p$ matrix \mathbf{V} , with elements $v_{ij} = z_{ij}/\sqrt{\sum_{l \leq m} z_{lm}^2}$, where $z_{ij} \sim N(0, 1)$, $1 \leq i \leq j \leq p$. The other elements of \mathbf{V} are defined by symmetry.

Step 6. Generate $\rho \sim N(0, 1)$ and set $\mathbf{W} = \Sigma_{k-1}^* + \rho\mathbf{V}$. Decompose $\mathbf{W} = \mathbf{H}'\mathbf{C}^*\mathbf{H}$, where $\mathbf{C}^* = \text{diag}(c_1^*, \dots, c_p^*)$, $c_1^* > c_2^* > \dots > c_p^*$, and $\mathbf{H}'\mathbf{H} = \mathbf{I}$. Calculate $\mathbf{S}_k(\Phi_k) = \mathbf{S}_k\mathbf{Q} = (\mathbf{Y} - \mathbf{X}\Phi_k)'\mathbf{Q}_k(\mathbf{Y} - \mathbf{X}\Phi_k)$. Compute

$$\begin{aligned} \alpha_k &= \log(\pi(\exp(\mathbf{W})|\mathbf{S}_k\mathbf{Q})) \\ &\quad - \log(\pi(\exp(\Sigma_{k-1}^*)|\mathbf{S}_k\mathbf{Q})) \\ &= \frac{T}{2} \sum_{i=1}^p (\lambda_i^* - c_i^*) \\ &\quad + \frac{1}{2} \text{trace}\{((\exp \Sigma_{k-1}^*)^{-1} - (\exp \mathbf{W})^{-1})\mathbf{S}_k\mathbf{Q}\} \\ &\quad + \sum_{i < j} \log(\lambda_i^* - \lambda_j^*) - \sum_{i < j} \log(c_i^* - c_j^*). \end{aligned}$$

Step 7. Generate $u \sim \text{unif}(0, 1)$. If $u \leq \min(1, \exp(\alpha_k))$, let $\Sigma_k^* = \mathbf{W}$ and $\Sigma_k = \mathbf{HCH}'$, where $\mathbf{C} = \text{diag}(e^{c_1}, \dots, e^{c_p})$; otherwise, let $\Sigma_k^* = \Sigma_{k-1}^*$ and $\Sigma_k = \Sigma_{k-1}$.

Step 8. Simulate w conditional on ϕ_k , Σ_k , and \mathbf{Q}_k from distribution (26) using the adaptive rejection sampling scheme of Gilks and Wild (1992), and let $v_k = 2w$.

When the shrinkage prior is used to replace the constant prior for ϕ , the algorithms for Bayesian computation need to be modified by adding one step for drawing ϕ using Fact 8. In cycle k , ϕ_k is drawn in two steps. First, parameter δ_k is drawn from an inverse gamma distribution, which depends on ϕ_{k-1} . Then ϕ_k is drawn from a multivariate normal distribution that depends on δ_k , Σ_k , and the data sample.

Algorithms for simulating the posterior of (Φ, Σ) with the normal errors are similar to the algorithms presented earlier, except that in the normal model the steps for simulating \mathbf{Q} and v are omitted and \mathbf{Q} is fixed at \mathbf{I}_T .

[Received March 2002. Revised January 2004.]

REFERENCES

- Altissimo, F., and Giovanni, L. V. (2001), "The Non-Linear Dynamics of Output and Unemployment in the U.S.," *Journal of Applied Econometrics*, 16, 461–486.
- Anderson, T. W. (1984), *An Introduction to Multivariate Statistical Analysis* (2nd ed.), New York: Wiley.
- Baranchik, A. J. (1964), "Multiple Regression and Estimation of the Mean of Multivariate Normal Distribution," Technical Report 51, Stanford University, Dept. of Statistics.
- Barberis, N. (2000), "Investing for the Long Run When the Returns Are Predictable," *Journal of Finance*, 55, 225–264.
- Berger, J. O., and Bernardo, J. M. (1992), "On the Development of Reference Priors," in *Bayesian Analysis IV*, eds. J. M. Bernardo et al., Oxford, U.K.: Oxford University Press, pp. 61–77.
- Berger, J. O., and Strawderman, W. E. (1996), "Choice of Hierarchical Priors: Admissibility in Estimation of Normal Means," *The Annals of Statistics*, 24, 931–951.
- Bernardo, J. M. (1979), "Reference Posterior Distributions for Bayesian Inference," *Journal of the Royal Statistical Society*, Ser. B, 41, 113–147.
- Bowman, K. O., and Shenton, L. R. (1988), *Properties of Estimators for the Gamma Distribution*, New York: Marcel Dekker.
- Chib, S., Osiewalski, J., and Steel, M. F. J. (1991), "Posterior Inference on the Degrees of Freedom Parameter in Multivariate- t Regression Models," *Economics Letters*, 37, 391–397.
- Christiano, L. J., Eichenbaum, M., and Evans C. (1999), "Monetary Policy Shocks: What Have We Learned and to What End?" in *Handbook of Macroeconomics*, Vol. 1, eds. J. B. Taylor and M. Woodford, Amsterdam: North-Holland, pp. 65–147.
- Doan, T., Litterman, R. B., and Sims, C. A. (1984), "Forecasting and Conditional Projection Using Realistic Prior Distributions," *Econometric Reviews*, 3, 1–100.
- Fernandez, C., Osiewalski, J., and Steel, M. F. J. (1997), "Classical and Bayesian inference Robustness in Multivariate Regression Models," *Journal of American Statistical Association*, 92, 1434–1444.
- Fernandez, C., and Steel, M. F. J. (1999), "Multivariate Student- t Regression Models: Pitfalls and Inference," *Biometrika*, 86, 153–167.
- Geisser, S. (1965), "Bayesian Estimation in Multivariate Analysis," *The Annals of Mathematical Statistics*, 36, 150–159.
- Gelfand, A. E., and Smith, A. F. M. (1990), "Sampling-Based Approaches to Calculating Marginal Densities," *Journal of the American Statistical Association*, 85, 398–409.
- Geweke, J. (1993), "Bayesian Treatment of the Independent Student- t Linear Model," *Journal of Applied Econometrics*, 8, 19–40.
- Gilks, W. R., and Wild, P. (1992), "Adaptive Rejection Sampling for Gibbs Sampling," *Applied Statistics*, 41, 337–348.
- Gordon, D. B., and Leeper, E. M. (1994), "The Dynamic Impacts of Monetary Policy: An Exercise in Tentative Identification," *Journal of Political Economy*, 102, 1228–1247.
- Hollifield, B., Koop, G., and Li, K. (2001), "A Bayesian Analysis of a Variance Decomposition of Stock Returns," working paper, Carnegie Mellon University.
- James, W., and Stein, C. (1961), "Estimation With Quadratic Loss," in *Proceedings of the Fourth Berkeley Symposium on Mathematics, Statistics, and Probability*, Vol. 1, Berkeley, CA: University of California Press, pp. 361–390.
- Jeffreys, H. (1961) *Probability Theory*, New York: Oxford University Press.
- Johnson, N. L., and Kotz, S. (1972), *Distributions in Statistics: Continuous Multivariate Distributions*, New York: Wiley.
- Kadiyala, K. R., and Karlsson, S. (1997), "Numerical Methods for Estimation and Inference in Bayesian VAR-Models," *Journal of Applied Econometrics*, 12, 99–132.
- Kandel, S., McCulloch R., and Stambaugh, R. F. (1995), "Bayesian Inference and Portfolio Efficiency," *Review of Financial Studies*, 8, 1–53.
- Kandel, S., and Stambaugh, R. F. (1996), "On the Predictability of Stock Returns," *Journal of Finance*, 51, 385–424.
- Kass, R. E., and Wasserman, L. (1996), "The Selection of Prior Distributions by Formal Rules," *Journal of American Statistical Association*, 91, 1343–1370.
- Lamoureux, C. G., and Zhou, G. (1996), "Temporary Components of Stock Returns: What Do the Data Tell Us?" *Review of Financial Studies*, 9, 1033–1059.
- Litterman, R. B. (1986), "Forecasting With Bayesian Vector Autoregression: Five Years of Experience," *Journal of Business & Economic Statistics*, 4, 25–38.
- MacKinnon, J. G., and Smith, A. A., Jr. (1998), "Approximate Bias Correction in Econometrics," *Journal of Econometrics*, 85, 205–230.
- Ni, S., and Sun, D. (2003), "Noninformative Priors and Frequentist Risks of Bayesian Estimators of Vector-Autoregressive Models," *Journal of Econometrics*, 115, 159–197.
- Pastor, L. (2000), "Portfolio Selection and Asset Pricing Models," *Journal of Finance*, 55, 179–223.
- Pastor, L., and Stambaugh, R. F. (2000), "Comparing Asset Pricing Models: An Investment Perspective," *Journal of Financial Economics*, 56, 335–381.
- Osiewalski, J., and Steel, M. F. J. (1993), "Robust Bayesian Inference in Elliptical Regression Models," *Journal of Econometrics*, 57, 345–363.
- Phillips, P. C. B. (1991), "To Criticize the Critics: An Objective Bayesian Analysis of Stochastic Trends," *Journal of Applied Econometrics*, 6, 423–434.
- Sims, C. A. (1980), "Macroeconomics and Reality," *Econometrica*, 48, 1–48.
- (1991), Comment on "To Criticize the Critics: An Objective Bayesian Analysis of Stochastic Trends," by Peter C. B. Phillips, *Journal of Applied Econometrics*, 6, 423–434.
- (1992), "Interpreting the Macroeconomic Time Series Facts: The Effects of Monetary Policy," *European Economic Review*, 38, 975–1000.
- Sims, C. A., and Zha, T. (1998), "Does Monetary Policy Generate Recessions?" Working Paper 98-12, Federal Reserve Bank of Atlanta.
- (1999), "Error Bands for Impulse Responses," *Econometrica*, 67, 1113–1155.

- Stambaugh, R. F. (1999), "Predictive Regressions," *Journal of Financial Economics*, 54, 375–421.
- Stein, C. (1956), "Inadmissibility of the Usual Estimator for the Mean of a Multivariate Normal Distribution," in *Proceedings of the Third Berkeley Symposium*, Vol. 1, Berkeley, CA: University of California Press, pp. 197–206.
- Sun, D., and Berger, J. O. (1998), "Reference Priors Under Partial Information," *Biometrika*, 85, 55–71.
- Sun, D., and Ni, S. (2004), "Bayesian Analysis of Vector-Autoregressive Models With Noninformative Priors," *Journal of Statistical Planning and Inference*, 121, 291–309.
- Tiao, G. C., and Zellner, A. (1964), "On the Bayesian Estimation Analysis of Multivariate Regression," *Journal of Royal Statistical Society, Ser. B*, 26, 389–399.
- Yang, R., and Berger, J. O. (1994), "Estimation of a Covariance Matrix Using the Reference Prior," *The Annals of Statistics*, 22, 1195–1211.
- Zellner, A. (1971), *An Introduction to Bayesian Inference in Econometrics*, New York: Wiley.
- (1976), "Bayesian and Non-Bayesian Analysis of the Regression Model With Multivariate Student-*t* Error Terms," *Journal of the American Statistical Association*, 71, 400–405.
- (1986), "Bayesian Estimation and Prediction Using Asymmetric Loss Functions," *Journal of the American Statistical Association*, 81, 446–451.
- Zellner, A., and Hong, C. (1989), "Forecasting International Growth Rates Using Bayesian Shrinkage and Other Procedures," *Journal of Econometrics*, 40, 183–202.