Centre for Central Banking Studies

Technical Handbook – No. 4 Applied Bayesian econometrics for central bankers

Andrew Blake and Haroon Mumtaz

BANK OF ENGLAND

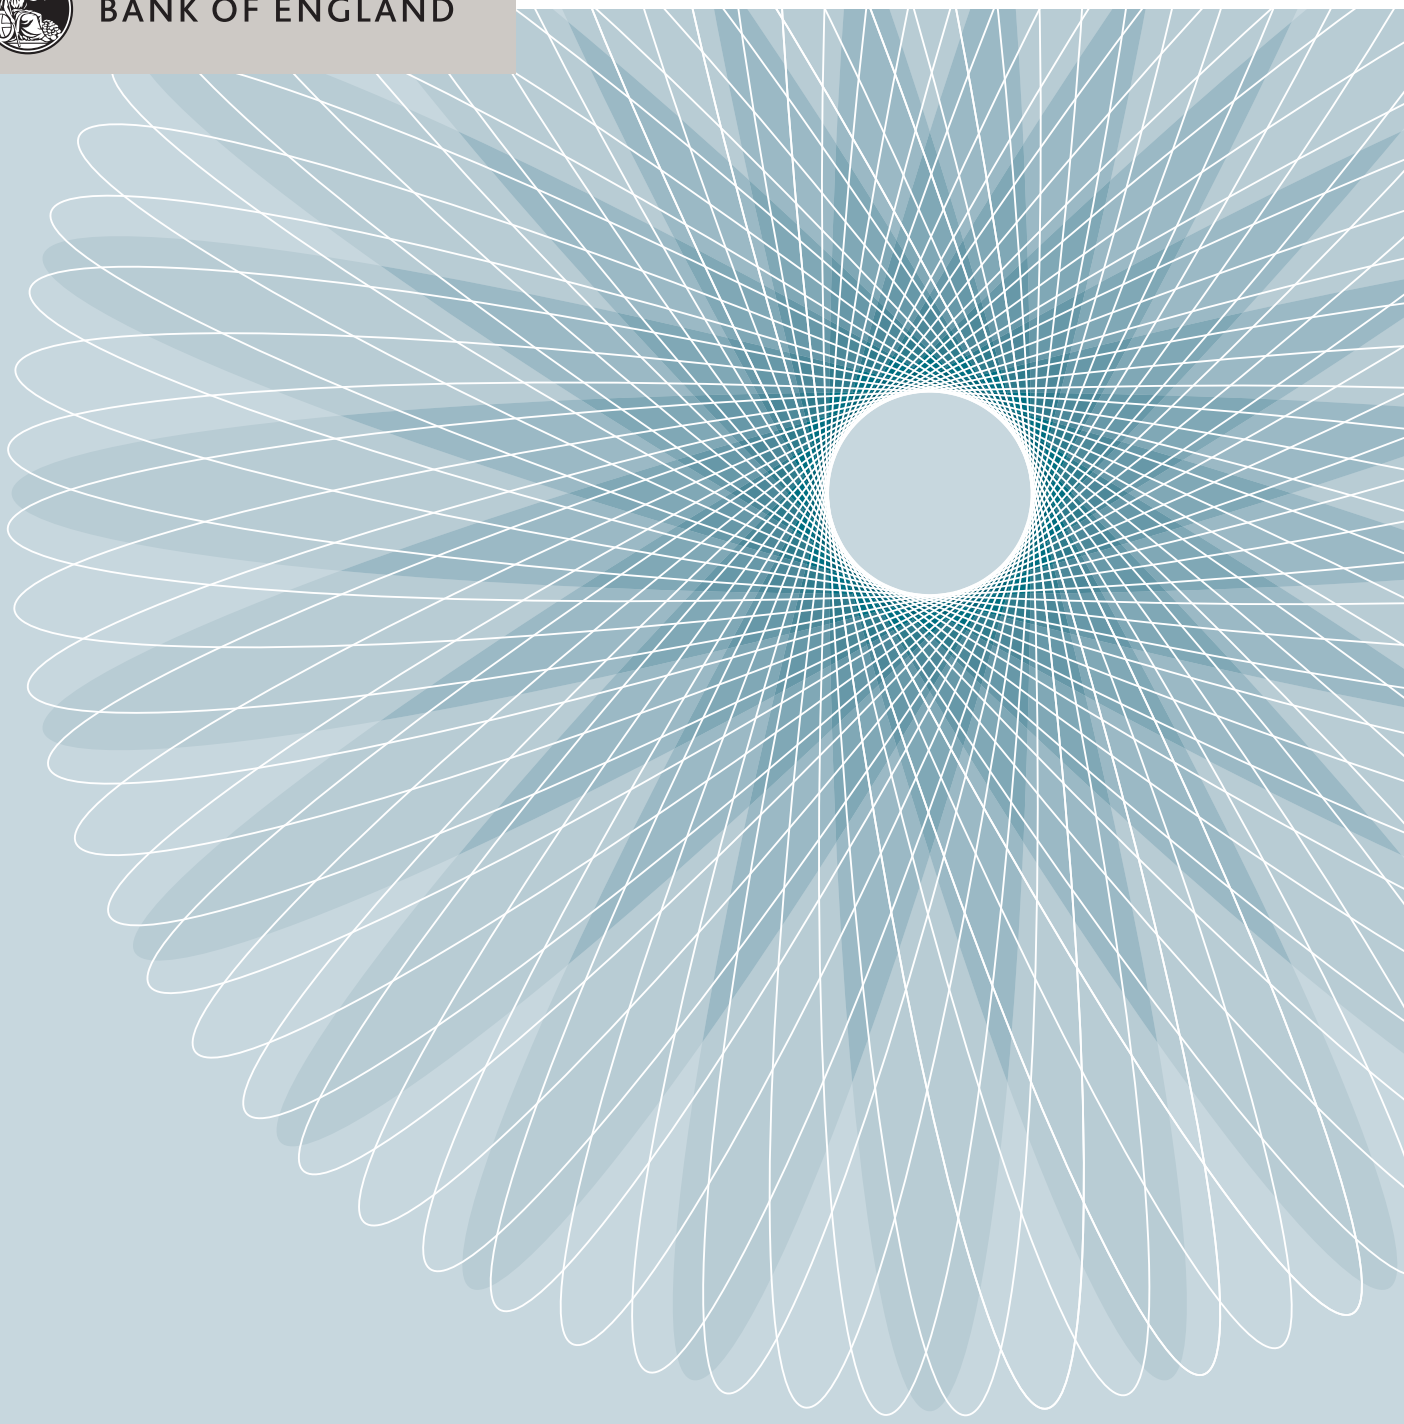

BANK OF ENGLAND

CCBS Technical Handbook No. 4

Applied Bayesian econometrics for central bankers

Andrew Blake and Haroon Mumtaz[1]

[1] Andrew Blake, Bank of England, Threadneedle Street, London, EC2R 8AH.
Email Andrew.blake@bankofengland.co.uk
Haroon Mumtaz, Bank of England, Threadneedle Street, London, EC2R 8AH.
Email Haroon.mumtaz@bankofengland.co.uk

Series Editor: Andrew Blake: Andrew.blake@bankofenlgand.co.uk

This Research Paper was finalized in September 2012

Centre for Central Banking Studies, Bank of England, Threadneedle Street, London, EC2R 9AH
Email: ccbsinfo@bankofengland.co.uk

ISSN (online) 1756-7297

Foreword

The series of *Handbooks in Central Banking* form a key part of the activities of the Centre for Central Banking Studies (CCBS) at the Bank of England. The CCBS has been in existence since 1990, delivering seminars, workshops and expert advice to central banks all over the world. The *Handbooks* cover topics that concern the technical and analytical aspect of central banking.

The *Handbooks* are aimed primarily at central bankers, and have proced extremely popular and useful reference works for all those looking for materials that provide both a clear analytical framework together with the practical application of these ideas.

Most of the CCBS *Handbooks* are available from our website www.bankofengland.co.uk/education/ccbs/handbooks_lectures.htm. Several have been translated into Spanish, Russian and Arabic, and these versions are also available on the website.

Our aim is to continue to add to the series, covering new areas of interest and also updating existing *Handbooks* to take account of recent developments. Some of the new *Technical Handbooks* include econometric exercises developed in our workshops, thus making these available to a wider audience.

We hope you find the new additions to the series useful, and would welcome any comments on the *Handbooks* and any suggestions for future topics.

We should note that all views expressed are those of the authors and not necessarily those of the Bank of England or Monetary Policy Committee members.

Andrew Blake
Series Editor

Applied Bayesian Econometrics for Central Bankers

Andrew Blake

Haroon Mumtaz

CentRE For Central Banking Studies, Bank of England
E-mail address: Andrew.Blake@bankofengland.co.uk

CentRE For Central Banking Studies, Bank of England
E-mail address: Haroon.Mumtaz@bankofengland.co.uk

Preface

This aim of this handbook is to introduce key topics in Bayesian econometrics from an applied perspective.

The handbook assumes that readers have a fair grasp of basic classical econometrics (e.g. maximum likelihood estimation). It is recommended that readers familiarise themselves with the Matlab© programming language to derive the maximum benefit from this handbook. A basic guide to Matlab© is attached at the end of the Handbook.

The first chapter of the handbook introduces basic concepts of Bayesian analysis. In particular, the chapter focuses on the technique of Gibbs sampling and applies it to a linear regression model. The chapter shows how to code this algorithm via several practical examples. The second chapter introduces Bayesian vector autoregressions (VARs) and discusses how Gibbs sampling can be used for these models. The third chapter shows how Gibbs sampling can be applied to popular econometric models such as time-varying VARS and dynamic factor models. The final chapter introduces the Metropolis Hastings algorithm. We intend to introduce new topics in revised versions of this handbook on a regular basis.

The handbook comes with a set of Matlab© codes that can be used to replicate the examples in each chapter. The code (provided in code.zip) is organised by chapter. For example, the folder 'Chapter1' contains all the examples referred to in the first chapter of this handbook.

The views expressed in this handbook are those of the authors, and not necessarily those of the Bank of England. The reference material and computer codes are provided without any guarantee of accuracy. The authors would appreciate feedback on possible coding errors and/or typos.

Contents

Part 1

A Practical Introduction to Gibbs Sampling

CHAPTER 1

Gibbs Sampling for Linear Regression Models

1. Introduction

This chapter provides an introduction to the technique of estimating linear regression models using Gibbs sampling. While the linear regression model is a particularly simple case, the application of Gibbs sampling in this scenario follows the same principles as the implementation in a more complicated models (considered in later chapters) and thus serves as a useful starting point. We draw heavily on the seminal treatment of this topic in Kim and Nelson (1999). A more formal (but equally accessible) reference is Koop (2003).

The reader should aim to become familiar with the following with the help of this chapter

- The prior distribution, the posterior distribution and Bayes Theorem.
- Bayesian treatment of the linear regression model.
- Why Gibbs sampling provides a convenient estimation method.
- Coding the Gibbs sampling algorithm for a linear regression in Matlab

2. A Bayesian approach to estimating a linear regression model

Consider the task of estimating the following regression model

$$\begin{aligned} Y_t &= BX_t + v_t \\ v_t &\sim N\left(0, \sigma^2\right) \end{aligned} \tag{2.1}$$

where Y_t is a $T \times 1$ matrix of the dependent variable, Y_t is a $T \times K$ matrix of the independent variables and deterministic terms. We are concerned with estimating the $K \times 1$ vector of coefficients B and the variance of the error term σ^2.

A classical econometrician proceeds by obtaining data on Y_t and X_t and writes down the likelihood function of the model

$$F\left(Y_t \backslash B, \sigma^2\right) = \left(2\pi\sigma^2\right)^{-T/2} \exp\left(-\frac{\left(Y_t - BX_t\right)'\left(Y_t - BX_t\right)}{2\sigma^2}\right) \tag{2.2}$$

and obtains estimates $\hat{B}$ and $\hat{\sigma}^2$ by maximising the likelihood. In this simple case these deliver the familiar OLS estimator for the coefficients $\hat{B}_{OLS} = \left(X_t'X_t\right)^{-1}\left(X_t'Y_t\right)$ and the (biased) maximum likelihood estimator for the error variance $\hat{\sigma}^2 = \frac{v_t'v_t}{T}$. For our purpose, the main noteworthy feature of the classical approach is the fact that the estimates of the parameters of the model are solely based on information contained in data.

Bayesian analysis departs from this approach by allowing the researcher to incorporate her prior beliefs about the parameters B and σ^2 into the estimation process. To be exact, the Bayesian econometrician when faced with the task of estimating equation 2.1 would proceed in the following steps.

Step 1. The researcher forms a prior belief about the parameters to be estimated. This prior belief usually represents information that the researcher has about B and σ^2 which is not derived using the data Y_t and X_t. These prior beliefs may have been formed through past experience or by examining studies (estimating similar models) using other datasets. (We will discuss the merits of this approach for specific examples in the chapters below). The key point is that these beliefs are expressed in the form of a probability distribution. For example, the prior on the coefficients B is expressed as

$$P\left(B\right) \sim N\left(B_0, \Sigma_0\right) \tag{2.3}$$

where the mean B_0 represents the actual beliefs about the elements of B.

EXAMPLE 1. *In the case of two explanatory variables, the vector* $B_0 = \begin{pmatrix} 1 \\ -1 \end{pmatrix}$ *represents the belief that the first coefficient equals* 1 *and the second equals* -1. *The variance of the prior distribution* Σ_0 *controls how strong this prior belief is. A large number for* Σ_0 *would imply that the researcher is unsure about the numbers she has chosen for* B_0 *and wants to place only a small weight on them. In contrast, a very small number for* Σ_0 *implies that the researcher is very sure about the belief expressed in* B_0.*In the case of two explanatory variables* Σ_0 *may equal* $\Sigma_0 = \begin{pmatrix} 10 & 0 \\ 0 & 10 \end{pmatrix}$ *representing a 'loose prior' or an uncertain prior belief.*

Step 2. The researcher collects data on Y_t and X_t and write down the likelihood function of the model

$$F\left(Y_t \backslash B, \sigma^2\right) = \left(2\pi\sigma^2\right)^{-T/2} \exp\left(-\frac{\left(Y_t - BX_t\right)'\left(Y_t - BX_t\right)}{2\sigma^2}\right)$$

This step is identical to the approach of the classical econometrican and represents the information about the model parameters contained in the data.

Step 3. The researcher updates her prior belief on the model parameters (formed in step 1) based on the information contained in the data (using the likelihood function in step 2). In other words, the researcher combines the prior distribution $P\left(B, \sigma^2\right)$ and the likelihood function $F\left(Y_t \backslash B, \sigma^2\right)$ to obtain the *posterior* distribution $H\left(B, \sigma^2 \backslash Y_t\right)$.

More formally, the Bayesian econometrician is interested in the posterior distribution $H\left(B, \sigma^2 \backslash Y_t\right)$ which is defined by the *Bayes Law*

$$H\left(B, \sigma^2 \backslash Y_t\right) = \frac{F\left(Y_t \backslash B, \sigma^2\right) \times P\left(B, \sigma^2\right)}{F(Y)} \tag{2.4}$$

Equation 2.4 simply states that the posterior distribution is a product of the likelihood $F\left(Y_t \backslash B, \sigma^2\right)$ and the prior $P\left(B, \sigma^2\right)$ divided by the density of the data $F(Y)$ (also referred to as the marginal likelihood or the marginal data density). Note that $F(Y)$ is a scalar and will not have any operational significance as far as estimation is concerned (although it is crucial for model comparison, a topic we return to). Therefore the Bayes Law can be written as

$$H\left(B, \sigma^2 \backslash Y_t\right) \propto F\left(Y_t \backslash B, \sigma^2\right) \times P\left(B, \sigma^2\right) \tag{2.5}$$

Equation 2.5 states that the posterior distribution is proportional to the likelihood times the prior. In practice, we will consider equation 2.5 when considering the estimation of the linear regression model.

As an aside note that the Bayes law in equation 2.4 can be easily derived by considering the joint density of the data Y_t and parameters B, σ^2, $G\left(Y_t, B, \sigma^2\right)$ and observing that ir can be factored in two ways

$$G\left(Y_t, B, \sigma^2\right) = F(Y_t) \times H(B, \sigma^2 \backslash Y_t) = F(Y_t \backslash B, \sigma^2) \times P(B, \sigma^2)$$

That is the joint density $G\left(Y_t, B, \sigma^2\right)$ is the product of the marginal density of Y_t and the conditional density of the parameters $H(B, \sigma^2 \backslash Y_t)$. Or equivalently the joint density is the product of the conditional density of the data and the marginal density of the parameters. Rearranging the terms after the first equality leads to equation 2.4.

These steps in Bayesian analysis have a number of noteworthy features. First, the Bayesian econometrician is interested in the posterior *distribution* and not the mode of the likelihood function. Second, this approach combines prior information with the information in the data. In contrast, the classical econometrician focusses on information contained in the data about the parameters as summarised by the likelihood function.

To motivate the use of Gibbs sampling for estimating $H(B, \sigma^2 \backslash Y_t)$ we will consider the derivation of the posterior distribution in three circumstances. First we consider estimating the posterior distribution of B under the assumption that σ^2 is known. Next we consider estimating the posterior distribution of σ^2 under the assumption that B is known and finally we consider the general case when both sets of parameters are unknown.

2.1. Case 1: The posterior distribution of B assuming σ^2 is known. Consider the scenario where the econometrician wants to estimate β in equation 2.1 but knows the value of σ^2 already. As discussed above, the posterior distribution is derived using three steps.

Setting the prior. In the first step the researcher sets the prior distribution for β. A normally distributed prior $P(B) \sim N\left(B_0, \Sigma_0\right)$ for the coefficients is a *conjugate* prior. That is, when this prior is combined with the likelihood function this results in a posterior with the same distribution as the prior. Since the form of the posterior is known when using conjugate priors these are especially convenient from a practical point of view. The prior distribution is given by the following equation

$$\begin{aligned} & \left(2\pi\right)^{-K/2} \left|\Sigma_0\right|^{-\frac{1}{2}} \exp\left[-0.5\left(B - B_0\right)' \Sigma_0^{-1}\left(B - B_0\right)\right] \\ \propto \quad & \exp\left[-0.5\left(B - B_0\right)' \Sigma_0^{-1}\left(B - B_0\right)\right] \end{aligned} \tag{2.6}$$

The equation in 2.6 simply defines a normal distribution with mean B_0 and variance Σ_0. Note that for practical purposes we only need to consider terms in the exponent (second line of equation 2.6) as the first two terms in 2.6 $\left(\left(2\pi\right)^{-K/2} \left|\Sigma_0\right|^{-\frac{1}{2}}\right)$ are constants.

Setting up the likelihood function. In the second step, the researcher collects the data and forms the likelihood function:

$$\begin{aligned} F\left(Y_t \backslash B, \sigma^2\right) &= \left(2\pi\sigma^2\right)^{-T/2} \exp\left(-\frac{\left(Y_t - BX_t\right)'\left(Y_t - BX_t\right)}{2\sigma^2}\right) \\ &\propto \exp\left(-\frac{\left(Y_t - BX_t\right)'\left(Y_t - BX_t\right)}{2\sigma^2}\right) \end{aligned} \tag{2.7}$$

As σ^2 is assumed to be known in this example, we can drop the first term in equation 2.7 $\left(2\pi\sigma^2\right)^{-T/2}$

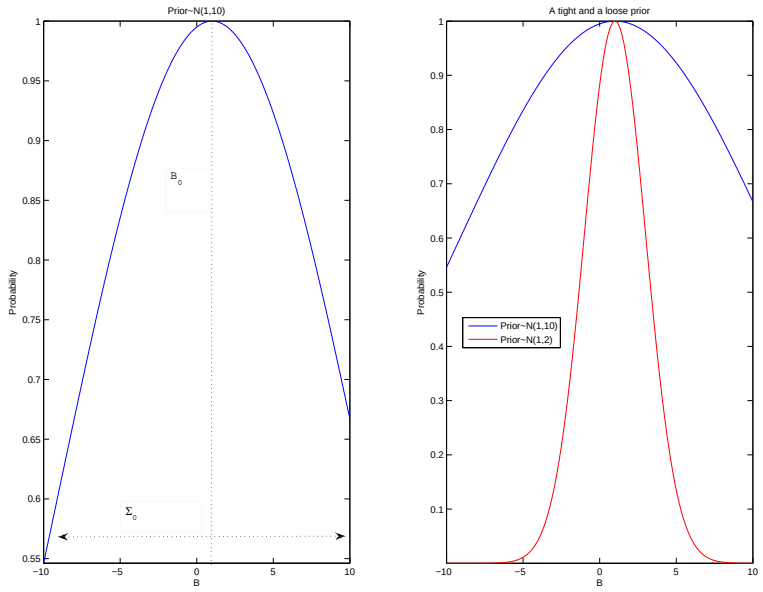

FIGURE 1. Loose and tight prior for the coefficients. An example

Calculating the posterior distribution. Recall from equation 2.5 that *the posterior distribution is proportional to the likelihood times the prior*. Therefore to find the posterior distribution for B (conditional on knowing σ^2) the researcher multiplies equation 2.6 and 2.7 to obtain

$$H\left(B\backslash\sigma^2, Y_t\right) \propto \exp\left[-0.5\left(B-B_0\right)'\Sigma_0^{-1}\left(B-B_0\right)\right] \times \exp\left(-\frac{\left(Y_t-BX_t\right)'\left(Y_t-BX_t\right)}{2\sigma^2}\right) \tag{2.8}$$

Equation 2.8 is simply a product of two normal distributions and the result is also a normal distribution. Hence the posterior distribution of B conditional on σ^2 is given by:

$$H\left(B\backslash\sigma^2, Y_t\right) \tilde{} N\left(M^*, V^*\right) \tag{2.9}$$

As shown in Hamilton (1994) pp 354 and Koop (2003) pp 61 the mean and the variance of this normal distribution are given by the following expressions

$$\begin{aligned} M^* &= \left(\Sigma_0^{-1}+\frac{1}{\sigma^2}X_t'X_t\right)^{-1}\left(\Sigma_0^{-1}B_0+\frac{1}{\sigma^2}X_t'Y_t\right) \\ V^* &= \left(\Sigma_0^{-1}+\frac{1}{\sigma^2}X_t'X_t\right)^{-1} \end{aligned} \tag{2.10}$$

Consider the expression for the mean of the conditional posterior distribution $M^* = \left(\Sigma_0^{-1}+\frac{1}{\sigma^2}X_t'X_t\right)^{-1}\left(\Sigma_0^{-1}B_0+\frac{1}{\sigma^2}X_t'Y_t\right)$. Note that the final term $X_t'Y_t$ can be re-written as $X_t'X_tB_{ols}$ where $B_{ols} = \left(X_t'X_t\right)^{-1}X_t'Y_t$. That is

$$M^* = \left(\Sigma_0^{-1}+\frac{1}{\sigma^2}X_t'X_t\right)^{-1}\left(\Sigma_0^{-1}B_0+\frac{1}{\sigma^2}X_t'X_tB_{ols}\right) \tag{2.11}$$

The second term of the expression in equation 2.11 shows that the mean of the conditional posterior distribution is weighted average of the prior mean B_0 and the maximum likelihood estimator B_{ols} with the weights given by the reciprocal of the variances of the two (in particular Σ_0^{-1} and $\frac{1}{\sigma^2}X_t'X_t$). A large number for Σ_0 would imply a very small weight on the prior and hence M^* would be dominated by the OLS estimate. A very small number for Σ_0, on the other hand, would imply that the conditional posterior mean is dominated by the prior. Note also that if the prior is removed from the expressions in equation 2.10 (i.e. if one removes B_0 and Σ_0^{-1} from the expressions) , one is left with the maximum likelihood estimates.

EXAMPLE 2. *Figure 1 shows a simple example about a prior distribution for a regression model with* 1 *coefficient* B. *The X-axis of these figures show a range of values of* B. *The Y-axis plots the value of the normal prior distribution associated with these values of* B. *The left panel shows a a prior distribution with a mean of* 1 *and a variance of* 10. *As expected, the prior distribution is centered at* 1 *and the width of the distribution reflects the variance. The right panel compares this prior distribution with a tighter prior centered around the same mean. In particular, the new prior distribution (shown as the red line) has a variance of* 2 *and is much more tightly concentrated around the mean.*

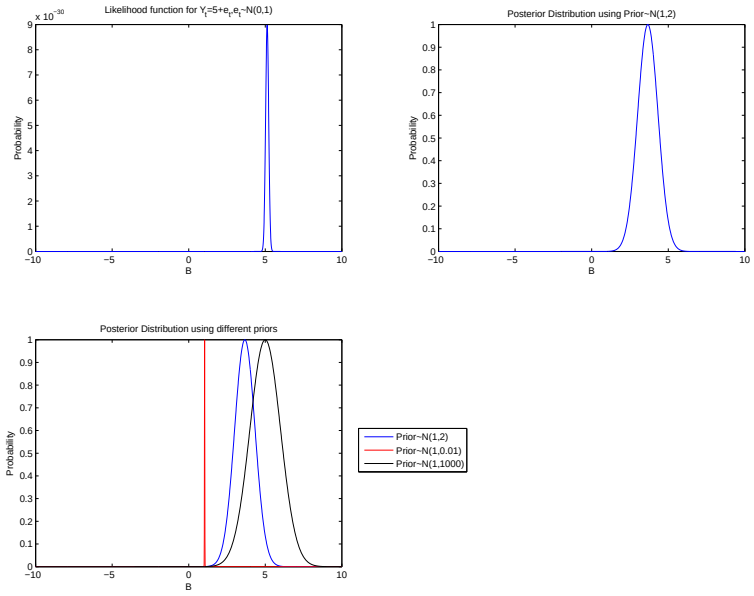

FIGURE 2. The posterior distribution for the model $Y_t = 5X_t + e_t, e_t \tilde{} N(0,1)$ using different priors

The tope left panel of figure 2 plots the likelihood function for the simple regression model $Y_t = BX_t + e_t, e_t \tilde{} N(0,1), B = 5$. *As expected the likelihood function has its peak at* $B = 5$. *The top right panel shows the posterior distribution which combines the prior distribution in figure 1* ($N(1,2)$ *shown as the red line) with the likelihood function. Note that as the posterior combines the prior information (with a mean of* 1*) and the likelihood function, the posterior distribution is not exactly centered around* 5*, but around a value slightly less than* 5*, reflecting the influence of the prior. Note that if the prior is tightened significantly and its variance reduced to* 0.01*, this has the affect of shifting the posterior distribution with the mass concentrated around* 1 *(red line in the bottom left panel). In contrast, a loose prior with a prior variance of* 1000 *is concentrated around* 5.

2.2. Case 2: The posterior distribution of σ^2 assuming B is known. In the second example we consider the estimation of σ^2 in equation 2.1 assuming that the value of B is known. The derivation of the (conditional) posterior distribution of σ^2 proceeds in exactly the same three steps

Setting the prior. The normal distribution allows for negative numbers and is therefore not appropriate as a prior distribution for σ^2. A conjugate prior for σ^2 is the inverse Gamma distribution or *equivalently* a conjugate prior for $1/\sigma^2$ is the Gamma distribution.

DEFINITION 1. ***(Gamma Distribution)****: Suppose we have* T *iid numbers from the normal distribution* v_t

$$v_t \tilde{} N\left(0, \frac{1}{\theta}\right)$$

If we calculate the sum of squares of $W = \sum_{t=1}^{T} v_t^2$*, then* W *is distributed as a Gamma distribution with* T *degrees of freedom and a scale parameter* θ

$$W \tilde{} \Gamma\left(\frac{T}{2}, \frac{\theta}{2}\right) \tag{2.12}$$

The probability density function for the Gamma distribution has a simple form and is given by

$$g(W) \propto W^{\frac{T}{2}-1} \exp\left(\frac{-W\theta}{2}\right) \tag{2.13}$$

where the mean of the distribution is defined as $E(W) = \frac{T}{\theta}$.

Setting the prior (continued). We set a Gamma prior for $1/\sigma^2$. That is $p\left(1/\sigma^2\right) \sim \Gamma\left(\frac{T_0}{2}, \frac{\theta_0}{2}\right)$ where T_0 denotes the prior degrees of freedom and θ_0 denotes the prior scale parameter. As discussed below, the choice of T_0 and θ_0 affects the mean and the variance of the prior. The prior density, therefore, has the following form (see equation 2.13.)

$$\frac{1}{\sigma^2}^{\frac{T_0}{2}-1} \exp\left(\frac{-\theta_0}{2\sigma^2}\right) \tag{2.14}$$

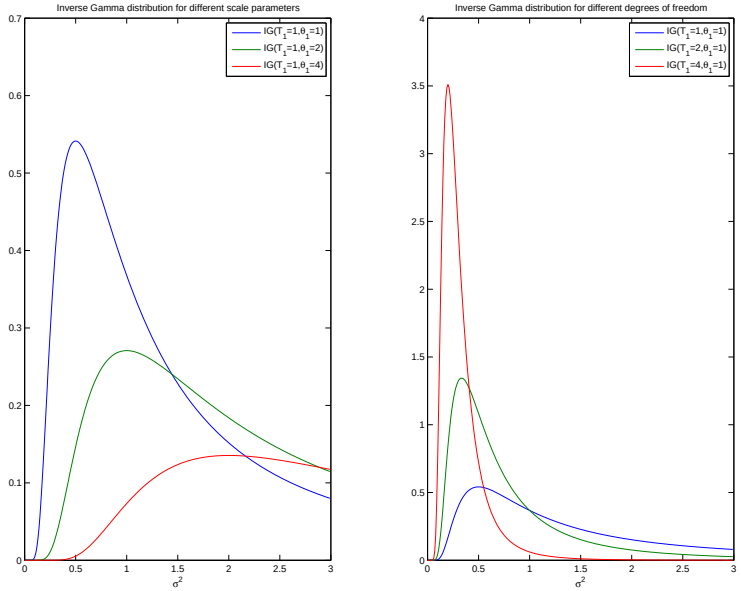

FIGURE 3. The inverse Gamma distribution for different degrees of freedom and scale parameters.

Setting up the likelihood function. In the second step, the researcher collects the data and forms the likelihood function:

$$\begin{aligned} F\left(Y_t \backslash B, \sigma^2\right) &= \left(2\pi\sigma^2\right)^{-T/2} \exp\left(-\frac{\left(Y_t - BX_t\right)'\left(Y_t - BX_t\right)}{2\sigma^2}\right) \\ &\propto \left(\sigma^2\right)^{-T/2} \exp\left(-\frac{\left(Y_t - BX_t\right)'\left(Y_t - BX_t\right)}{2\sigma^2}\right) \end{aligned} \tag{2.15}$$

As σ^2 is assumed to be unknown in this example, we cannot drop the entire first term in equation 2.15.

Calculating the posterior distribution. To calculate the posterior distribution of $1/\sigma^2$ (conditional on B) we multiply the prior distribution in equation 2.14 and the likelihood function 2.15 to obtain

$$H\left(\frac{1}{\sigma^2} \backslash B, Y_t\right) \propto \frac{1}{\sigma^2}^{\frac{T_0}{2}-1} \exp\left(\frac{-\theta_0}{2\sigma^2}\right) \times \sigma^{2-\frac{T}{2}} \exp\left(-\frac{1}{2\sigma^2}\left(Y_t - BX_t\right)'\left(Y_t - BX_t\right)\right)$$

$\rightarrow$

$$\frac{1}{\sigma^2}^{\frac{T_0}{2}-1-\frac{T}{2}} \exp\left(-\frac{1}{2\sigma^2}\left[\theta_0 + \left(Y_t - BX_t\right)'\left(Y_t - BX_t\right)\right]\right)$$

$\rightarrow$

$$\frac{1}{\sigma^2}^{\frac{T_1}{2}-1} \exp\left(-\frac{\theta_1}{2\sigma^2}\right) \tag{2.16}$$

The resulting conditional posterior distribution for $1/\sigma^2$ in equation 2.16 can immediately be recognised as a Gamma distribution with degrees of freedom $T_1 = \frac{T_0+T}{2}$ and $\theta_1 = \frac{\theta_0+(Y_t-BX_t)'(Y_t-BX_t)}{2}$. Note that the conditional posterior distribution for σ^2 is inverse Gamma with degrees of freedom T_1 and scale parameter θ_1.

Consider the mean of the conditional posterior distribution (given by $\frac{T_1}{\theta_1}$)

$$\frac{T_0 + T}{\theta_0 + \left(Y_t - BX_t\right)'\left(Y_t - BX_t\right)} \tag{2.17}$$

It is interesting to note that without the prior parameters T_0 and θ_0, equation 2.17 simply defines the reciprocal of the maximum likelihood estimator of σ^2.

EXAMPLE 4. *The left panel of figure 3 plots the inverse Gamma distribution with the degrees of freedom held fixed at $T_1 = 1$, but for scale parameter $\theta_1 = 1, 2, 4$. Note that as the scale parameter increases, the distribution becomes skewed to the right and the mean increases. This suggests that an inverse Gamma prior with a larger scale parameter incorporates a prior belief of a larger value for σ^2. The right panel of the figure plots the inverse Gamma distribution for $\theta_1 = 1$, but for degrees of freedom $T_1 = 1, 2, 4$. As the degrees of freedom increase, the inverse Gamma distribution is more tightly centered around the mean. This suggests that a higher value for the degrees of freedom implies a tighter set of prior beliefs.*

2.3. Case 3: The posterior distribution of σ^2 and B. We now turn to the empirically relevant case when both the coefficient vector B and the variance $1/\sigma^2$ (in equation 2.1) is unknown. We proceed in exactly the same three steps

Setting the prior. We set the joint prior density for

$$p\left(B, \frac{1}{\sigma^2}\right) = P\left(\frac{1}{\sigma^2}\right) \times P\left(B \backslash \frac{1}{\sigma^2}\right) \tag{2.18}$$

where $P\left(B \backslash \frac{1}{\sigma^2}\right) \tilde{} N(B_0, \sigma^2 \Sigma_0)$ and $P\left(\frac{1}{\sigma^2}\right) \tilde{} \Gamma\left(\frac{T_0}{2}, \frac{\theta_0}{2}\right)$. That is: $P\left(\frac{1}{\sigma^2}\right) = \frac{1}{\sigma^2}^{\frac{T_0}{2}-1} \exp\left(\frac{-\theta_0}{2\sigma^2}\right)$ as in section 2.2 and $P\left(B \backslash \frac{1}{\sigma^2}\right) = (2\pi)^{-K/2} \left|\sigma^2 \Sigma_0\right|^{-\frac{1}{2}} \exp\left[-0.5 \left(B - B_0\right)' \left(\sigma^2 \Sigma_0\right)^{-1} \left(B - B_0\right)\right]$. Note that the prior for B is set conditional on σ^2. This prior is referred to as the natural conjugate prior for the linear regression model. A natural conjugate prior is a conjugate prior which has the same functional form as the likelihood.

Setting up the likelihood function. As above, the likelihood function is given by

$$F\left(Y_t \backslash B, \sigma^2\right) = \left(2\pi\sigma^2\right)^{-T/2} \exp\left(-\frac{\left(Y_t - BX_t\right)' \left(Y_t - BX_t\right)}{2\sigma^2}\right) \tag{2.19}$$

Calculating the posterior distribution. The *joint* posterior distribution of B and the variance $1/\sigma^2$ is obtained by combining 2.18 and 2.19

$$H\left(\frac{1}{\sigma^2}, B \backslash Y_t\right) \propto p\left(B, \frac{1}{\sigma^2}\right) \times F\left(Y_t \backslash B, \sigma^2\right) \tag{2.20}$$

Note that equation 2.20 is a *joint* posterior distribution involving $\frac{1}{\sigma^2}$ and B. Its form is more complicated than the conditional distributions for B and $\frac{1}{\sigma^2}$ shown in sections 2.1 and 2.2. To proceed further in terms of inference, the researcher has to 'isolate' the component of the posterior relevant to B or $\frac{1}{\sigma^2}$. For example, to conduct inference about B, the researcher has to derive the *marginal* posterior distribution for B. Similarly, inference on $\frac{1}{\sigma^2}$ is based on the marginal posterior distribution for $\frac{1}{\sigma^2}$. The marginal posterior for B is defined as

$$H\left(B \backslash Y_t\right) = \int_0^\infty H\left(\frac{1}{\sigma^2}, B \backslash Y_t\right) d\frac{1}{\sigma^2} \tag{2.21}$$

while the marginal posterior for $\frac{1}{\sigma^2}$ is given by

$$H\left(\frac{1}{\sigma^2} \backslash Y_t\right) = \int_0^\infty H\left(\frac{1}{\sigma^2}, B \backslash Y_t\right) dB \tag{2.22}$$

In the case of this simple linear regression model under *the natural conjugate prior*, analytical results for these integrals are available. As shown in Hamilton (1994) pp 357, the marginal posterior distribution for B is a multivariate T distribution, while the marginal posterior for $\frac{1}{\sigma^2}$ is a Gamma distribution. An intuitive description of these analytical results can also be found in Koop (2003) Chapter 2.

However, for the linear regression model with other prior distributions (for example where the prior for the coefficients is set independently from the prior for the variance) analytical derivation of the joint posterior and then the marginal posterior distribution is not possible. Similarly, in more complex models with a larger set of unknown parameters (i.e. models that may be more useful for inference and forecasting) these analytical results may be difficult to obtain. This may happen if the form of the joint posterior is unknown or is too complex for analytical integration.

Readers should pause at this point and reflect on two key messages from the three cases considered above:

EXAMPLE 3

- As shown by Case 1 and Case 2, *conditional* posterior distributions are relatively easy to derive and work with.
- In contrast, as shown by Case 3, derivation of the marginal posterior distribution (from a joint posterior distribution) requires analytical integration which may prove difficult in complex models.

This need for analytical integration to calculate the marginal posterior distribution was the main stumbling block of Bayesian analysis making it difficult for applied researchers.

3. Gibbs Sampling for the linear regression model

It was the development of simulation method such as Gibbs sampling which greatly simplified the integration step discussed above and made it possible to easily extend Bayesian analysis to a variety of econometric models.

DEFINITION 2. *Gibbs sampling is a numerical method that uses draws from* ***conditional distributions*** *to approximate joint and marginal distributions.*

As discussed in case 3 above, researchers are interested in marginal posterior distributions which may be difficult to derive analytically. In contrast, the conditional posterior distribution of each set of parameters is readily available. According to definition 2, one can approximate the marginal posterior distribution by sampling from the conditional distributions.

We describe this algorithm in detail below, first in a general setting and then applied specifically to the linear regression model. *Most importantly, we then describe how to code the algorithm for linear regression models. Note that all the files referred to below are saved in the sub-folder called chapter 1 in the main folder called code.*

3.1. Gibbs Sampling a general description. Suppose we have a joint distribution of k variables

$$f(x_1, x_2...x_k) \tag{3.1}$$

This may, for example, be a joint posterior distribution.

and we are interested in obtaining the marginal distributions

$$f(x_i), i = 1...k \tag{3.2}$$

The standard way to do this is to integrate the joint distribution in 3.1. However, as discussed above, this integration may be difficult or infeasible in some cases. It may be that the exact form of 3.1 is unknown or is to complicated for direct analytical integration.

Assume that the form of the conditional distributions $f(x_i \backslash x_j), i \neq j$ is known. A Gibbs sampling algorithm with the following steps can be used to approximate the marginal distributions.

(1) Set starting values for $x_1....x_k$
$$x_1^0, ...x_K^0$$
where the superscript 0 denotes the starting values.

(2) Sample x_1^1 from the distribution of x_1 conditional on current values of $x_2....x_k$
$$f\left(x_1^1 \backslash x_2^0, ..x_k^0\right)$$

(3) Sample x_2^1 from the distribution of x_2 conditional on current values of $x_1, x_3...x_k$
$$f\left(x_2^1 \backslash x_1^1, x_3^0..x_k^0\right)$$

.
.
.

k. Sample x_k^1 from the distribution of x_k conditional on current values of $x_1, x_2...x_{k-1}$
$$f\left(x_k^1 \backslash x_1^1, x_2^1..x_{k-1}^1\right)$$
to complete 1 *iteration of the Gibbs sampling algorithm.*

As the number of Gibbs iterations increases to infinity, the samples or draws from the conditional distributions converge to the joint and marginal distributions of x_i at an exponential rate (for a proof of convergence see Casella and George (1992)). Therefore after a large enough number of iterations, the marginal distributions can be approximated by the empirical distribution of x_i.

In other words, one repeats the Gibbs iterations M times (ie a number of iterations large enough for convergence) and saves the last H draws of x_i (for eg $H = 1000$). This implies that the researcher is left with H values for $x_1....x_k$. The histogram for $x_1....x_k$ (or any other estimate of the empirical density) is an approximation for the marginal density of $x_1....x_k$.

Thus an estimate of the mean of the marginal posterior distribution for x_i is simply the sample mean of the H retained draws

$$\frac{1}{H}\sum_{b=1}^{H} x_i^b$$

where the superscript b indexes the (retained) Gibbs iterations. Similarly, the estimate of the variance of the marginal posterior distribution is given by

How many Gibbs iterations are required for convergence? We will deal with this question in detail in section section 3.7 below.

One crucial thing to note is that the implementation of the Gibbs sampling algorithm requires the researcher to know the form of the conditional distributions $f(x_i \backslash x_j)$. In addition, it must be possible to take random draws from these conditional distributions.

3.2. Gibbs Sampling for a linear regression. We now proceed to our first practical example involving a linear regression model. We first describe the application of the Gibbs sampling algorithm to the regression. This is followed immediately by a line by line description of Matlab code needed to implement the algorithm.

Consider the estimation of the following AR(2) model via Gibbs sampling

$$Y_t = \alpha + B_1 Y_{t-1} + B_2 Y_{t-2} + v_t, v_t \tilde{} N(0, \sigma^2) \tag{3.3}$$

where Y_t is annual CPI inflation for the US over the period 1948Q1 to 2010Q3. Let $X_t = \{1, Y_{t-1,1} Y_{t-2}\}$ denote the RHS variables in equation 3.3 and $B = \{\alpha, B_1, B_2\}$ the coefficient vector. Our aim is to approximate the marginal posterior distribution of α, B_1, B_2 and σ^2. As discussed above it is difficult to derive these marginal distributions analytically. Note, however, that we readily derived the posterior distribution of $B = \{\alpha, B_1, B_2\}$ conditional on σ^2

(see section 2.1) and the posterior distribution of σ^2 conditional on $B = \{\alpha, B_1, B_2\}$ (see section 2.2) Estimation of this model proceeds in the following steps

Step 1 Set priors and starting values. We set a normal prior for the coefficients B.

$$p(B) \tilde{} N\left(\underbrace{\begin{pmatrix} \alpha^0 \\ B_1^0 \\ B_2^0 \end{pmatrix}}_{B_0}, \underbrace{\begin{pmatrix} \Sigma_\alpha & 0 & 0 \\ 0 & \Sigma_{B1} & 0 \\ 0 & 0 & \Sigma_{B2} \end{pmatrix}}_{\Sigma_0}\right) \tag{3.4}$$

In other words, we specify the prior means for each coefficient in B (denoted as B_0 in 3.4) and the prior variance Σ_0. For this example (with three coefficients) B_0 is a 3×1 vector, while Σ_0 is 3×3 matrix with each diagonal element specifying the prior variance of the corresponding element of B_0.

We set an inverse Gamma prior for σ^2 and set the prior degrees of freedom T_0 and the prior scale matrix θ_0 (see equation 3.5). We will therefore work with the inverse Gamma distribution in the Gibbs sampler below. Note that this is equivalent to working with Gamma distribution and $1/\sigma^2$.

$$p\left(\sigma^2\right) \tilde{} \Gamma^{-1}\left(\frac{T_0}{2}, \frac{\theta_0}{2}\right) \tag{3.5}$$

To initialise the Gibbs sampler we need a starting value for either σ^2 or B. In this example we will assume that the starting value for $\sigma^2 = \sigma^2_{OLS}$ where σ^2_{OLS} is the OLS estimate of σ^2. In linear models (such as linear regressions and Vector Autoregressions) the choice of starting values has, in our experience, little impact on the final results given that the number of Gibbs iterations is large enough.

Step 2 Given a value for σ^2 we sample from the conditional posterior distribution of B. As discussed in section 2.1, this is a normal distribution with a known mean and variance given

$$H\left(B \backslash \sigma^2, Y_t\right) \tilde{} N\left(M^*, V^*\right) \tag{3.6}$$

where

$$\begin{aligned} \underset{(3\times1)}{M^*} &= \left(\Sigma_0^{-1} + \frac{1}{\sigma^2} X_t' X_t\right)^{-1} \left(\Sigma_0^{-1} B_0 + \frac{1}{\sigma^2} X_t' Y_t\right) \\ \underset{(3\times3)}{V^*} &= \left(\Sigma_0^{-1} + \frac{1}{\sigma^2} X_t' X_t\right)^{-1} \end{aligned} \tag{3.7}$$

Note that we have all the ingredients to calculate M^* and V^* which in this example are 3×1 and 3×3 matrices respectively. We now need a sample from the normal distribution with mean M^* and variance V^*. For this we can use the following algorithm.

ALGORITHM 1. *To sample a $k \times 1$vector denoted by z from the $N(m, v)$ distribution, first generate $k \times 1$ numbers from the standard normal distribution (call these z^0. Note that all computer packages will provide a routine to do this). These standard normal numbers can then be transformed such that the mean is equal to m and variance equals v using the following transformation*

$$z = m + z^0 \times v^{1/2}$$

Thus one adds the mean and multiplies z^0 by the square root of the variance.

Step 2 (continued) The procedure in algorithm 1 suggests that once we have calculated M^* and V^*, the draw for B is obtained as

$$\underset{(3\times1)}{B^1} = \underset{(3\times1)}{M^*} + \left[\underset{(1\times3)}{\bar{B}} \times \underset{(3\times3)}{\left(V^*\right)^{1/2}}\right]' \tag{3.8}$$

where $\bar{B}$ is a 1×3 vector from the standard normal distribution. Note that the superscript 1 in B^1 denotes the first Gibbs iteration.

Step 3 Given the draw B^1, we draw σ^2 form its conditional posterior distribution. As shown in section 2.2 the conditional posterior distribution for σ^2 is inverse Gamma

$$H\left(\sigma^2 \backslash B, Y_t\right) \tilde{} \Gamma^{-1}\left(\frac{T_1}{2}, \frac{\theta_1}{2}\right) \tag{3.9}$$

where

$$\begin{aligned} T_1 &= T_0 + T \\ \theta_1 &= \theta_0 + \left(Y_t - B^1 X_t\right)' \left(Y_t - B^1 X_t\right) \end{aligned} \tag{3.10}$$

A crucial thing to note about the posterior scale parameter of this distribution θ_1 is the fact that the second term $\left(\left(Y_t - B^1 X_t\right)' \left(Y_t - B^1 X_t\right)\right)$ is calculated using the previous draw of the coefficient vector (in this case B^1). To draw from the inverse Gamma distribution in equation 3.9 we first calculate the parameters in equation 3.10 and then use the following algorithm to draw $\left(\sigma^2\right)^1$ from the inverse Gamma distribution (note that $\left(\sigma^2\right)^i$ denotes the *ith* Gibbs draw) .

ALGORITHM 2. *To sample a scalar z from the Inverse Gamma distribution with degrees of freedom* $\frac{T}{2}$ *and scale parameter* $\frac{D}{2}$ *i.e.* $\Gamma^{-1}(\frac{T}{2}, \frac{D}{2})$*: Generate T numbers form the standard normal distribution* $z^0 \tilde{} N(0,1)$*. Then*

$$z = \frac{D}{z^{0\prime} z^0}$$

is a draw from the $\Gamma^{-1}(\frac{T}{2}, \frac{D}{2})$ *distribution.*

Step 4 Repeat steps 2 and 3 M times to obtain $B^1 ... B^M$ and $\left(\sigma^2\right)^1 .... \left(\sigma^2\right)^M$. The last H values of B and σ^2 from these iterations is used to form the empirical distribution of these parameters. Note that this empirical distribution is an approximation to the marginal posterior distribution. Note also that the first $M - H$ iterations which are discarded are referred to as burn-in iterations. These are the number of iterations required for the Gibbs sampler to converge.

Its worth noting that it makes no difference which order steps 2 and 3 are repeated. For example one could start the Gibbs sampler by drawing σ^2 conditional on starting values for B (rather than the other way around as we have done here)

3.2.1. *Inference using output from the Gibbs sampler.* The Gibbs sampler applied to the linear regression model produces a sequence of draws from the approximate marginal posterior distribution of B and σ^2. The mean of these draws is an approximation to the posterior mean and provides a point estimate of of B and σ^2. The percentiles calculated from these draws can be used to produce posterior density intervals. For example, the 5^{th} and the 95^{th} percentiles approximate the 10% highest posterior density intervals (HPDI) or 10% credible sets which can be used for simple hypothesis testing. For example, if the highest posterior density interval for B does not contain zero, this is evidence that the hypothesis that $B = 0$ can be rejected.

More formal methods for model comparison involve the marginal likelihood $F(Y)$ mentioned in section 2. The marginal likelihood is defined as

$$F(Y) = \int F\left(Y \backslash B, \sigma^2\right) p\left(B, \sigma^2\right) d\Xi$$

where $\Xi = B, \sigma^2$. In other words, the marginal likelihood represents the posterior distribution with the parameters integrated out. Consider two models M_1 and M_2. Model M_1 is preferred if $F_{M_1}(Y) > F_{M_2}(Y)$ or the Bayes factor $\frac{F_{M_1}(Y)}{F_{M_2}(Y)}$ is larger than 1. In comparison to HPDIs, inference based on marginal likelihoods or Bayes factors is more complicated from a computational and statistical point of view. First, while an analytical expression for $F(Y)$ is available for the linear regression model under the natural conjugate prior, numerical methods are generally required to calculate the integral in the expression for $F(Y)$ above. In the appendix to this chapter, we provide an example of how Gibbs sampling can be used to compute the marginal likelihood for the linear regression model. Second, model comparison using marginal likelihoods requires the researchers to use proper priors (i.e. prior distributions that integrate to 1). In addition, using non-informative priors may lead to problems when interpreting Bayes Factors. An excellent description of these issues can be found in Koop (2003) pp 38.

3.3. Gibbs Sampling for a linear regression in Matlab (example1.m). We now go through the Matlab code needed for implementing the algorithm described in the section above. Note that the aim is to estimate the following AR(2) model via Gibbs sampling.

$$Y_t = \alpha + B_1 Y_{t-1} + B_2 Y_{t-2} + v_t, v_t \tilde{} N(0, \sigma^2) \tag{3.11}$$

where Y_t is annual CPI inflation for the US over the period 1948Q1 to 2010Q3 and $B = \{\alpha, B_1, B_2\}$. The code presented below is marked with comments for convenience. The same code without comments accompanies this monograph and is arranged by chapter. The code for the example we consider in this section is called example1.m and is saved in the folder

Consider the code for this example presented in 4 and 5. Line 2 of the code adds functions that are needed as utilities–eg for taking lags or differences. We will not discuss these further. On line 5, we load data for US inflation from an excel file. Line 7 creates the regressors, a constant and two lags of inflation (using the function lag0 in the folder functions). Line 11 specifies the total number of time series observations after removing the missing values generated after taking lags. Line 14 sets prior mean for the regression coefficients.

$$\begin{pmatrix} \alpha^0 \\ B_1^0 \\ B_2^0 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}$$

The prior mean for each coefficient is set to zero in this example. The prior variance is set to an identity matrix on line 15 in this example.

$$\begin{pmatrix} \Sigma_\alpha & 0 & 0 \\ 0 & \Sigma_{B1} & 0 \\ 0 & 0 & \Sigma_{B2} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

Line 17 sets the prior degrees of freedom for the inverse Gamma distribution while line sets θ_0 the prior scale parameter. Line 20 sets the starting value of B, while line 21 sets the starting value for σ^2. Line 22 specifies the total

```
1 clear
2 addpath('functions'); %this line adds functions to take lags etc
3 %an AR 2 model for US inflation
4 %load inflation data
5 Y=xlsread('\data\inflation.xls');
6 T=rows(Y);
7 X=[ones(T,1) lag0(Y,1) lag0(Y,2)];
8 %remove missing obs
9 Y=Y(3:end);
10 X=X(3:end,:);
11 T=rows(X);
12 %step 1 set priors and starting values
13 %priors for B
```

$$\begin{pmatrix} \alpha^0 \\ B_1^0 \\ B_2^0 \end{pmatrix}, \begin{pmatrix} \Sigma_\alpha & 0 & 0 \\ 0 & \Sigma_{B1} & 0 \\ 0 & 0 & \Sigma_{B2} \end{pmatrix}$$

B_0 Σ_0

```
14 B0=[0;0;0];
15 Sigma0=eye(3);
16 %priors for sigma2
17 T0=1;
```

θ_0

```
18 D0=0.1;
19 %starting values
20 B=B0;
21 sigma2=1;
22 reps=5000;   %total numbers of Gibbs iterations
23 burn=4000;   %percent of burn-in iterations
24 out1=[];
25 out2=[];
26 for i=1:reps
27 %step 2 Sample B conditional on sigma N(M*,V*)
28
M=inv(inv(Sigma0)+(1/sigma2)*(X'*X))*(inv(Sigma0)*B0+(1/sigma2)*X'*Y);
```

$$\underset{(3\times1)}{M^*} = \left(\Sigma_0^{-1} + \tfrac{1}{\sigma^2}X_t'X_t\right)^{-1}\left(\Sigma_0^{-1}B_0 + \tfrac{1}{\sigma^2}X_t'Y_t\right)$$

```
29 V=inv(inv(Sigma0)+(1/sigma2)*(X'*X));
```

$$\underset{(3\times3)}{V^*} = \left(\Sigma_0^{-1} + \tfrac{1}{\sigma^2}X_t'X_t\right)^{-1}$$

```
30 chck=-1;
31 while chck<0                    %check for stability
32 B=M+(randn(1,3)*chol(V))';
```

$$\underset{(3\times1)}{B^1} = \underset{(3\times1)}{M^*} + \left[\underset{(1\times3)}{\bar{B}} \times \underset{(3\times3)}{(V^*)^{1/2}}\right]'$$

```
33 b=[B(2) B(3);1   0];
34 ee=max(abs(eig(b)));
35 if ee<=1
36     chck=1;
```

FIGURE 4. Example 1: Matlab code

number of Gibbs iterations, while line 23 specifies the number to discard (in this example we save 1000 iterations for inference). out1 and out2 on line 24 and 25 are empty matrices that will save the draws of B and σ^2 respectively. Line 26 starts the main loop that carries out the Gibbs iterations. On line 28, we begin the first step of the Gibbs algorithm and calculate the mean of the conditional posterior distribution of B ($M^* = \left(\Sigma_0^{-1} + \frac{1}{\sigma^2}X_t'X_t\right)^{-1}\left(\Sigma_0^{-1}B_0 + \frac{1}{\sigma^2}X_t'Y_t\right)$) and on line 29 we calculate the variance of this conditional posterior distribution. Line 32 draws from the normal distribution with this mean and variance. Note that it is standard to restrict the draw of the AR coefficients to be stable. This is why line 31 has a while loop which keeps on drawing from the coefficients from the normal distribution if the draws are unstable. Stability is checked on line 33 by computing the eigenvalues of the coefficient matrix written

```
end
end
%step 3 sample sigma2 conditional on B from IG(T1,D1);
%compute residuals
resids=Y-X*B;
%compute posterior df and scale matrix
T1=T0+T;
```

$$T_1 = T_0 + T$$

```
D1=D0+resids'*resids;
```

$$\theta_1 = \theta_0 + (Y_t - B^1 X_t)'(Y_t - B^1 X_t)$$

```
%draw from IG
z0=randn(T1,1);
z0z0=z0'*z0;
sigma2=D1/z0z0;
```

$$z = \frac{D}{z^{0\prime} z^0}$$

```
if i>burn
    out1=[out1;B'];
    out2=[out2;sigma2];
end
end
%plot marginal posterior distributions
subplot(2,2,1);
hist(out1(:,1),50);
axis tight
title('Constant');
subplot(2,2,2);
hist(out1(:,2),50);
axis tight
title('AR(1) coefficient');
subplot(2,2,3);
hist(out1(:,3),50);
axis tight
title('AR(2) coefficient');
subplot(2,2,4);
hist(out2(:,1),50);
axis tight
title('\sigma^{2}');
%compute mean of the marginal posterior distribution of B
MB=mean(out1);
%compute standard error
VB=std(out1);
%compute 95% error band
EB=prctile(out1,[5 95]);
```

Published with MATLAB® 7.9

FIGURE 5. Example 1: Matlab Code (continued)

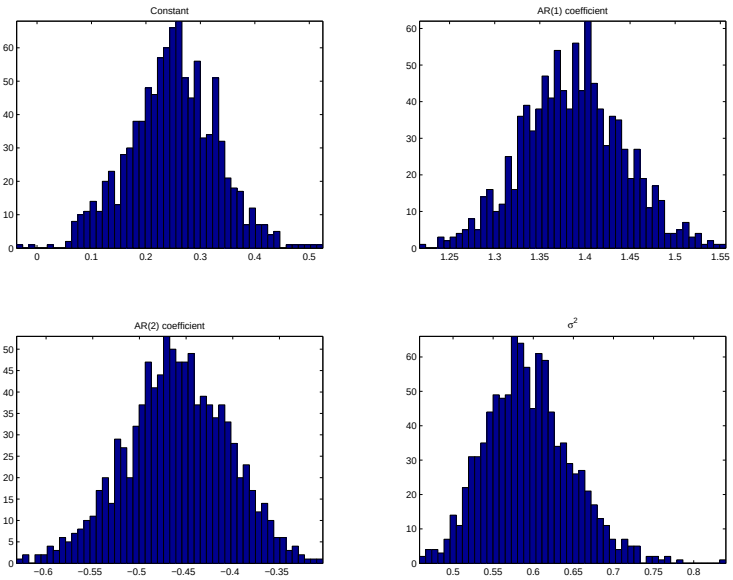

FIGURE 6. Results using example1.m

	Posterior Mean	Standard Deviation	5th and 95th percentiles
α	0.2494	0.0799	$(0.1104, 0.3765)$
B_1	1.3867	0.0557	$(1.2922, 1.4806)$
B_2	-0.4600	0.0550	$(-0.5532, -0.3709)$

TABLE 1. Results using example1.m

in first order companion form. That is the AR(2) model is re-written as (this is the companion form)

$$\begin{pmatrix} Y_t \\ Y_{t-1} \end{pmatrix} = \begin{pmatrix} \alpha \\ 0 \end{pmatrix} + \begin{pmatrix} B_1 & B_2 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} Y_{t-1} \\ Y_{t-2} \end{pmatrix} + \begin{pmatrix} v_t \\ 0 \end{pmatrix}$$

Then the AR model is stable if the eigenvalues of $\begin{pmatrix} B_1 & B_2 \\ 1 & 0 \end{pmatrix}$ are less than or equal to 1 in absolute value. Note that this check for stability is not required for the Gibbs sampling algorithm but usually added by researchers for practical convenience. Line 41 computes the residuals using the last draw of the coefficients. Line 43 computes the posterior degrees of freedom for the inverse Gamma distribution $T_1 = T_0 + T$. Line 44 computes the posterior scale parameter $\theta_1 = \theta_0 + \left(Y_t - B^1 X_t\right)' \left(Y_t - B^1 X_t\right)$. Line 46 to 48 draw from the inverse Gamma distribution using algorithm 2. Lines 49 to 51 save the draws of B and σ^2 once the number of iterations exceed the burn-in period. Running this file produces the histograms shown in figure 6 (see lines 54 to 70 in example1.m–these histograms are drawn using the retained draws in out1 and out2). These histograms are the Gibbs sampling estimate of the marginal posterior distribution of the coefficients and the variance. Note that the mean of the posterior distribution is easily calculated as the sample mean of these saved draws. Similarly, the sample standard deviation and percentiles provide measures of uncertainty. Researchers usually report the posterior mean, the standard deviation and the 5th and 95th percentiles of the posterior distribution. Example1.m produces the following moments for the coefficients (see table 1).

Note that the percentiles of the distribution are a useful measure of uncertainty. These represent HPDIs, or the posterior belief that the parameter lies within a range (see Canova (2007) page 337 and Koop (2003) pp 43). Suppose that the lower bound for α was less than 0. Then this would indicate that one cannot exclude the possibility that the posterior mean for α is equal to zero.

3.4. Gibbs Sampling for a linear regression in Matlab and forecasting (example2.m). The file example2.m considers the same model as in the previous subsection. However, we know use the AR model to forecast inflation and build the distribution of the forecast. This example shows that one can easily obtain the distribution of functions of the regression coefficients. Note that the forecast from an AR(2) model is easily obtained via simulation. In other words, given a value for the current and lagged data and the regression coefficients, the 1 period ahead forecast is

$$\hat{Y}_{t+1} = \alpha + B_1 Y_t + B_2 Y_{t-1} + (\sigma v^*) \tag{3.12}$$

```
1 clear
2 addpath('functions'); %this line adds functions to take lags etc
3 %an AR 2 model for US inflation
4 %load inflation data
5 Y=xlsread('\data\inflation.xls');
6 T=rows(Y);
7 X=[ones(T,1) lag0(Y,1) lag0(Y,2)];
8 %remove missing obs
9 Y=Y(3:end);
10 X=X(3:end,:);
11 T=rows(X);
12 %step 1 set priors and starting values
13 %priors for B
14 B0=[0;0;0];
15 Sigma0=eye(3);
16 %priors for sigma2
17 T0=1;
18 D0=0.1;
19 %starting values
20 B=B0;
21 sigma2=1;
22 reps=5000;   %total numbers of Gibbs iterations
23 burn=4000;   %percent of burn-in iterations
24 out1=[];
25 out2=[];
26 out3=[];
27 for i=1:reps
28 %step 2 Sample B conditional on sigma N(M*,V*)
29
M=inv(inv(Sigma0)+(1/sigma2)*(X'*X))*(inv(Sigma0)*B0+(1/sigma2)*X'*Y);
30 V=inv(inv(Sigma0)+(1/sigma2)*(X'*X));
31 chck=-1;
32 while chck<0                    %check for stability
33 B=M+(randn(1,3)*chol(V))';
34 b=[B(2) B(3);1   0];
35 ee=max(abs(eig(b)));
36 if ee<=1
37     chck=1;
38 end
39 end
40 %step 3 sample sigma2 conditional on B from IG(T1,D1);
41 %compute residuals
42 resids=Y-X*B;
43 %compute posterior df and scale matrix
44 T1=T0+T;
45 D1=D0+resids'*resids;
46 %draw from IG
47 z0=randn(T1,1);
48 z0z0=z0'*z0;
49 sigma2=D1/z0z0;
50 if i>burn
51     out1=[out1;B'];
52     out2=[out2;sigma2];
53
54     %compute forecast for 2 years
55     yhat=zeros(14,1);
56     yhat(1:2)=Y(end-1:end); %starting values
57     cfactor=sqrt(sigma2);  %standard deviation of the shocks
58     for m=3:14
59         yhat(m)=[1 yhat(m-1) yhat(m-2)]*B+(randn(1,1)*cfactor);
```

FIGURE 7. Example 2: Matlab code

where v^* is a scalar drawn from the standard normal distribution. Similarly, the 2 period ahead forecast is

$$\hat{Y}_{t+2} = \alpha + B_1\hat{Y}_{t+1} + B_2 Y_t + (\sigma v^*) \tag{3.13}$$

and so forth. Note that we incorporate future shock uncertainty by adding the term σv^* i.e. a draw from the normal distribution with mean 0 and variance σ^2.

The code shown in figures 7 and 8 is identical to example 1 until line 54. Once past the burn in stage, we not only save the draws from the conditional distributions of the coefficients and the variance, but we use these draws to compute a two year ahead forecast for inflation. Line 55 intialises an empty matrix yhat which will save the forecast.

```
60      end
61      %save
62      out3=[out3 [Y;yhat(3:end)]];
63  end
64  end
65  %plot marginal posterior distributions
66  figure(1)
67  subplot(2,2,1);
68  hist(out1(:,1),50);
69  axis tight
70  title('Constant');
71  subplot(2,2,2);
72  hist(out1(:,2),50);
73  axis tight
74  title('AR(1) coefficient');
75  subplot(2,2,3);
76  hist(out1(:,3),50);
77  axis tight
78  title('AR(2) coefficient');
79  subplot(2,2,4);
80  hist(out2(:,1),50);
81  axis tight
82  title('\sigma^{2}');
83  %compute mean of the marginal posterior distribution of B
84  MB=mean(out1);
85  %compute standard error
86  VB=std(out1);
87  %compute 95% error band
88  EB=prctile(out1,[5 95]);
89  %plot forecast distribution
90  figure(2)
91  TT=1947.25:0.25:2012.5;
92  outx=prctile(out3',[10 20 30 40 50 60 70 80 90])'; %here prctile
calculates percentiles of the forecast distribution
93  plot(TT,outx);
94  xlim([2000 2013])
```

Published with MATLAB® 7.9

FIGURE 8. Example 2: Matlab code (continued)

Line 56 fills the first two values of yhat as actual values of inflation in the last two periods of the sample. Line 58 to 60 carries out the recursion shown in equations 3.12 and 3.13 for 12 periods. Line 62 saves actual inflation and the forecast in a matrix out3. The crucial thing to note is that this done for each Gibbs iteration after the burn-in period. Therefore in the end we have a set of 1000 forecasts. This represents an estimate of the posterior density. On line 92 we calculate the percentiles of the 1000 forecasts. The result gives us a fan chart for the inflation forecast shown in figure 9.

3.5. Gibbs Sampling for a linear regression with serial correlation. We now proceed to our second main example involving the linear regression model. We illustrate the power of the Gibbs sampler by considering

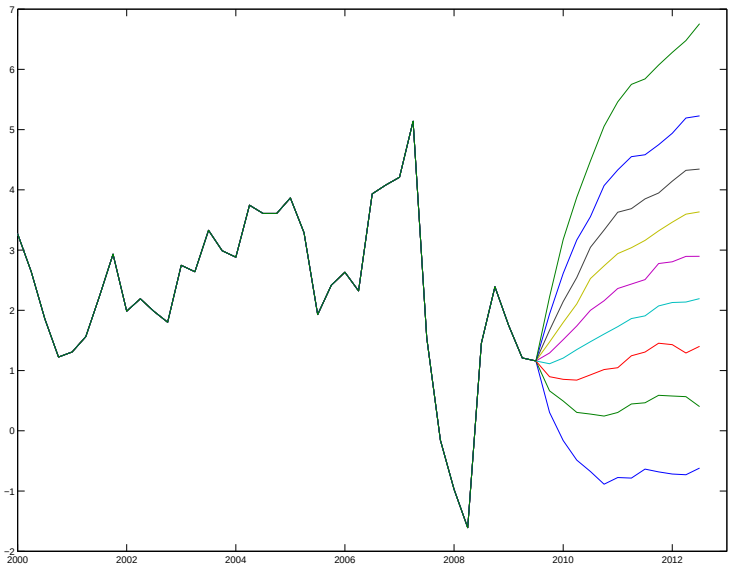

FIGURE 9. The distribution of the forecast of inflation using example2.m

the model in 3.3 but allowing for first order serial correlation in the residuals. We first describe the application of the Gibbs sampling algorithm to the regression. This is followed immediately by a line by line description of Matlab code needed to implement the algorithm. This algorithm was first developed in Chib (1993).

Consider the estimation of the following AR(2) model via Gibbs sampling

$$\begin{aligned} Y_t &= \alpha + B_1 Y_{t-1} + B_2 Y_{t-2} + v_t \\ v_t &= \rho v_{t-1} + \varepsilon_t, \varepsilon_t \sim N(0, \sigma^2) \end{aligned} \tag{3.14}$$

where Y_t is annual CPI inflation for the US over the period 1948Q1 to 2010Q3. Let $X_t = \{1, Y_{t-1,1} Y_{t-2}\}$ denote the RHS variables in equation 3.3 and $B = \{\alpha, B_1, B_2\}$ the coefficient vector. Our aim is to approximate the marginal posterior distribution of α, B_1, B_2 and σ^2 *and* ρ.

The key to seeting up the Gibbs sampler for this model is to make the following two observations

- Suppose we knew the value of ρ. Then the model in equation 3.14 can be transformed to remove the serial correlation. In particular we can re-write the model as

$$(\underbrace{Y_t - \rho Y_{t-1}}_{Y_t^*}) = \alpha(1-\rho) + B_1(\underbrace{Y_{t-1} - \rho Y_{t-2}}_{Y_{t-1}^*}) + B_2(\underbrace{Y_{t-2} - \rho Y_{t-3}}_{Y_{t-2}^*}) + (\underbrace{v_t - \rho v_{t-1}}_{\varepsilon_t}) \tag{3.15}$$

  That is we subtract the lag of each variable times the serial correlation coefficient ρ. Note that the transformed error term $v_t - \rho v_{t-1}$ is serially uncorrelated. Therefore after this transformation we are back to the linear regression framework we saw in the first example (see section 3.2). In other words, after removing the serial correlation, the conditional distribution of the coefficients and of the error variance is exactly as described for the standard linear regression model in section 3.2.
- Suppose we know α, B_1 and B_2. Then we can compute $v_t = Y_t - (\alpha + B_1 Y_{t-1} + B_2 Y_{t-2})$ and treat the equation $v_t = \rho v_{t-1} + \varepsilon_t, \varepsilon_t \sim N(0, \sigma^2)$ as a linear regression model in v_t. Again, this is just a standard linear regression model with an iid error term and the standard formulas for the conditional distribution of the regression coefficient ρ and the error variance σ^2 applies.

These two observations clearly suggest that to estimate this model, the Gibbs sampler needs three steps (instead of two in the previous example). We draw α, B_1 and B_2 conditional on knowing σ^2 and ρ after transforming the model to remove serial correlation (as in equation 3.15). Conditional on α, B_1 and B_2 and σ^2 we draw ρ. Finally, conditional on α, B_1,B_2 and ρ we draw σ^2. The steps are as follows

Step 1 Set priors and starting values. We set a normal prior for the coefficients B.

$$p(B) \tilde{} N\left(\underbrace{\begin{pmatrix} \alpha^0 \\ B_1^0 \\ B_2^0 \end{pmatrix}}_{B_0}, \underbrace{\begin{pmatrix} \Sigma_\alpha & 0 & 0 \\ 0 & \Sigma_{B1} & 0 \\ 0 & 0 & \Sigma_{B2} \end{pmatrix}}_{\Sigma_0} \right) \tag{3.16}$$

In other words, we specify the prior means for each coefficient in B (denoted as B_0 in 3.4) and the prior variance Σ_0. For this example (with three coefficients) B_0 is a 3×1 vector, while Σ_0 is 3×3 matrix with each diagonal element specifying the prior variance of the corresponding element of B_0.

We set a normal prior for the serial correlation coefficient ρ

$$p(\rho) \tilde{} N\left(\rho^0, \Sigma_\rho\right) \tag{3.17}$$

We set an inverse Gamma prior for σ^2 and set the prior degrees of freedom T_0 and the prior scale matrix θ_0 (see equation 3.18).

$$p\left(\sigma^2\right) \tilde{} \Gamma^{-1}\left(\frac{T_0}{2}, \frac{\theta_0}{2}\right) \tag{3.18}$$

To initialise the Gibbs sampler we need a starting value for σ^2 and ρ. In this example we will assume that the starting value for $\sigma^2 = \sigma^2_{OLS}$ where σ^2_{OLS} is the OLS estimate of σ^2. We assume that the starting value for $\rho = 0$.

Step 2 Given a value for σ^2 and ρ we sample from the conditional posterior distribution of B. As discussed above, this is done by first transforming the dependent and independent variables in the model to remove serial correlation. Once this is done we are back to the standard linear regression framework. We create the following transformed variables

$$\begin{aligned} Y_t^* &= Y_t - \rho Y_{t-1} \\ X_t^* &= X_t - \rho X_{t-1} \end{aligned}$$

where X_t^* represent the right hand side variables in our AR model. The conditional distribution of the regression coefficients is then given as

$$H\left(B \backslash \sigma^2, \rho, Y_t\right) \tilde{} N\left(M^*, V^*\right) \tag{3.19}$$

where

$$\begin{aligned} \underset{(3\times1)}{M^*} &= \left(\Sigma_0^{-1} + \frac{1}{\sigma^2} X_t^{*\prime} X_t^*\right)^{-1} \left(\Sigma_0^{-1} B_0 + \frac{1}{\sigma^2} X_t^{*\prime} Y_t^*\right) \\ \underset{(3\times3)}{V^*} &= \left(\Sigma_0^{-1} + \frac{1}{\sigma^2} X_t^{*\prime} X *_t\right)^{-1} \end{aligned} \tag{3.20}$$

Note that the mean and variance in equation 3.20 is identical to the expressions in equation 3.7. We have simply replaced the dependent and independent variables with our transformed data.

Step 3 Conditional on σ^2 and B we sample from the conditional distribution of ρ. Given the previous draw of B we can calculate the model residuals $v_t = Y_t - (\alpha + B_1 Y_{t-1} + B_2 Y_{t-2})$ and treat the equation $v_t = \rho v_{t-1} + \varepsilon_t, \varepsilon_t \sim N(0, \sigma^2)$ as an AR(1) model in v_t. Therefore, the conditional distribution for ρ is simply a normal distribution with the mean and variance derived in section 2.1. That is, the conditional distribution is

$$H\left(\rho \backslash \sigma^2, B, Y_t\right) \tilde{} N\left(\rho^*, z^*\right) \tag{3.21}$$

where

$$\begin{aligned} \underset{(1\times1)}{\rho^*} &= \left(\Sigma_\rho^{-1} + \frac{1}{\sigma^2} x_t' x_t\right)^{-1} \left(\Sigma_\rho^{-1} \rho^0 + \frac{1}{\sigma^2} x_t' y_t\right) \\ \underset{(1\times1)}{z^*} &= \left(\Sigma_\rho^{-1} + \frac{1}{\sigma^2} x_t' x_t\right)^{-1} \end{aligned} \tag{3.22}$$

where $y_t = v_t$ and $x_t = v_{t-1}$. With a value for ρ^* and z^* in hand, we simply draw ρ from the normal distribution with this mean and variance

$$\underset{(1\times1)}{\rho^1} = \underset{(1\times1)}{\rho^*} + \left[\underset{(1\times1)}{\bar{\rho}} \times \underset{(1\times1)}{\left(z^*\right)^{1/2}}\right]$$

where $\bar{\rho}$ is a draw from the standard normal distribution.

Step 4 Given a draw for B and ρ we draw σ^2 form its conditional posterior distribution. As shown in section 2.2 the conditional posterior distribution for σ^2 is inverse Gamma

$$H\left(\sigma^2 \backslash B, Y_t\right) \tilde{} \Gamma^{-1}\left(\frac{T_1}{2}, \frac{\theta_1}{2}\right) \tag{3.23}$$

where

$$\begin{aligned} T_1 &= T_0 + T \\ \theta_1 &= \theta_0 + \left(Y_t^* - B^1 X_t^*\right)' \left(Y_t^* - B^1 X_t^*\right) \end{aligned} \tag{3.24}$$

Note that the term $\left(\left(Y_t^* - B^1 X_t^*\right)' \left(Y_t^* - B^1 X_t^*\right)\right)$ is calculated using the *iid* residuals $Y_t^* - B^1 X_t^*$ (where B^1 is the previous draw of the coefficient vector).

```
clear
addpath('functions');
%an AR 2 model for US inflation with autocorrelated AR(1) disturbances
%load inflation data
Y=xlsread('\data\inflation.xls');
T=rows(Y);
X=[ones(T,1) lag0(Y,1) lag0(Y,2)];
%remove missing obs
Y=Y(3:end);
X=X(3:end,:);
T=rows(X);
%step 1 set priors and starting values
%priors for B
```

$$\underbrace{\begin{pmatrix} \alpha^0 \\ B_1^0 \\ B_2^0 \end{pmatrix}}_{B_0}, \underbrace{\begin{pmatrix} \Sigma_\alpha & 0 & 0 \\ 0 & \Sigma_{B1} & 0 \\ 0 & 0 & \Sigma_{B2} \end{pmatrix}}_{\Sigma_0}$$

```
B0=[0;0;0];
Sigma0=eye(3);
%priors for sigma2
T0=1;
```

$$p(\sigma^2) \sim \Gamma^{-1}\left(\frac{T_0}{2}, \frac{\theta_0}{2}\right)$$

```
D0=0.1;
%priors for rho
```

$$\rho^0$$

```
rho0=0;
```

$$\Sigma_\rho$$

```
Sigma0r=1;
%starting values
B=B0;
rho=rho0;
sigma2=1;
reps=15000;
burn=12000;
out1=[]; %will save the inflation forecast
out2=[];
out3=[];
for i=1:reps
%step 2 Sample B conditional on sigma N(M*,V*)
%remove serial correlation
```

$$Y_t^* = Y_t - \rho Y_{t-1}$$

```
ystar=Y-lag0(Y,1)*rho;
```

$$X_t^* = X_t - \rho X_{t-1}$$

```
xstar=X-lag0(X,1)*rho;
ystar=ystar(2:end,:);
xstar=xstar(2:end,:);

M=inv(inv(Sigma0)+(1/sigma2)*(xstar'*xstar))*(inv(Sigma0)*B0+(1/sigma2)*
xstar'*ystar);
```

$$\underset{(3\times1)}{M^*} = \left(\Sigma_0^{-1} + \frac{1}{\sigma^2} X_t^{*\prime} X_t^*\right)^{-1} \left(\Sigma_0^{-1} B_0 + \frac{1}{\sigma^2} X_t^{*\prime} Y_t^*\right)$$

FIGURE 10. Example 3: Matlab code

Step 5 Repeat steps 2 and 4 M times to obtain $B^1 ... B^M$, $\rho^1 ... \rho^M$ and $\left(\sigma^2\right)^1 .... \left(\sigma^2\right)^M$. The last H values of B, ρ and σ^2 from these iterations is used to form the empirical distribution of these parameters. This example shows that we reduce a relatively complicated model into three steps, each of which are simple and based on the linear regression framework. As seen in later chapters, Gibbs sampling will operate in exactly the same way in more complicated models–i.e. by breaking the problem down into smaller simpler steps.

3.6. Gibbs Sampling for a linear regression with serial correlation in Matlab (example3.m). The matlab code for this example is a simple extension of example1.m and shown in figures 10, 11 and 12. Note that the

```
39 V=inv(inv(Sigma0)+(1/sigma2)*(xstar'*xstar));
```

$$V^* = \left(\Sigma_0^{-1} + \frac{1}{\sigma^2} X_t^{*\prime} X *_t\right)^{-1}$$
$$(3\times 3)$$

```
40 chck=-1;
41 while chck<0
42 B=M+(randn(1,3)*chol(V))';
43 b=[B(2) B(3);1   0];
44 ee=max(abs(eig(b)));
45 if ee<=1
46     chck=1;
47 end
48 end
49 %step 3 compute rho
```

$$v_t = Y_t - (\alpha + B_1 Y_{t-1} + B_2 Y_{t-2})$$

```
50 y=Y-X*B;
51 x=lag0(y,1);
52 y=y(2:end);
53 x=x(2:end);

54
MM=inv(inv(Sigma0r)+(1/sigma2)*(x'*x))*(inv(Sigma0r)*rho0+(1/sigma2)*x'*
```

$$\rho^* = \left(\Sigma_\rho^{-1} + \frac{1}{\sigma^2} x_t' x_t\right)^{-1} \left(\Sigma_\rho^{-1} \rho^0 + \frac{1}{\sigma^2} x_t' y_t\right)$$
$$(1\times 1)$$

```
y);
55 VV=inv(inv(Sigma0r)+(1/sigma2)*(x'*x));
```

$$z^* = \left(\Sigma_\rho^{-1} + \frac{1}{\sigma^2} x_t' x_t\right)^{-1}$$
$$(1\times 1)$$

```
56 %draw rho but again ensure stationarity
57 chck=-1;
58 while chck<0
59 rho=MM+(randn(1,1)*chol(VV))';
60 ee=abs(rho);
61 if ee<=1
62     chck=1;
63 end
64 end
65 %step 3 sample sigma2 conditional on B from IG(T1,D1);
66 %compute residuals
```

$$Y_t^* - B^1 X_t^*$$

```
67 resids=ystar-xstar*B;
68 %compute posterior df and scale matrix
69 T1=T0+T;
70 D1=D0+resids'*resids;
71 %draw from IG
72 z0=randn(T1,1);
73 z0z0=z0'*z0;
74 sigma2=D1/z0z0;
75 if i>burn
76     %compute forecast for 12 quarters
77     yhat=zeros(14,1);
78     vhat=zeros(14,1);
79     yhat(1:2)=Y(end-1:end); %starting values
80         cfactor=sqrt(sigma2);  %standard deviation of the shocks
81     for m=3:14
82         vhat(m)=vhat(m-1)*rho+randn(1,1)*cfactor;
83         yhat(m)=[1 yhat(m-1) yhat(m-2)]*B+vhat(m);
```

FIGURE 11. Example 3: Matlab code (continued)

underlying data is exactly as before. This is loaded and lags etc created using the commands from lines 5 to 11. Lines 14 and 15 set the prior mean and variance for B for lines 17 and lines 18 sets the prior scale parameter and degrees of freedom for the inverse Gamma prior for σ^2. Lines 20 and 21 set the mean and variance for the normal prior for ρ, i.e. $p(\rho)\tilde{}N\left(\rho^0, \Sigma_\rho\right)$ Lines 23 to 25 set starting values for the parameters. The first step of the Gibbs sampling algorithm starts on line 34 and 35 where we create $Y_t^* = Y_t - \rho Y_{t-1}, X_t^* = X_t - \rho X_{t-1}$, the data transformed to remove serial correlation. Lines 38 and 39 calculate the mean and the variance of the conditional distribution of B using this tranformed data. As in the previous example, lines 40 to 48 draw B from its conditional distribution, but ensure that the draw is stable. Line 50 calculates the (serially correlated) residuals $v_t = Y_t - (\alpha + B_1 Y_{t-1} + B_2 Y_{t-2})$

```
84
85      end
86      %save
87      out1=[out1 [Y;yhat(3:end)]];
88      out2=[out2;B'];
89      out3=[out3;rho];
90  end
91  end
92  figure(1)
93  TT=1947.25:0.25:2012.5;
94  out2x=prctile(out1',[10 20 30 40 50 60 70 80 90])';
95  plot(TT,out2x);
96  xlim([2000 2013])
```

Published with MATLAB® 7.9

FIGURE 12. Example 3: Matlab code (continued)

using the previous draw of α, B_1 and B_2 and lines 50 and 51 create $y_t = v_t$ and $x_t = v_{t-1}$. Line 54 calculates the mean of the conditional distribution of ρ, $\underset{(1\times1)}{\rho^*} = \left(\Sigma_\rho^{-1} + \frac{1}{\sigma^2}x_t'x_t\right)^{-1}\left(\Sigma_\rho^{-1}\rho^0 + \frac{1}{\sigma^2}x_t'y_t\right)$ while line 55 calculates the variance of the conditional distribution $\underset{(1\times1)}{z^*} = \left(\Sigma_\rho^{-1} + \frac{1}{\sigma^2}x_t'x_t\right)^{-1}$. Line 59 draws ρ from the normal distribution using $\underset{(1\times1)}{\rho^1} = \underset{(1\times1)}{\rho^*} + \left[\underset{(1\times1)}{\bar{\rho}} \times \underset{(1\times1)}{(z^*)^{1/2}}\right]$ and the while loop ensures that ρ is less than or equal to 1 in absolute value. Line 67 calculates the serially uncorrelated residuals $Y_t^* - B^1 X_t^*$. These are used on lines 69 to 74 to draw σ^2 from

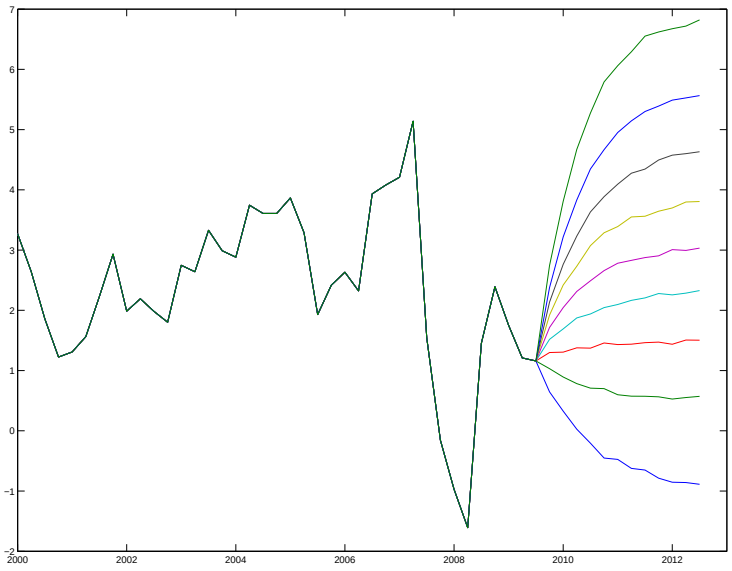

FIGURE 13. The distribution of the inflation forecast using example3.m.

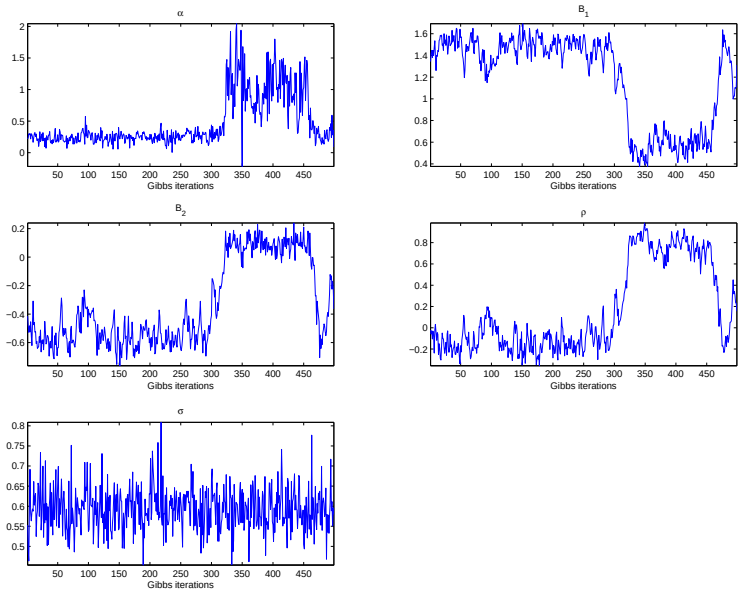

FIGURE 14. Sequence of retained Gibbs draws for the AR(2) model with serial correlation using 500 iterations

the inverse Gamma distribution. After the burn-in stage, the code computes the forecast from this AR(2) model with serial correlation. Line 82 projects forward the equation for the error term i.e. $v_{t+i} = \rho v_{t+i-1} + \sigma v^*$ where v^* is a standard normal shock. Line 83 calculates the projected value of inflation given v_{t+i}. This is done for each retained draw of the Gibbs sampler with the results (along with actual data) stored in the matrix out1 (line 87). The resulting distribution of the forecast is seen in 13.

3.7. Convergence of the Gibbs sampler. A question we have ignored so far is: How many draws of the Gibbs sampling algorithm do we need before we can be confident that the draws from the conditional posterior distributions have converged to the marginal posterior distribution? Generally researchers proceed in two steps

- Choose a minimum number of draws M and run the Gibbs sampler
- Check if the algorithm has converged (using the procedures introduced below). If there is insufficient evidence for convergence, increase M and try again.

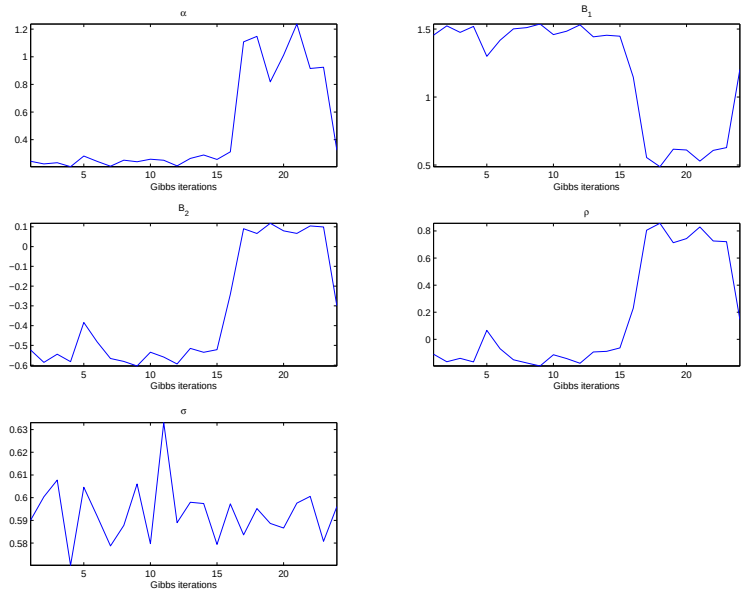

FIGURE 15. Recursive means of the retained Gibbs draws for the AR(2) model with serial correlation using 500 iterations

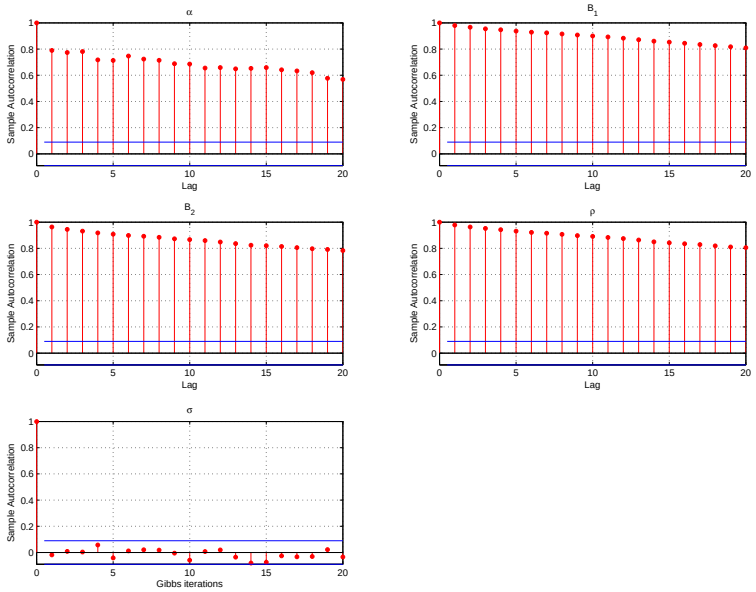

FIGURE 16. Autocorrelation of the retained Gibbs draws for the AR(2) model with serial correlation using 500 iterations

The simplest way to check convergence is to examine the sequence of retained draws. If the Gibbs sampler has converged to the target distibution, then the retained draws should fluctuate randomly around a stationary mean and not display any trend. This visual inspection is usually easier if one plots the recursive mean of the retained draws. If the Gibbs sampler has converged, then the recursive mean should show little fluctuation. A related method to examine convergence is plot the autocorrelation of the retained draws. If convergence has occurred, the sequence of draws should display little autocorrelation (i.e. they should be fluctuating randomly around a stationary mean).

In order to illustrate these ideas, we plot the sequence of retained draws, the recursive means of those draws and the autocorrelation functions of the retained draws for the parameters of the model examined in section 3.6. In particular, we estimate the AR(2) model with serial correlation using 500 Gibbs iterations (using the file example3.m) and retain all of these draws. Figures 14, 15 and 16 examine the convergence of the model. Figures 14 and 15 clearly show that the Gibbs draws are not stationary with the recursive mean for α, B_1, B_2 and ρ showing a large change

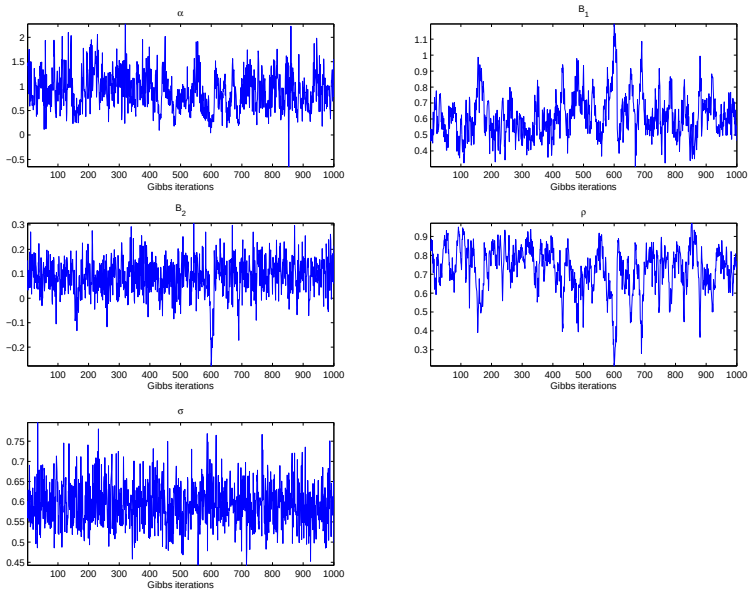

FIGURE 17. Sequence of retained Gibbs draws for the AR(2) model with serial correlation using 25000 iterations

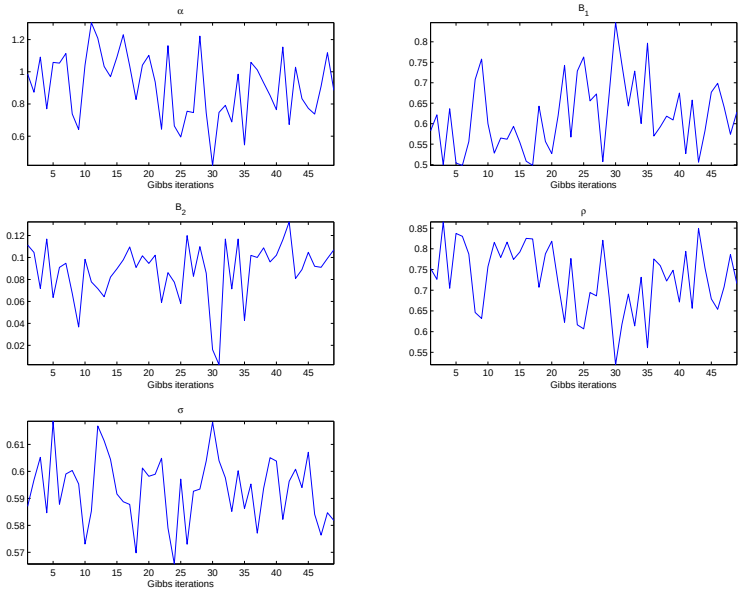

FIGURE 18. Recursive means of retained Gibbs draws for the AR(2) model with serial correlation using 25000 iterations

after 300 iterations (but σ^2 appears to have converged with the draws fluctuating around a stationary mean). This also shows up in the autocorrelation functions, with the autocorrelation high for α, B_1, B_2 and ρ. These figures can be produced using the file example4.m. These results would indicate that a higher number of Gibbs iterations are required. Figures 17, 18 and 19 plot the same objects when 25000 Gibbs iterations are used (with 24000 as the number of burn-in iterations). The sequence of retained draws and the recursive means appear substantially more stable. The autocorrelations for α, B_1, B_2 and ρ decay much faster in figure 19.

These graphical methods to assess convergence are widely used in applied work. A more formal test of convergence has been proposed by Geweke (1991). The intuition behind this test is related to the idea behind the recursive mean plot: If the Gibbs sampler has converged then the mean over different sub-samples of the retained draws should be similar. Geweke (1991) suggests the following procedure:

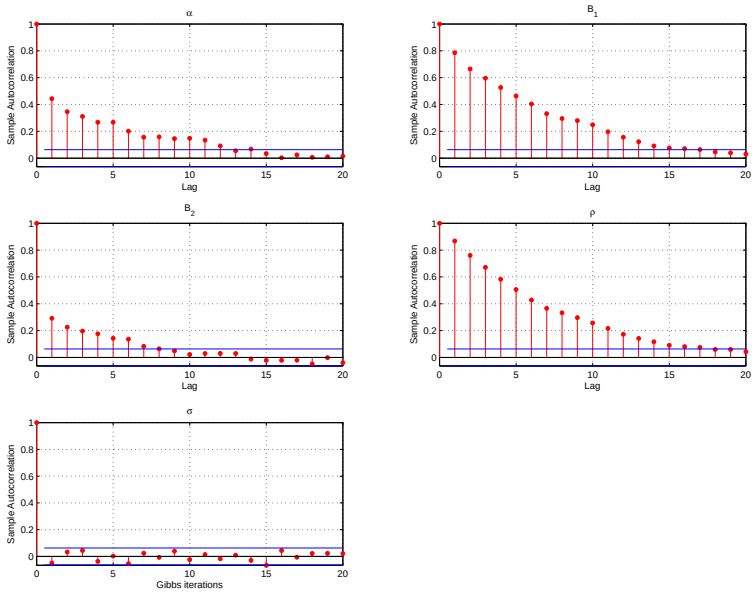

FIGURE 19. Autocorrelation of retained Gibbs draws for the AR(2) model with serial correlation using 25000 iterations

(1) Divide the retained Gibbs draws of the model parameters θ into two subsamples N_1 N_2 where Geweke (1991) recommends $N_1 = 0.1N, N_2 = 0.5N$ where N denotes the total number of retained draws.
(2) Compute averages $M_1 = \sum_{i=1}^{N_1} \frac{\theta_i}{N_1}$ and $M_2 = \sum_{i=N_2+1}^{N} \frac{\theta_i}{N_2}$
(3) Compute the asymptotic variance $\frac{S_1(0)}{N_1}$ and $\frac{S_2(0)}{N_2}$ where $S(\varpi)$ is the spectral density at frequency ϖ. Note that this estimate of the variance takes into account the possibility that the Gibbs sequence may be autocorrelated. For a description of spectral analysis see Hamilton (1994) and Canova (2007).
(4) Then the test statistic

$$Z = \frac{M_1 - M_2}{\sqrt{\frac{S_1(0)}{N_1} + \frac{S_2(0)}{N_2}}} \tag{3.25}$$

is asymptotically distributed as $N(0,1)$. Large values of this test statistic indicate a significant difference in the mean across the retained draws and suggests that one should increase the number of initial Gibbs iterations (i.e. increase the number of burn-in draws).

Geweke (1991) suggests a related statistic to judge the efficiency of the Gibbs sampler and to gauge the total number of Gibbs iterations to be used. The intuition behind this measure of relative numerical efficiency (RNE) is as follows. Suppose one could take iid draws of $\theta_i \in \{\theta_1, \theta_2 .... \theta_N\}$ directly from the posterior. Then the variance of the posterior mean $E(\theta_i) = \frac{1}{N}\sum_i \theta_i$ is given by

$$\begin{aligned} VAR(E(\theta_i)) &= \frac{1}{N^2}VAR(\theta_1) + \frac{1}{N^2}VAR(\theta_2) + ...\frac{1}{N^2}VAR(\theta_N) \\ &= VAR(\theta_i)/N \end{aligned}$$

However, in practice one uses the Gibbs sampler to approximate draws from the posterior. These Gibbs draws are likely to be autocorrelated and a measure of their variance which takes this into account is $S(0)/N$. Thus a measure of the RNE is

$$RNE = \frac{\widehat{VAR(\theta_i)}}{S(0)} \tag{3.26}$$

where $\widehat{VAR(\theta_i)}$ is the sample variance of the Gibbs draws $\theta_1, \theta_2 .... \theta_N$. If the Gibbs sampler has converged then RNE should be close to 1 as the variance of the iid draws $\widehat{VAR(\theta_i)}$ should be similar to the measure of the variance that takes any possible autocorrelation into account.

The file example5.m illustrates the calculation of the statistics in equation 3.25 and 3.26.

4. Further Reading

- An intuitive description of the Gibbs sampling algorithm for the linear regression model can be found in Kim and Nelson (1999) Chapter 7. Gauss codes for the examples in Kim and Nelson (1999) are available at http://www.econ.washington.edu/user/cnelson/markov/prgmlist.htm.

- A more formal treatment of the linear regression model from a Bayesian perspective can be found in Koop (2003), Chapters 2, 3 and 4.
- The appendix in Zellner (1971) provides a detailed description of the Inverse Gamma and Gamma distributions. See Bauwens *et al.* (1999) for a detailed description of algorithms to draw from these distributions.

5. Appendix: Calculating the marginal likelihood for the linear regression model using the Gibbs sampler.

Consider the following linear regression model

$$Y_t = BX_t + v_t, v_t \tilde{} N(0, \sigma^2)$$

The prior distributions are assumed to be

$$P(B) \tilde{} N(B_0, \Sigma_0)$$
$$P\left(\sigma^2\right) \tilde{} IG(V_0, T_0)$$

The posterior distribution of the model parameters $\Phi = B, \sigma^2$ is defined via the Bayes rule

$$H\left(\Phi \backslash Y\right) = \frac{F\left(Y \backslash \Phi\right) \times P\left(\Phi\right)}{F\left(Y\right)} \tag{5.1}$$

where $F\left(Y \backslash \Phi\right) = \left(2\pi\sigma^2\right)^{-\frac{T}{2}} \exp\left(-\frac{1}{2\sigma^2}\left(Y_t - BX_t\right)'\left(Y_t - BX_t\right)\right)$ is the likelihood function, $P\left(\Phi\right)$ is the joint prior distribution while $F\left(Y\right)$ is the marginal likelihood that we want to compute. Chib (1995) suggests computing the marginal likelihood by re-arranging equation 5.1. Note that in logs we can re-write equation 5.1 as

$$\ln F\left(Y\right) = \ln F\left(Y \backslash \Phi\right) + \ln P\left(\Phi\right) - \ln H\left(\Phi \backslash Y\right) \tag{5.2}$$

Note that equation 5.2 can be evaluated at any value of the parameters Φ to calculate $\ln F\left(Y\right)$. In practice a high density point Φ^* such as the posterior mean or posterior mode is used.

The first two terms on the right hand side of equation 9.3 are easy to evaluate at Φ^*. The first term is the log likelihood function. The second term is the joint prior which is the product of a normal density for the coefficients and an inverse Gamma density for the variance (see example below). Evaluating the third term $\ln H\left(\Phi^* \backslash Y\right)$ is more complicated as the posterior distribution is generally not known in closed form. Chib (1995) shows how this term can be evaluated using the output from the Gibbs sampling algorithm used to approximate the posterior distribution for Φ. Recall that $H\left(\Phi^* \backslash Y\right) = H\left(B^*, \sigma^{2*}\right)$ where have dropped the conditioning on y on the right hand side for simplicity. The marginal, conditional decomposition of this distribution is

$$H\left(B^*, \sigma^{2*}\right) = H\left(B^* \backslash \sigma^{2*}\right) \times H\left(\sigma^{2*}\right) \tag{5.3}$$

The first term $H\left(B^* \backslash \sigma^{2*}\right)$ is the conditional posterior distribution for the regression coefficients. Recall that this a normal distribution with mean and variance given by

$$\begin{aligned} M^* &= \left(\Sigma_0^{-1} + \frac{1}{\sigma^{2*}} X_t' X_t\right)^{-1} \left(\Sigma_0^{-1} B_0 + \frac{1}{\sigma^{2*}} X_t' Y_t\right) \\ V^* &= \left(\Sigma_0^{-1} + \frac{1}{\sigma^{2*}} X_t' X_t\right)^{-1} \end{aligned}$$

and therefore can be easily evaluated at B^* and σ^{2*}.

The second term in equation 5.3 $H\left(\sigma^{2*}\right)$ can be evaluated using the weak law of large numbers (see Koop (2003) Appendix B). That is

$$H\left(\sigma^{2*}\right) \approx \frac{1}{S} \sum_{s=1}^{S} H\left(\sigma^{2*} \backslash B_s\right)$$

where B_s denotes $s = 1, 2....S$ draws of the Gibbs sampler. Note that the conditional distribution is simply the Inverse Gamma distribution derived for section 2.2 above.

The marginal likelihood is then given by

$$\ln F\left(Y\right) = \ln F\left(Y \backslash \Phi\right) + \ln P\left(\Phi\right) - \ln H\left(B^* \backslash \sigma^{2*}\right) - \ln H\left(\sigma^{2*}\right) \tag{5.4}$$

As an example we consider the following linear regression model based on 100 artificial observations

$$y_t = 1 + 0.5x_t + v_t, VAR(v_t) = 0.2$$

where $x_t \tilde{} N(0, 1)$. We assume a natural conjugate prior of the form $P(B \backslash \sigma^2) \tilde{} N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, 4\sigma^2 \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}\right)$ and $P\left(\sigma^2\right) \tilde{} IG(2.5, 3)$.

The matlab code for this example is shown in figures 20 and 21. The code on Lines 5 to 9 generates the artificial data. We set the priors on lines 11 to 14. On line 16 we calculate the marginal likelihood for this model analytically using the formula on page 41 in Koop (2003). We can now compare this estimate with the estimate produced using Chib's method. The Gibbs sampler used to estimate the model is coded on lines 19 to 43. Line 46 calculates

```
1 clear;
2 clc
3 addpath('functions')
4 %generate artificial data
5 T=100;
6 X=[ones(T,1) randn(T,1)];
7 btrue=[1;0.5];
8 sigmatrue=0.2;
9 Y=X*btrue+randn(T,1)*sqrt(sigmatrue);
10 %set priors
11 T0=3;
12 D0=2.5;
13 B0=zeros(2,1);
14 Sigma0=eye(2)*(4);
15 %analytical computation of the marginal likelihood
16 mlika=mlikols(B0,Sigma0,T0,D0,Y,X);
17 disp('Analytical log Marginal Likelihood');
18 disp(log(mlika));
19 sigma2=1;
20 reps=15000;   %total numbers of Gibbs iterations
21 burn=4000;   %percent of burn-in iterations
22 out1=[];
23 out2=[];
24 for i=1:reps
25 %Sample B conditional on sigma N(M*,V*)
26 M=inv(inv(Sigma0)+(1/sigma2)*(X'*X))*(inv(Sigma0)*B0+(1/sigma2)*X'*Y);
27 V=inv(inv(Sigma0)+(1/sigma2)*(X'*X));
28 B=M+(randn(1,2)*chol(V))';
29 %sample sigma2 conditional on B from IG(T1,D1);
30 %compute residuals
31 resids=Y-X*B;
32 %compute posterior df and scale matrix
33 T1=T0+T;
34 D1=D0+resids'*resids;
35 %draw from IG
36 z0=randn(T1,1);
37 z0z0=z0'*z0;
38 sigma2=D1/z0z0;
39 if i>burn
40     out1=[out1;B'];
41     out2=[out2;sigma2];
42 end
43 end
44 %calculate the marginal likelihood using Chib's method
45 %posterior mean
46 bstar=mean(out1);
47 sigmastar=mean(out2);
48 Hstar=mean(1./out2);
49 % Step1 evaluate the prior distributions at the posterior mean
50 %P(B)~N(B0,Sigma0) in logs
51 Pb=log(mvnpdf(bstar',B0,Sigma0));
52 %P(1/sigma2)~Gamma(D0,T0) in logs
53 PH=gampdf1(T0,D0,Hstar);
54 %Step 2 evaluate the log likelihood
55 loglik=-(T/2)*log(2*pi*sigmastar)-0.5*(((Y-X*bstar')'*(Y-
X*bstar'))/sigmastar);
56 %step 3 evaluate the posterior density
57 %H(bstar,sigmastar)=H(bstar\sigmastar)*H(sigmastar)
58 %step 3a H(bstar\sigmastar)~N(M1,V1) in logs
```

FIGURE 20. Matlab code for calculating the marginal likelihood

the posterior mean of the coefficients, line 47 calculates the posterior mean of the variance while line 48 calculates the posterior mean of $1/\sigma^2$. For computational convenience, when considering the prior $\ln P(\Phi)$ and the posterior distribution $\ln H(\Phi \backslash Y)$ in the expression for the marginal likelihood (see equation 5.2) we consider the precision $1/\sigma^2$ and use the Gamma distribution. This allows us to use built in matlab functions to evaluate the Gamma PDF. On line 51, we evaluate the log of the prior distribution of the VAR coefficients $P(B \backslash \sigma^2)$ at the posterior mean. Line 53 evaluates the Gamma posterior for the precision. The function gampdf1 converts the two parameters of the distribution: the degrees of freedom T_0 and scale parameter D_0 into the parameters $A = T_0/2$ and $B = 2/D_0$ as expected by the parameterisation of the Gamma distribution used by Matlab in its built in function gampdf.

```
59
M1=inv(inv(Sigma0)+(1/sigmastar)*(X'*X))*(inv(Sigma0)*B0+(1/sigmastar)*X'
*Y);
60 V1=inv(inv(Sigma0)+(1/sigmastar)*(X'*X));
61 H1=log(mvnpdf(bstar',M1,V1));
62 %step 3b evaluate H(sigmastar) using the Gibbs draws
63 H2log=[];
64 for i=1:size(out1,1)
65     bgibbs=out1(i,:)';
66     res=Y-X*bgibbs;
67    H2i=gampdf1(rows(res)+T0,(res'*res)+D0,Hstar);
68     H2log=[H2log;H2i];
69 end
70 %take exponential and mean in a way that prevents underflow
71 factor=max(H2log);
72 H2exp=exp(H2log-factor);
73 H2mean=log(mean(H2exp))+factor;
74 %calculate marginal likelihood
75 mlik=loglik+Pb+PH-H1-H2mean;
76 disp('Chib log Marginal Likelihood');
77 disp(mlik);
```

Published with MATLAB® 7.9

FIGURE 21. Matlab code for calculating the marginal likelihood continued

Line 55 evaluates the log likelihood at the posterior mean. Lines 56 to 61 evaluate the term $H\left(B^*\backslash\sigma^{2*}\right)$ in the factorisation of the posterior $H\left(B^*,1/\sigma^{2*}\right)=H\left(B^*\backslash\sigma^{2*}\right)\times H\left(1/\sigma^{2*}\right)$. Lines 63 to 69 evaluate the term $H\left(1/\sigma^{2*}\right)$. Each iteration in the loop evaluates $H\left(1/\sigma^{2*}\backslash B_j\right)$. Note that this is simply the Gamma distribution with degrees of freedom T_0+T and scale parameter $D_0+v_t'v_t$ where the residuals v_t are calculated using each Gibbs draw of the regression coefficients B_j.T_0 and D_0 denote the prior degrees of freedom and prior scale parameter respectively. Line 73 constructs $H\left(1/\sigma^{2*}\right)\approx\frac{1}{J}\sum_{j=1}^{J}H\left(1/\sigma^{2*}\backslash B_J\right)$. The marginal likelihood is calculated using equation 5.2 on line 75 of the code.

CHAPTER 2

Gibbs Sampling for Vector Autoregressions

This chapter introduces Bayesian simulation methods for Vector Autoregressions (VARs). The estimation of these models typically involves a large number of parameters. As a consequence, estimates of objects of interest such as impulse response functions and forecasts can become imprecise in large scale models. By incorporating prior information into the estimation process, the estimates obtained using Bayesian methods are generally more precise than those obtained using the standard classical approach. In addition, bayesian simulation methods such as Gibbs sampling provide an efficient way not only to obtain point estimates but also to characterise the uncertainty around those point estimates. Therefore we focus on estimation of VARs *via Gibbs sampling* in this chapter.

Note, however, that under certain prior distributions, analytical expressions exist for the marginal posterior distribution of the VAR parameters. A more general treatment of Bayesian VARs can be found in Canova (2007) amongst others. See http://apps.eui.eu/Personal/Canova/Courses.html for F.Canova's BVAR code.

This chapter focusses on two key issues

- It states the conditional posterior distributions of the VAR parameters required for Gibbs sampling and discussed the Gibbs sampling algorithm for VARs
- We go through the practical details of setting different type of priors for VAR parameters
- We focus on implementation of Gibbs sampling for VARs in Matlab.
- We discuss how to estimate structural VARs with sign restrictions using Matlab.

1. The Conditional posterior distribution of the VAR parameters and the Gibbs sampling algorithm

Consider the following VAR(p) model

$$\begin{aligned} Y_t &= c + B_1 Y_{t-1} + B_2 Y_{t-2} ... B_P Y_{t-p} + v_t \\ E\left(v_t' v_s\right) &= \Sigma \text{ if } t = s \\ E\left(v_t' v_s\right) &= 0 \text{ if } t \neq s \\ E\left(v_t\right) &= 0 \end{aligned} \tag{1.1}$$

where Y_t is a $T \times N$ matrix of endogenous variables, c denotes a constant term. The VAR can be written compactly as

$$Y_t = X_t B + v_t \tag{1.2}$$

with $X_t = \{c_i, Y_{it-1}, Y_{it-2} ..., Y_{it-p}\}$.Note that as each equation in the VAR has *identical* regressors, it can be re-written as

$$y = \left(I_N \otimes X\right) b + V \tag{1.3}$$

where $y = vec(Y_t)$ and $b = vec(B)$ and $V = vec(v_t)$.

Assume that the prior for the VAR coefficients b is normal and given by

$$p(b) \tilde{} N\left(\tilde{b}_0, H\right) \tag{1.4}$$

where $\tilde{b}_0$ is a $(N \times (N \times P + 1)) \times 1$ vector which denotes the prior mean while H is a is a $[N \times (N \times P + 1)] \times [N \times (N \times P + 1)]$ matrix where the diagonal elements denote the variance of the prior. We discuss different ways of setting $\tilde{b}_0$ and H in detail below.

It can be shown that the *posterior distribution of the VAR coefficients* conditional on Σ is normal (see Kadiyala and Karlsson (1997)) . That is the conditional posterior for the coefficients is given by $H\left(b \backslash \Sigma, Y_t\right) \tilde{} N\left(M^*, V^*\right)$ where

$$\begin{aligned} M^* &= \left(H^{-1} + \Sigma^{-1} \otimes X_t' X_t\right)^{-1} \left(H^{-1} \tilde{b}_0 + \Sigma^{-1} \otimes X_t' X_t \hat{b}\right) \\ V^* &= \left(H^{-1} + \Sigma^{-1} \otimes X_t' X_t\right)^{-1} \end{aligned} \tag{1.5}$$

where $\hat{b}$ is a $(N \times (N \times P + 1)) \times 1$ vector which denotes the OLS estimates of the VAR coefficients in vectorised format $\hat{b} = vec\left(\left(X_t' X_t\right)^{-1} \left(X_t' Y_t\right)\right)$. The format of the conditional posterior mean in equation 1.5 is very similar to that discussed for the linear regression model (see section 2.1 in the previous chapter). That is the mean of the conditional posterior distribution is a weighted average of the OLS estimator $\hat{b}$ and the prior $\tilde{b}_0$ with the weights given by the inverse of the variance of each ($\Sigma^{-1} \otimes X_t' X_t$ is the inverse of $\hat{b}$ while H^{-1} is the inverse of the variance of the prior).

The conjugate prior for the VAR covariance matrix is an *inverse Wishart distribution* with prior scale matrix $\bar{S}$ and prior degrees of freedom α.

$$p(\Sigma) \tilde{} IW\left(\bar{S}, \alpha\right) \tag{1.6}$$

DEFINITION 3. *If Σ is a $n \times n$ positive definite matrix, it is distributed as an inverse Wishart with the following density $P(\Sigma) = k\frac{|H|^{v/2}}{|\Sigma|^{(v+n+1)/2}} \exp\left(-0.5tr\Sigma^{-1}H\right)$ where $k^{-1} = 2^{vn/2}\pi^{n(n-1)/4}\prod_{i=1}^{n}\Gamma\left[(v+1-i)/2\right]$, H is the scale matrix and v denotes the degrees of freedom. See Zellner (1971) pp395 for more details.*

Informally, one can think of the inverse Wishart distribution as a multivariate version of the inverse Gamma distribution introduced in the context of the linear regression model in the previous chapter. Given the prior in equation 1.6, the posterior for Σ conditional on b is also inverse Wishart $H(\Sigma \backslash b, Y_t) \tilde{} IW\left(\bar{\Sigma}, T+\alpha\right)$ where T is the sample size and

$$\bar{\Sigma} = \bar{S} + (Y_t - X_t B)'(Y_t - X_t B) \tag{1.7}$$

Note that B denotes the VAR coefficients reshaped into $(N \times P+1)$ by N matrix.

1.1. Gibbs sampling algorithm for the VAR model. The Gibbs sampling algorithm for the VAR model consists of the following steps:

Step 1 Set priors for the VAR coefficients and the covariance matrix. As discussed above, the prior for the VAR coefficients is normal and given by $p(b) \tilde{} N\left(\tilde{b}_0, H\right)$. The prior for the covariance matrix of the residuals Σ is inverse Wishart and given by $IW\left(\bar{S}, \alpha\right)$. Set a starting value for Σ (e.g. the OLS estimate of Σ).

Step 2 Sample the VAR coefficients from its conditional posterior distribution $H(b \backslash \Sigma, Y_t) \tilde{} N(M^*, V^*)$ where

$$\underset{(N\times(N\times P+1))\times 1}{M^*} = \left(H^{-1} + \Sigma^{-1} \otimes X_t'X_t\right)^{-1}\left(H^{-1}\tilde{b}_0 + \Sigma^{-1} \otimes X_t'X_t\hat{b}\right) \tag{1.8}$$

$$\underset{(N\times(N\times P+1))\times(N\times(N\times P+1))}{V^*} = \left(H^{-1} + \Sigma^{-1} \otimes X_t'X_t\right)^{-1} \tag{1.9}$$

Once M^* and V^* are calculated, the VAR coefficients are drawn from the normal distribution (see algorithm 1 in Chapter 1)

$$\underset{((N\times(N\times P+1))\times 1)}{b^1} = \underset{((N\times(N\times P+1))\times 1)}{M^*} + \left[\underset{(1\times(N\times(N\times P+1)))}{\bar{b}} \times \underset{(N\times(N\times P+1))\times(N\times(N\times P+1))}{(V^*)^{1/2}}\right] \tag{1.10}$$

Step 3 Draw Σ from its conditional distribution $H(\Sigma \backslash b, Y_t) \tilde{} IW\left(\bar{\Sigma}, T+\alpha\right)$ where $\bar{\Sigma} = \bar{S} + \left(Y_t - X_t B^1\right)'\left(Y_t - X_t B^1\right)$ where B^1 is the previous draw of the VAR coefficients reshaped into a matrix with dimensions $(N \times P+1) \times N$ so it is conformable with X_t.

ALGORITHM 3. *To draw a matrix $\hat{\Sigma}$ from the IW distribution with v degrees of freedom and scale parameter S, draw a matrix Z with dimensions $v \times n$, from the multivariate normal $N(0, S^{-1})$. Then the draw from the inverse Wishart distribution is given by the following transformation:*

$$\hat{\Sigma} = \left(\sum_{i=1}^{v} Z_i Z_i'\right)^{-1}$$

Step 3 (continued) With the parameters of inverse Wishart distribution in hand ($\bar{\Sigma} = \bar{S} + \left(Y_t - X_t B^1\right)'\left(Y_t - X_t B^1\right)$ and $T+\alpha$) one can use algorithm 3 to draw Σ from the inverse Wishart distribution.

Repeat Steps 2 to 3 M times to obtain $B^1 ... B^M$ and $(\Sigma)^1 .... (\Sigma)^M$. The last H values of B and Σ from these iterations is used to form the empirical distribution of these parameters. Note that the draws of the model parameters (after the burn-in period) are typically used to calculate forecasts or impulse response functions and build the distribution for these statistics of interest.

In general, the Gibbs sampling algorithm for VARs is very similar to that employed for the linear regression model in the previous chapter. The key difference turns out to be the fact that setting up the prior in the VAR model is a more structured process than the linear regression case.

We now turn to a few key prior distributions for VARs that have been proposed in the literature and the implementation of the Gibbs sampling algorithm in Matlab. To discuss the form of the priors we will use the following bi-variate VAR(2) model as an example:

$$\begin{pmatrix} y_t \\ x_t \end{pmatrix} = \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} + \begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{pmatrix}\begin{pmatrix} y_{t-1} \\ x_{t-1} \end{pmatrix} + \begin{pmatrix} d_{11} & d_{12} \\ d_{21} & d_{22} \end{pmatrix}\begin{pmatrix} y_{t-2} \\ x_{t-2} \end{pmatrix} + \begin{pmatrix} v_1 \\ v_2 \end{pmatrix} \tag{1.11}$$

where $var\begin{pmatrix} v_1 \\ v_2 \end{pmatrix} = \Sigma = \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{12} & \Sigma_{22} \end{pmatrix}$

2. The Minnesota prior

The Minnesota prior (named after its origins at the Federal Reserve Bank of Minnesota) incorporates the prior belief that the endogenous variables included in the VAR follow a random walk process or an AR(1) process. In other words, the *mean* of the Minnesota prior for the VAR coefficients in equation 1.11 implies the following form for the VAR

$$\begin{pmatrix} y_t \\ x_t \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} + \begin{pmatrix} b_{11}^0 & 0 \\ 0 & b_{22}^0 \end{pmatrix} \begin{pmatrix} y_{t-1} \\ x_{t-1} \end{pmatrix} + \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} y_{t-2} \\ x_{t-2} \end{pmatrix} + \begin{pmatrix} v_1 \\ v_2 \end{pmatrix} \tag{2.1}$$

Equation 2.1 states that the Minnesota prior incorporates the belief that both y_t and x_t follow an AR(1) process or a random walk if $b_{11}^0 = b_{22}^0 = 1$. If y_t and x_t are stationary variables then it may be more realistic to incorporate the prior that they follow an AR(1) process. For this example, the mean of the Minnesota prior distribution for the VAR coefficients (i.e. $\tilde{b}_0$ from $p(b)\tilde{}N\left(\tilde{b}_0, H\right)$) is given by the vector

$$\tilde{b}_0 = \begin{pmatrix} 0 \\ b_{11}^0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ b_{22}^0 \\ 0 \\ 0 \end{pmatrix} \tag{2.2}$$

where the first five rows correspond to the coefficients for the first equation and the second five rows correspond to the coefficients for the second equation. The variance of the prior H is a set in a more structured manner (as compared to the examples in chapter 1) and is given by the following relations for the VAR coefficients b_{ij}

$$\begin{aligned} \left(\frac{\lambda_1}{l^{\lambda_3}}\right)^2 \; if \; i &= j \\ \left(\frac{\sigma_i \lambda_1 \lambda_2}{\sigma_j l^{\lambda_3}}\right)^2 \; if \; i &\neq j \end{aligned} \tag{2.3}$$

$$(\sigma_1 \lambda_4)^2 \text{ for the constant}$$

where i refers to the dependent variable in the i^{th} equation and j to the independent variables in that equation. Therefore, if $i = j$ then we are referring to the coefficients on the own lags of variable i. σ_i and σ_j are variances of error terms from AR regressions estimated via OLS using the variables in the VAR. The ratio of σ_i and σ_j in the formulas above controls for the possibility that variable i and j may have different scales. Note that l is the lag length. The $\lambda's$ are parameters set by the researcher that control the tightness of the prior:

- λ_1 controls the standard deviation of the prior on own lags. As $\lambda_1 \to 0$ $b_{11}, b_{22} \to b_{11}^0, b_{22}^0$ respectively and all other lags go to zero in our example VAR in equation 1.11.
- λ_2 controls the standard deviation of the prior on lags of variables other than the dependent variable i.e. b_{12}, b_{21} etc. As $\lambda_2 \to 0$ b_{ij}, d_{ij} go to zero. With $\lambda_2 = 1$ there is no distinction between lags of the dependent variable and other variables.
- λ_3 controls the degree to which coefficients on lags higher than 1 are likely to be zero. As λ_3 increases coefficients on higher lags are shrunk to zero more tightly.
- The prior variance on the constant is controlled by λ_4. As $\lambda_4 \to 0$ the constant terms are shrunk to zero.

It is instructive to look at how the prior variance matrix looks for our example VAR(2) in equation 1.11. This is shown below in equation 2.4

$$H = \begin{pmatrix} (\sigma_1\lambda_4)^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & (\lambda_1)^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \left(\frac{\sigma_1\lambda_1\lambda_2}{\sigma_2}\right)^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \left(\frac{\lambda_1}{2^{\lambda_3}}\right)^2 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \left(\frac{\sigma_1\lambda_1\lambda_2}{\sigma_2 2^{\lambda_3}}\right)^2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & (\sigma_2\lambda_4)^2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \left(\frac{\sigma_2\lambda_1\lambda_2}{\sigma_1}\right)^2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & (\lambda_1)^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \left(\frac{\sigma_2\lambda_1\lambda_2}{\sigma_1 2^{\lambda_3}}\right)^2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \left(\frac{\lambda_1}{2^{\lambda_3}}\right)^2 \end{pmatrix} \tag{2.4}$$

The matrix H in equation 2.4 is a 10×10 matrix, because for this example we have 10 total coefficients in the VAR model. The diagonal elements of the matrix H are the prior variances for each corresponding coefficient. Consider the the first five elements on the main diagonal correspond to the first equation the VAR model and is re-produced in equation 2.5.

$$\begin{pmatrix} (\sigma_1\lambda_4)^2 & 0 & 0 & 0 & 0 \\ 0 & (\lambda_1)^2 & 0 & 0 & 0 \\ 0 & 0 & \left(\frac{\sigma_1\lambda_1\lambda_2}{\sigma_2}\right)^2 & 0 & 0 \\ 0 & 0 & 0 & \left(\frac{\lambda_1}{2^{\lambda_3}}\right)^2 & 0 \\ 0 & 0 & 0 & 0 & \left(\frac{\sigma_1\lambda_1\lambda_2}{\sigma_2 2^{\lambda_3}}\right)^2 \end{pmatrix} \tag{2.5}$$

The first diagonal element $(\sigma_1\lambda_4)^2$ controls the prior on the constant term. The second element $(\lambda_1)^2$ controls the prior on b_{11} the coefficient on the first lag of y_t. Note that this element comes from the first expression in equation 2.3 $\left(\frac{\lambda_1}{l^{\lambda_3}}\right)^2$ with the lag length $l = 1$ as we are dealing with the first lag. The third diagonal element controls the prior on b_{12} the coefficient on the first lag of x_t in the equation for y_t. Note that this element comes from the second expression in equation 2.3 i.e. $\left(\frac{\sigma_i\lambda_1\lambda_2}{\sigma_j l^{\lambda_3}}\right)^2$ with $l = 1$. The third and the fourth diagonal elements control the prior on the coefficients d_{11} and d_{12} respectively (and again come from the first and second expression in equation 2.3 with $l = 2$).

Under a strict interpretation of the Minnesota prior, the covariance matrix of the residuals of the VAR Σ is assumed to be diagonal with the diagonal entries fixed using the error variances from AR regressions σ_i. Under this assumption, the mean of the posterior distribution for the coefficients is available in closed form. For the exact formula, see Kadiyala and Karlsson (1997) Table 1. However, it is common practice amongst some researchers to incorporate the Minnesota prior into the Gibbs sampling framework and draw Σ from the inverse Wishart distribution. We turn to the practical implementation of this algorithm next.

An important question concerns the values of the hyperparameters that control the priors. Canova (2007) pp 380 reports the following values for these parameters typically used in the literature.

$$\begin{aligned} \lambda_1 &= 0.2 \\ \lambda_2 &= 0.5 \\ \lambda_3 &= 1 \textit{ or } 2 \\ \lambda_4 &= 10^5 \end{aligned}$$

Some researchers set the value of these parameters by comparing forecast performance of the VAR across a range of values for these parameters. In addition, the marginal likelihood can be used to select the value of these hyperparameters. The appendix to this chapter shows how to use the procedure in Chib (1995) to calculate the marginal likelihood for a VAR model.

2.1. Gibbs sampling and the Minnesota prior. Matlab code. We consider the estimation of a bi-variate VAR(2) model using quarterly data on annual GDP growth and CPI inflation for the US from 1948Q2 to 2010Q4. We employ a Minnesota prior which incorporates the belief that both variables follow a random walk. Note that while annual CPI inflation may be non-stationary (and hence the random walk prior reasonable), annual GDP growth is likely to be less persistent. Hence one may want to consider incorporating the belief that this variable follows an AR(1) process in actual applications. Note that, we also incorporate a inverse Wishart prior for the covariance matrix and hence depart from the strict form of this model where the covariance matrix is fixed and diagonal. The

```
1 clear
2 addpath('functions');
3 % a bi-variate VAR with a Minnesota Prior and Gibbs Sampling
4 %load data
5 data=xlsread('\data\datain.xls'); %data for US GDP growth and
inflation 1948q1 2010q4
6 N=size(data,2);
7 L=2;   %number of lags in the VAR
8 Y=data;
9 X=[ones(size(Y,1),1) lag0(data,1) lag0(data,2) ];
10 Y=Y(3:end,:);
11 X=X(3:end,:);
12 T=rows(X);
13 %compute standard deviation of each series residual via an ols
regression
14 %to be used in setting the prior
15 %first variable
16 y=Y(:,1);
17 x=X(:,1:2);
18 b0=inv(x'*x)*(x'*y);
19 s1=sqrt(((y-x*b0)'*(y-x*b0))/(rows(y)-2));  %std of residual standard
error
20 %second variable
21 y=Y(:,2);
22 x=X(:,[1 3]);
23 b0=inv(x'*x)*(x'*y);
24 s2=sqrt(((y-x*b0)'*(y-x*b0))/(rows(y)-2));
25 %specify parameters of the minnesota prior
26 lamda1=1;
27 lamda2=1;
28 lamda3=1;
29 lamda4=1;
30 %specify the prior mean of the coefficients of the Two equations of
the VAR
31 B01=[0;1;0;0;0];
32 B02=[0;0;1;0;0];
33 B0=[B01;B02];
34 %Specify the prior variance of vec(B)
```

λ_1 controls the standard deviation of the prior on own lags.

λ_2 controls the standard deviation of the prior on lags of variables other than the dependent variable

λ_3 controls the degree to which coefficients on lags higher than 1 are likely to be zero

The prior variance on the constant is controlled by λ_4

$$\tilde{b}_0 = \begin{pmatrix} 0 \\ b_{11}^0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ b_{22}^0 \\ 0 \\ 0 \end{pmatrix}$$

FIGURE 1. Matlab code for example 1

model is estimated using the Gibbs sampling algorithm described in section 1.1. The code for this model is in the file example1.m in the subfolder chapter 2 under the folder code. The code is also shown in figures 1, 2 and 3. We now go through this code line by line.

Line 5 of the code loads the data for the two variables from an excel file and lines 8 and 9 prepare the matrices Y_t, X_t.Lines 16 to 24 compute σ_1 and σ_2 (to be used to form the Minnesota prior) using AR(1) regressions for each variable. In this example we use the full sample to compute these regressions. Some researchers use a pre-sample (or a training sample) to compute σ_1 and σ_2 and then estimate the VAR on the remaining data points. The argument for using a pre-sample is that the full sample should not really be used to set parameters that affect the prior.

```
35 H=zeros(10,10);
```

$$H = \begin{pmatrix} (\sigma_1\lambda_4)^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & (\lambda_1)^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \left(\frac{\sigma_1\lambda_1\lambda_2}{\sigma_2}\right)^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \left(\frac{\lambda_1}{2^{\lambda_3}}\right)^2 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \left(\frac{\sigma_1\lambda_1\lambda_2}{\sigma_2 2^{\lambda_3}}\right)^2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & (\sigma_2\lambda_4)^2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \left(\frac{\sigma_2\lambda_1\lambda_2}{\sigma_1}\right)^2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & (\lambda_1)^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \left(\frac{\sigma_2\lambda_1\lambda_2}{\sigma_1 2^{\lambda_3}}\right)^2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \left(\frac{\lambda_1}{2^{\lambda_3}}\right)^2 \end{pmatrix}$$

```
36 %for equation 1  of the VAR
37 H(1,1)=(s1*lamda4)^2;  %constant
38 H(2,2)=(lamda1)^2;     %own lag
39 H(3,3)=((s1*lamda1*lamda2)/s2)^2;  %lag of other variable
40 H(4,4)=(lamda1/(2^lamda3))^2;    %own second lag
41 H(5,5)=((s1*lamda1*lamda2)/(s2*(2^lamda3)))^2;  %lag of other
variable
42 %for equation 2 of the VAR
43 H(6,6)=(s2*lamda4)^2;  %constant
44 H(7,7)=((s2*lamda1*lamda2)/s1)^2;  %lag of other variable
45 H(8,8)=(lamda1)^2;     %own lag
46 H(9,9)=((s2*lamda1*lamda2)/(s1*(2^lamda3)))^2;  %lag of other
variable
47 H(10,10)=(lamda1/(2^lamda3))^2;    %own second lag
48 %prior scale matrix for sigma the VAR covariance
```

$\bar{S}$

```
49 S=eye(N);
50 %prior degrees of freedom
```

α

```
51 alpha=N+1;
52 %starting values for the Gibbs sampling algorithm
53 Sigma=eye(N);
54 betaols=vec(inv(X'*X)*(X'*Y));
55 Reps=10000;
56 burn=5000;
57 out1=[]; %will store forecast of GDP growth
58 out2=[]; %will store forecast of inflation
59 i=1;
60 for j=1:Reps
61 %step 1 draw the VAR coefficients
```

$$M^* = (H^{-1} + \Sigma^{-1} \otimes X_t'X_t)^{-1}\left(H^{-1}\tilde{b}_0 + \Sigma^{-1} \otimes X_t'X_t\hat{b}\right)$$
$$(N\times(N\times P+1))\times 1$$

```
62
M=inv(inv(H)+kron(inv(Sigma),X'*X))*(inv(H)*B0+kron(inv(Sigma),X'*X)*bet
aols);
```

$$V^* = (H^{-1} + \Sigma^{-1} \otimes X_t'X_t)^{-1}$$
$$(N\times(N\times P+1))\times(N\times(N\times P+1))$$

```
63 V=inv(inv(H)+kron(inv(Sigma),X'*X));
```

FIGURE 2. Matlab code for example 1 (continued)

Lines 27 to 29 specify the parameters λ_1,λ_2, λ_3,λ_4 that control the tightness of the prior and are used to build the prior covariance. Line 33 specifies $\tilde{b}_0$ the prior mean. As mentioned above in this example we simply assume a prior mean of 1 for the coefficients on own first lags. In practice, this choice should depend on the stationarity properties of the series. Line 35 forms the 10×10 prior variance matrix H. Lines 37 to 47 fill the diagonal elements of this matrix as shown in equation 2.4. Line 49 specifies the prior scale matrix for the inverse Wishart distribution as an identity matrix but specifies the prior degrees of freedom as the minimum possible $N+1$ (line 51) hence making this a non-informative prior. Line 53 sets the starting value for Σ as an identity matrix. We use 10,000 Gibbs replications discarding the first 5000 as burn-in. Line 62 is the first step of the Gibbs sampler with the calculation of the mean of the conditional

```
64 beta=M+(randn(1,N*(N*L+1))*chol(V))';
65 %draw sigma from the IW distribution
66 e=Y-X*reshape(beta,N*L+1,N);
67 %scale matrix
68 scale=e'*e+S;
69 Sigma=IWPQ(T+alpha,inv(scale));
70 if j>burn
71     %forecast GDP growth and inflation for 3 years
72    yhat=zeros(14,2);
73    yhat(1:2,:)=Y(end-1:end,:);
74    for i=3:14
75     yhat(i,:)=[1 yhat(i-1,:) yhat(i-
2,:)]*reshape(beta,N*L+1,N)+randn(1,N)*chol(Sigma);
76 end
77 out1=[out1 [Y(:,1);yhat(3:end,1)]];
78 out2=[out2 [Y(:,2);yhat(3:end,2)]];
79 end
80 end
81 TT=1948.75:0.25:2014;
82 subplot(1,2,1)
83 plot(TT,prctile(out1,[50 10 20 30 70 80 90],2))
84 xlim([1995 2015])
85 title('GDP Growth');
86 subplot(1,2,2)
87 plot(TT,prctile(out2,[50 10 20 30 70 80 90],2))
88 xlim([1995 2015])
89 legend('Median Forecast','10th percentile','20th percentile','30th
percentile','70th percentile','80th percentile','90th percentile');
90 title('Inflation');
```

Published with MATLAB® 7.9

FIGURE 3. Matlab code: Example 1 continued

posterior distribution of the VAR coefficients $\underset{(N\times(N\times P+1))\times 1}{M^*} = \left(H^{-1} + \Sigma^{-1} \otimes X_t'X_t\right)^{-1}\left(H^{-1}\tilde{b}_0 + \Sigma^{-1} \otimes X_t'X_t\hat{b}\right)$ while line 63 compute the variance of this distribution as $\underset{(N\times(N\times P+1))\times(N\times(N\times P+1))}{V^*} = \left(H^{-1} + \Sigma^{-1} \otimes X_t'X_t\right)^{-1}$. On line 64 we draw the VAR coefficients from the normal distribution using M^* and V^*. Line 66 calculates the residuals of the VAR. Line 68 calculates the posterior scale matrix $\bar{\Sigma}$. Line 69 draws the covariance matrix from the inverse Wishart distribution where the function IWPQ uses the method in algorithm 3. Once past the burn-in period we build up the predictive density and save the forecast for each variable. The quantiles of the predictive density are shown in figure 4

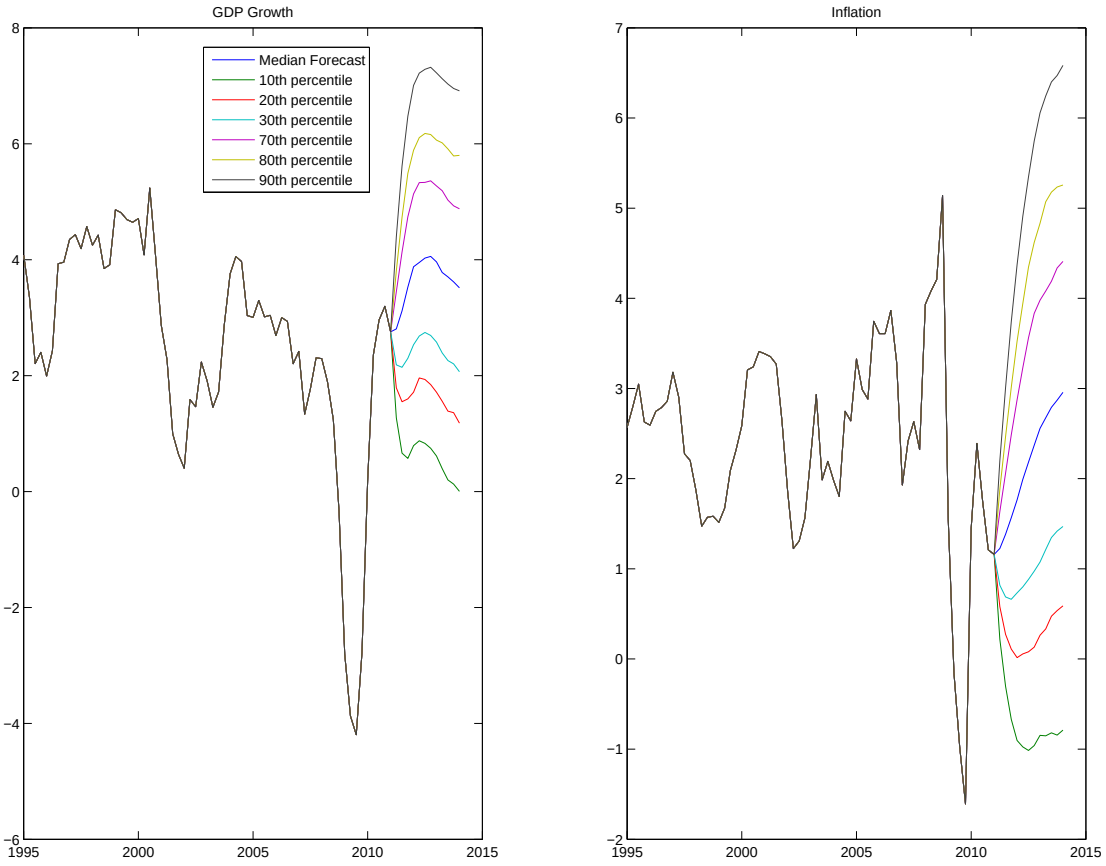

FIGURE 4. Forecast for annual GDP growth and inflation using a VAR with a Minnesota prior

3. The Normal inverse Wishart Prior

3.1. The natural conjugate prior. The normal inverse Wishart prior assumes a normal prior for the VAR coefficients and a inverse Wishart prior for the covariance matrix. This is a conjugate prior for the VAR model. This prior for the VAR parameters can be specified as follows

$$p\left(b\backslash\Sigma\right)\tilde{}N\left(\tilde{b}_0,\Sigma\otimes\bar{H}\right) \tag{3.1}$$

$$p(\Sigma)\tilde{}IW(\bar{S},\alpha) \tag{3.2}$$

where $\tilde{b}_0$ is specified exactly as in equation 2.1. The matrix $\bar{H}$ is a diagonal matrix where the diagonal elements are defined as

$$\left(\frac{\lambda_0\lambda_1}{l^{\lambda_3}\sigma_i}\right)^2 \text{ for the coefficients on lags} \tag{3.3}$$

$$\left(\lambda_0\lambda_4\right)^2 \text{ for the constant} \tag{3.4}$$

So, for our example VAR(2), this matrix is given as

$$\bar{H}=\begin{pmatrix}\left(\lambda_0\lambda_4\right)^2 & 0 & 0 & 0 & 0\\ 0 & \left(\frac{\lambda_0\lambda_1}{\sigma_1}\right)^2 & 0 & 0 & 0\\ 0 & 0 & \left(\frac{\lambda_0\lambda_1}{\sigma_2}\right)^2 & 0 & 0\\ 0 & 0 & 0 & \left(\frac{\lambda_0\lambda_1}{2^{\lambda_3}\sigma_1}\right)^2 & 0\\ 0 & 0 & 0 & 0 & \left(\frac{\lambda_0\lambda_1}{2^{\lambda_3}\sigma_2}\right)^2\end{pmatrix} \tag{3.5}$$

The matrix $\bar{S}$ is defined as a $N\times N$ diagonal matrix with diagonal elements given by

$$\left(\frac{\sigma_i}{\lambda_0}\right)^2 \tag{3.6}$$

For our example VAR this matrix is given by

$$\bar{S} = \begin{pmatrix} \left(\frac{\sigma_1}{\lambda_0}\right)^2 & 0 \\ 0 & \left(\frac{\sigma_2}{\lambda_0}\right)^2 \end{pmatrix} \tag{3.7}$$

The parameters that make up the diagonal elements of $\bar{H}$ and $\bar{S}$ have the following interpretation:

- λ_0 controls the overall tightness of the prior on the covariance matrix.
- λ_1 controls the tightness of the prior on the coefficients on the first lag. As $\lambda_1 \to 0$ the prior is imposed more tightly.
- λ_3 controls the degree to which coefficients on lags higher than 1 are likely to be zero. As λ_3 increases coefficients on higher lags are shrunk to zero more tightly.
- The prior variance on the constant is controlled by λ_4. As $\lambda_4 \to 0$ the constant is shrunk to zero.

To consider the interpretation of this prior (i.e. equations 3.1 and 3.2), consider calculating the prior covariance matrix for the coefficients. This will involve the following operation

$$\bar{S} \otimes \bar{H} \tag{3.8}$$

That is the matrix H or the prior variance of all the VAR coefficients is obtained by a kronecker product in 3.8. Consider calculating this kronecker product in our bi-variate VAR example

$$\begin{pmatrix} \left(\frac{\sigma_1}{\lambda_0}\right)^2 & 0 \\ 0 & \left(\frac{\sigma_2}{\lambda_0}\right)^2 \end{pmatrix} \otimes \begin{pmatrix} (\lambda_0\lambda_4)^2 & 0 & 0 & 0 & 0 \\ 0 & \left(\frac{\lambda_0\lambda_1}{\sigma_1}\right)^2 & 0 & 0 & 0 \\ 0 & 0 & \left(\frac{\lambda_0\lambda_1}{\sigma_2}\right)^2 & 0 & 0 \\ 0 & 0 & 0 & \left(\frac{\lambda_0\lambda_1}{2^{\lambda_3}\sigma_1}\right)^2 & 0 \\ 0 & 0 & 0 & 0 & \left(\frac{\lambda_0\lambda_1}{2^{\lambda_3}\sigma_2}\right)^2 \end{pmatrix}$$

This kronecker product involves each element of $\bar{S}$ being multiplied by the entire $\bar{H}$. If one does one obtains equation 3.9

$$H = \begin{pmatrix} (\sigma_1\lambda_4)^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & (\lambda_1)^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \left(\frac{\sigma_1\lambda_1}{\sigma_2}\right)^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \left(\frac{\lambda_1}{2^{\lambda_3}}\right)^2 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \left(\frac{\sigma_1\lambda_1}{\sigma_2 2^{\lambda_3}}\right)^2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & (\sigma_2\lambda_4)^2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \left(\frac{\sigma_2\lambda_1}{\sigma_1}\right)^2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & (\lambda_1)^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \left(\frac{\sigma_2\lambda_1}{\sigma_1 2^{\lambda_3}}\right)^2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \left(\frac{\lambda_1}{2^{\lambda_3}}\right)^2 \end{pmatrix} \tag{3.9}$$

Note that this is just the Minnesota prior variance with the parameter $\lambda_2 = 1$. Therefore the structure of the natural conjugate prior implies that we treat lags of dependent variable and lags of other variables in each equation of the VAR in exactly the same manner. This is in contrast to the Minnesota prior where the parameter λ_2 governs the tightness of the prior on lags of variables other than the dependent variable.

Given the natural conjugate prior, analytical results exist for the posterior distribution for the coefficients and the covariance matrix. Therefore one clear advantage of this set up over the Minnesota prior is that it allows the derivation of these analytical results without the need for a fixed and diagonal error covariance matrix. The exact formulas for the posteriors are listed in table 1 in Kadiyala and Karlsson (1997).

The Gibbs sampling algorithm for this model is identical to that described in section 2.1. As explained above the only difference is that the variance of the prior distribution is set equalt to H as described in equation 3.9.

3.2. The independent Normal inverse Wishart prior. The restrictions inherent in the natural conjugate prior may be restrictive in many practical circumstances. That is, in many practical applications one may want to treat the coefficients of the lagged dependent variables differently from those of other variables. An example is a situation where the researcher wants impose that some coefficients in a VAR equation are close to zero (e.g. to impose money neutrality or small open economy type restrictions). This can be acheived via the independent Normal

inverse Wishart prior. As the name suggests, this prior involves setting the prior for the VAR coefficients and the error covariance independently (unlike the natural conjugate prior)

$$p(b) \tilde{} N\left(\tilde{b}_0, H\right) \tag{3.10}$$

$$p(\Sigma) \tilde{} IW(\bar{S}, \alpha) \tag{3.11}$$

where the elements of $\tilde{b}_0$, H and $\bar{S}$ are set by the researcher to suit the empirical question at hand. Under this prior analytical expressions for the marginal posterior distributions are not available. Therefore, the Gibbs sampling algorithm outlined in section 1.1 has to be used.

As an example, consider estimating the following VAR(2) model for the US,

$$\begin{pmatrix} R_t \\ GB_t \\ U_t \\ P_t \end{pmatrix} = \begin{pmatrix} c_1 \\ c_2 \\ c_3 \\ c_4 \end{pmatrix} + \begin{pmatrix} b_{11} & b_{12} & b_{13} & b_{14} \\ b_{21} & b_{22} & b_{23} & b_{24} \\ b_{31} & b_{32} & b_{33} & b_{34} \\ b_{41} & b_{42} & b_{43} & b_{44} \end{pmatrix} \begin{pmatrix} R_{t-1} \\ GB_{t-1} \\ U_{t-1} \\ P_{t-1} \end{pmatrix} + \begin{pmatrix} d_{11} & d_{12} & d_{13} & d_{14} \\ d_{21} & d_{22} & d_{23} & d_{24} \\ d_{31} & d_{32} & d_{33} & d_{34} \\ d_{41} & d_{42} & d_{43} & d_{44} \end{pmatrix} \begin{pmatrix} R_{t-2} \\ GB_{t-2} \\ U_{t-2} \\ P_{t-2} \end{pmatrix} + \begin{pmatrix} v_{1t} \\ v_{2t} \\ v_{3t} \\ v_{4t} \end{pmatrix} \tag{3.12}$$

where

$$var \begin{pmatrix} v_{1t} \\ v_{2t} \\ v_{3t} \\ v_{4t} \end{pmatrix} = \Sigma$$

and R_t is the federal funds rate, GB_t is the 10 year government bond yield, U_t is the unemployment rate and P_t is annual CPI inflation. Suppose that one is interested in estimating the response of these variables to a decrease in the government bond yield. This shock may proxy the impact of quantitative easing polices recently adopted. Note, that given the recession in 2010/2011 it is reasonable to assume that the federal funds rate is unlikely to respond to changes in other variables. The standard way to impose this restriction on the contemporaneous period is to identify the yield shock using a Cholesky decomposition of Σ

$$\Sigma = A_0 A_0'$$

where A_0 is a lower triangular matrix. Note, however, that one may also want impose the restriction that the Federal Funds rate does not respond with a lag to changes in the other variables. Given that the Federal Funds rate is near the zero lower bound during the crisis period this restriction can be justified.

The independent Normal Wishart prior offers a convenient way to incorporate these restrictions into the VAR model. One can specify the prior mean for all coefficients equal to zero i.e. $\tilde{b}_0 = 0_{(N\times(N\times P+1))\times 1}$ and the covariance of this prior H as a diagonal marix with diagonal elements equal to a very large number *except for the elements corresponding to the coefficients* b_{12}, b_{13}, b_{14} and d_{12}, d_{13}, d_{14}. The elements of H corresponding to these coefficients are instead set to a very small number and the prior mean of zero is imposed very tightly for them. Therefore the posterior estimates of b_{12}, b_{13}, b_{14} and d_{12}, d_{13}, d_{14} will be very close to zero. We now turn to a matlab implementation of this example using Gibbs sampling.

3.2.1. *Gibbs sampling and the independent normal Wishart prior.* We estimate the VAR model in equation 3.12 using data for the US over the period 2007m1 to 2010m12, the period associated with the financial crisis. We employ a prior which sets the coefficients b_{12}, b_{13}, b_{14} and d_{12}, d_{13}, d_{14} close to zero– i.e. the prior mean for these equals zero and the prior variance is a very small number. Given the very short sample period, we also set a prior for the remaining VAR coefficients. For these remaining coefficients, we assume that the prior mean for coefficients on own first lags are equal to 0.95 and all others equal zero. The prior variance for these is set according to equation 3.9. We set a prior independently for error covariance. We use a Gibbs sampling algorithm to approximate the posterior. The matlab code (example2.m) can be seen in figures 5, 6 and 7.

Lines 16 to 34 of the code calculate σ_1, σ_2, σ_3, σ_4 the variances used to scale the prior variance for the VAR coefficients other than b_{12}, b_{13}, b_{14} and d_{12}, d_{13}, d_{14}. Lines 36 to 38 specify the parameters that will control the variance of the prior on these parameters. Lines 40 to 44 set the prior mean for the VAR coefficients. Under the prior the VAR has the following form:

$$\begin{pmatrix} R_t \\ GB_t \\ U_t \\ P_t \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} + \begin{pmatrix} 0.95 & 0 & 0 & 0 \\ 0 & 0.95 & 0 & 0 \\ 0 & 0 & 0.95 & 0 \\ 0 & 0 & 0 & 0.95 \end{pmatrix} \begin{pmatrix} R_{t-1} \\ GB_{t-1} \\ U_{t-1} \\ P_{t-1} \end{pmatrix} + \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} R_{t-2} \\ GB_{t-2} \\ U_{t-2} \\ P_{t-2} \end{pmatrix} + \begin{pmatrix} v_{1t} \\ v_{2t} \\ v_{3t} \\ v_{4t} \end{pmatrix}$$

```
1 clear
2 addpath('functions');
3 % a VAR for the US using the
4 %load data
5 data=xlsread('\data\dataUS.xls'); %data for US GDP growth and
inflation 1948q1 2010q4
6 N=size(data,2);
7 L=2;   %number of lags in the VAR
8 Y=data;
9 X=[ones(size(Y,1),1) lag0(data,1) lag0(data,2) ];
10 Y=Y(3:end,:);
11 X=X(3:end,:);
12 T=rows(X);
13 %compute standard deviation of each series residual via an ols
regression
14 %to be used in setting the prior
15 %first variable
16 y=Y(:,1);
17 x=X(:,1:2);
18 b0=inv(x'*x)*(x'*y);
19 s1=sqrt(((y-x*b0)'*(y-x*b0))/(rows(y)-2));  %std of residual standard
error
20 %second variable
21 y=Y(:,2);
22 x=X(:,[1 3]);
23 b0=inv(x'*x)*(x'*y);
24 s2=sqrt(((y-x*b0)'*(y-x*b0))/(rows(y)-2));
25 %third variable
26 y=Y(:,3);
27 x=X(:,[1 4]);
28 b0=inv(x'*x)*(x'*y);
29 s3=sqrt(((y-x*b0)'*(y-x*b0))/(rows(y)-2));
30 %fourth variable
31 y=Y(:,4);
32 x=X(:,[1 5]);
33 b0=inv(x'*x)*(x'*y);
34 s4=sqrt(((y-x*b0)'*(y-x*b0))/(rows(y)-2));
35 %parameters to control the prior
36 lamda1=0.1;  %tightness prior on the AR coefficients
37 lamda3=0.05;   %tightness of prior on higher lags
38 lamda4=1;  %tightness of prior on the constant term
39 %specify the prior mean of the coefficients of the Two equations of
the VAR
40 B0=zeros((N*L+1),N);
```

$$\begin{pmatrix} R_t \\ GB_t \\ U_t \\ P_t \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} + \begin{pmatrix} 0.95 & 0 & 0 & 0 \\ 0 & 0.95 & 0 & 0 \\ 0 & 0 & 0.95 & 0 \\ 0 & 0 & 0 & 0.95 \end{pmatrix} \begin{pmatrix} R_{t-1} \\ GB_{t-1} \\ U_{t-1} \\ P_{t-1} \end{pmatrix} + \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} R_{t-2} \\ GB_{t-2} \\ U_{t-2} \\ P_{t-2} \end{pmatrix} + \begin{pmatrix} v_{1t} \\ v_{2t} \\ v_{3t} \\ v_{4t} \end{pmatrix}$$

```
41 for i=1:N
```

FIGURE 5. Matlab code for example 2

Lines 48 to 53 set the variance around the prior for b_{12}, b_{13}, b_{14} and d_{12}, d_{13}, d_{14}. Note that the variance is set to a very small number implying that we incorporate the belief that these coefficients equal zero very strongly. Lines 56 to 88 set the prior variance for the remaining VAR coefficients according to equation 3.9. This is an ad hoc way of incorporating prior information about these coefficients but is important given the small sample. Lines 90 and 92 set the prior for the error covariance as in example 1. Given these priors the Gibbs algorithm is exactly the same as in the previous example. However, we incorporate one change usually adopted by researchers. On lines 108 to 115 we draw the VAR coefficients from its conditional posterior but ensure that the draw is stable. In other words the function stability re-writes the VAR coefficient matrix in companion form and checks if the eigenvalues of this

```
42      B0(i+1,i)=0.95;
43  end
44  B0=vec(B0);
45  %Specify the prior variance of vec(B)
46  H=eye(N*(N*L+1),N*(N*L+1));
47  %small for coefficients we want close to zero
48  H(3,3)=1e-9;    for b12
49  H(4,4)=1e-9;    for b13
50  H(5,5)=1e-9;    for b14
51  H(7,7)=1e-9;    for d12
52  H(8,8)=1e-9;    for d13
53  H(9,9)=1e-9;    for d14
54  %for others like the normal conjugate prior
55  %ist equation
56  H(1,1)=(s1*lamda4)^2;
57  H(2,2)=(lamda1)^2;
58  H(6,6)=(lamda1/(2^lamda3))^2;
59  %second equation
60  H(10,10)=(s2*lamda4)^2;
61  H(11,11)=((s2*lamda1)/s1)^2;
62  H(12,12)=(lamda1)^2;
63  H(13,13)=((s2*lamda1)/s3)^2;
64  H(14,14)=((s2*lamda1)/s4)^2;
65  H(15,15)=((s2*lamda1)/(s1*(2^lamda3)))^2;
66  H(16,16)=(lamda1/(2^lamda3))^2;
67  H(17,17)=((s2*lamda1)/(s3*(2^lamda3)))^2;
68  H(18,18)=((s2*lamda1)/(s4*(2^lamda3)))^2;
69  %third equation
70  H(19,19)=(s3*lamda4)^2;
71  H(20,20)=((s3*lamda1)/s1)^2;
72  H(21,21)=((s3*lamda1)/s2)^2;
73  H(22,22)=(lamda1)^2;
74  H(23,23)=((s3*lamda1)/s4)^2;
75  H(24,24)=((s3*lamda1)/(s1*(2^lamda3)))^2;
76  H(25,25)=((s3*lamda1)/(s2*(2^lamda3)))^2;
77  H(26,26)=(lamda1/(2^lamda3))^2;
78  H(27,27)=((s3*lamda1)/(s4*(2^lamda3)))^2;
79  %fourth equation
80  H(28,28)=(s4*lamda4)^2;
81  H(29,29)=((s4*lamda1)/s1)^2;
82  H(30,30)=((s4*lamda1)/s2)^2;
83  H(31,31)=((s4*lamda1)/s3)^2;
84  H(32,32)=(lamda1)^2;
85  H(33,33)=((s4*lamda1)/(s1*(2^lamda3)))^2;
86  H(34,34)=((s4*lamda1)/(s2*(2^lamda3)))^2;
87  H(35,35)=((s4*lamda1)/(s3*(2^lamda3)))^2;
88  H(36,36)=(lamda1/(2^lamda3))^2;
89  %prior scale matrix for sigma the VAR covariance
90  S=eye(N);
91  %prior degrees of freedom
92  alpha=N+1;
93  %starting values for the Gibbs sampling algorithm
94  Sigma=eye(N);
95  betaols=vec(inv(X'*X)*(X'*Y));
96  Reps=40000;
97  burn=30000;
```

FIGURE 6. example 2: Matlab code continued

matrix are less than or equal to1–i.e. that the VAR is stable (see Hamilton (1994) page 259). Once past the burn-in stage line 123 calculates the structural impact matrix A_0 as the Cholesky decomposition of the draw of Σ and lines 124 to 129 calculate the impulse response to a negative shock in the Government bond yield using this A_0. We save the impulse response functions for each remaining draw of the Gibbs sampler. Quantiles of the saved draws of the impulse response are error bands for the impulse responses.

The resulting median impulse responses and the 68% error bands are shown in figure 8. Note that 68% error bands are typically shown as the 90% or 95% bands can be misleading if the distribution of the impulse response function is skewed due to non-linearity. The response of the Federal Funds rate to this shock is close to zero as

```
98 out1=[]; %will store IRF of R
99 out2=[]; %will store IRF of GB
100 out3=[]; %will store IRF of U
101 out4=[]; %will store IRF of P
102 i=1;
103 for j=1:Reps
104 %step 1 draw the VAR coefficients
105
M=inv(inv(H)+kron(inv(Sigma),X'*X))*(inv(H)*B0+kron(inv(Sigma),X'*X)*bet
aols);
106 V=inv(inv(H)+kron(inv(Sigma),X'*X));
107 %check for stability of the VAR
108 check=-1;
109 while check<0
110 beta=M+(randn(1,N*(N*L+1))*chol(V))';
111 CH=stability(beta,N,L);
112 if CH==0
113     check=10;
114 end
115 end
116 %draw sigma from the IW distribution
117 e=Y-X*reshape(beta,N*L+1,N);
118 %scale matrix
119 scale=e'*e+S;
120 Sigma=IWPQ(T+alpha,inv(scale));
121 if j>burn
122     %impulse response using a cholesky decomposition
123     A0=chol(Sigma);
124     v=zeros(60,N);
125     v(L+1,2)=-1; %shock the government bondyield
126    yhat=zeros(60,N);
127    for i=3:60
128     yhat(i,:)=[0 yhat(i-1,:) yhat(i-
2,:)]*reshape(beta,N*L+1,N)+v(i,:)*A0;
129 end
130 out1=[out1 yhat(3:end,1)];
131 out2=[out2 yhat(3:end,2)];
132 out3=[out3 yhat(3:end,3)];
133 out4=[out4 yhat(3:end,4)];
134 end
135 end
136 subplot(2,2,1)
137 plot([prctile(out1,[50 16 84],2) zeros(size(out3,1),1)]);
138 title('Response of the Federal Funds rate');
139 axis tight
140 subplot(2,2,2)
141 plot([prctile(out2,[50 16 84],2) zeros(size(out3,1),1)]);
142 title('Response of the Government Bond Yield');
143 axis tight
144 subplot(2,2,3)
145 plot([prctile(out3,[50 16 84],2) zeros(size(out3,1),1)]);
146 title('Response of the Unemployment Rate');
147 axis tight
148 subplot(2,2,4)
149 plot([prctile(out4,[50 16 84],2) zeros(size(out3,1),1)]);
150 title('Response of Inflation');
151 axis tight
152 legend('Median Response','Upper 84%','Lower 16%','Zero Line');
```

Published with MATLAB® 7.9

FIGURE 7. example2: Matlab code (continued)

implied by the Cholesky decomposition and the prior on b_{12}, b_{13}, b_{14} and d_{12}, d_{13}, d_{14}. A 0.3% fall in the Government bond yield lowers unemployment by 0.1% after 10 months (but the impact is quite uncertain as evident from the wide error bands). The impact on inflation is much more imprecise with the zero line within the error bands for most of the impulse horizon.

4. Steady State priors

In some circumstances it is useful to incorporate priors about the long run behaviour of the variables included in the VAR. For example one may be interested in forecasting inflation using a VAR model. It can be argued that

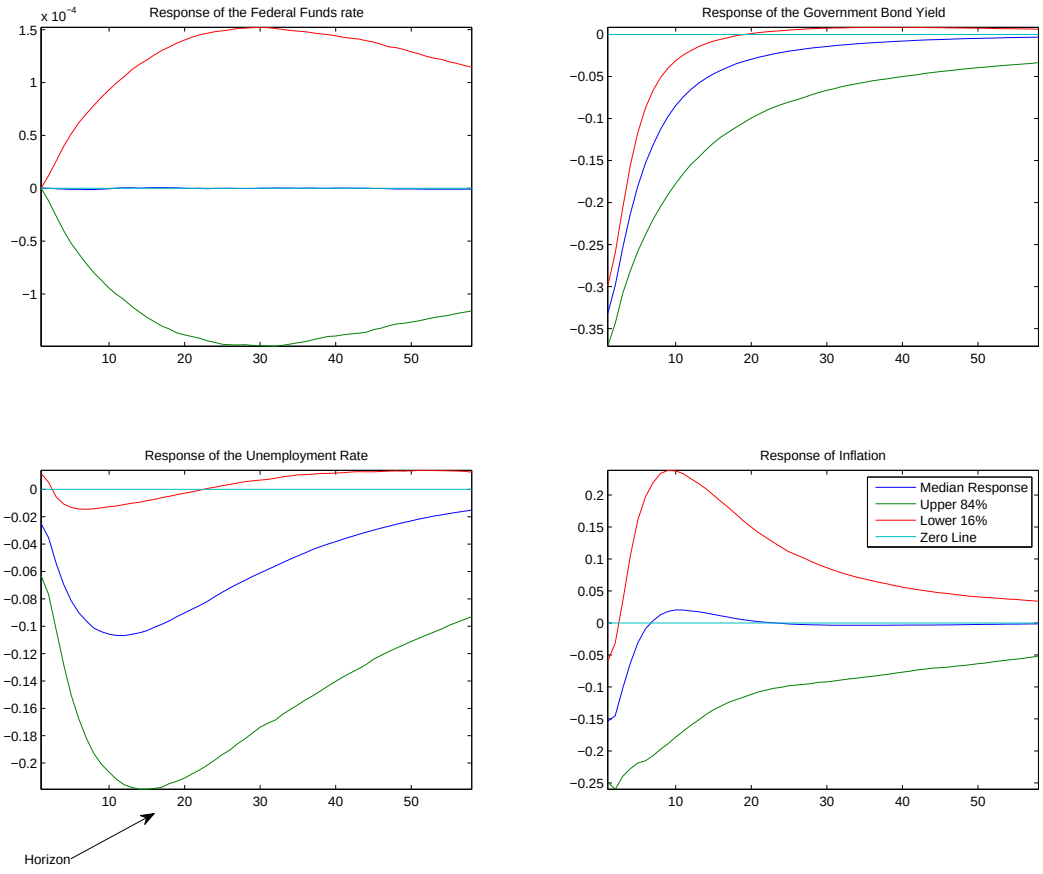

FIGURE 8. Impulse response to a fall in the Government bond yield

inflation in the long run will be close to the target set by the central bank. This information is a potentially useful input as a prior.

Note that while the priors introduced above allow the researcher to have an impact on the value of the constant terms in the VAR, there is no direct way to affect the long run mean (note that forecasts converge to the long run unconditional mean). Consider our example bi-variate VAR re-produced below

$$\begin{pmatrix} y_t \\ x_t \end{pmatrix} = \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} + \begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{pmatrix} \begin{pmatrix} y_{t-1} \\ x_{t-1} \end{pmatrix} + \begin{pmatrix} d_{11} & d_{12} \\ d_{21} & d_{22} \end{pmatrix} \begin{pmatrix} y_{t-2} \\ x_{t-2} \end{pmatrix} + \begin{pmatrix} v_1 \\ v_2 \end{pmatrix}, VAR \begin{pmatrix} v_1 \\ v_2 \end{pmatrix} = \Sigma \tag{4.1}$$

The Minnesota and the Normal inverse Wishart priors place a prior on the constants $\begin{pmatrix} c_1 \\ c_2 \end{pmatrix}$. The long run or steady state means for y_t and x_t denoted by μ_1 and μ_2 however, is defined as (see Hamilton (1994) page 258)

$$\begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix} = \left(\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} - \begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{pmatrix} - \begin{pmatrix} d_{11} & d_{12} \\ d_{21} & d_{22} \end{pmatrix} \right)^{-1} \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} \tag{4.2}$$

Villani (2009) proposes a prior distribution for the unconditional means $\mu = \{\mu_1, \mu_2\}$ along with coefficients of the VAR model. This requires one to re-write the model in terms of $\mu = \{\mu_1, \mu_2\}$ rather than the constants c_1 and c_2. This can be done in our example VAR by substituting for c_1 and c_2 in equation 4.1 using the values of these constants from equation 4.2 to obtain

$$\begin{aligned} \begin{pmatrix} y_t \\ x_t \end{pmatrix} &= \left(\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} - \begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{pmatrix} - \begin{pmatrix} d_{11} & d_{12} \\ d_{21} & d_{22} \end{pmatrix} \right) \begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix} \\ &\quad + \begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{pmatrix} \begin{pmatrix} y_{t-1} \\ x_{t-1} \end{pmatrix} + \begin{pmatrix} d_{11} & d_{12} \\ d_{21} & d_{22} \end{pmatrix} \begin{pmatrix} y_{t-2} \\ x_{t-2} \end{pmatrix} + \begin{pmatrix} v_1 \\ v_2 \end{pmatrix} \end{aligned}$$

or more compactly in terms of lag operators as

$$B(L)(Z_t - \mu) = v \tag{4.3}$$

where $Z_t = \{y_t, x_t\}$, $v = \{v_1, v_2\}$and $B(L) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} - \begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{pmatrix} L - \begin{pmatrix} d_{11} & d_{12} \\ d_{21} & d_{22} \end{pmatrix} L^2$

Villani (2009) proposes a normal prior for μ

$$p(\mu) \sim N(\mu_0, \Sigma_\mu) \tag{4.4}$$

The priors for the autoregressive coefficients and the error covariance are specified independently. For example, one can specify the Minnesota prior for the autoregressive coefficients and an inverse Wishart prior for the error covariance.

Note that there are three sets of parameters to be estimated in this VAR model: (1) The VAR coefficients, the error covariance and the long run means μ.Villani (2009) describes a Gibbs sampling algorithm to estimate the model and we turn to this next.

4.1. Gibbs sampling algorithm. The Gibbs sampling algorithm for this model is an extension of the algorithm described in section 1.1. Conditional on knowing μ the reparametrised model is a just a standard VAR and standard methods apply. The algorithm works in the following steps

Step 1 Set a normal prior for the VAR coefficients $p(\bar{b}) \tilde{} N\left(\tilde{b}_0, H\right)$ where $\bar{b}$ the (vectorised) VAR coefficients except for the constant terms. The prior for the covariance matrix of the residuals Σ is inverse Wishart and given by $IW\left(\bar{S}, \alpha\right)$. The prior for the long run means is $p(\mu) \sim N(\mu_0, \Sigma_\mu)$. Set a starting value for μ. A starting value can be set via OLS estimates of the VAR coefficients as

$$\mu_{ols} = \left(I - \tilde{B}\right)^{-1} \hat{C}$$

where $\tilde{B}$ are the OLS estimates of the VAR coefficients in companion form and C denotes the OLS estimates of the constant terms in a comformable matrix. For the bi-variate VAR in equation 4.1 this looks as follows

$$\begin{pmatrix} \mu_{OLS,1} \\ \mu_{OLS,2} \end{pmatrix} = \left(\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} - \begin{pmatrix} \hat{b}_{11} & \hat{b}_{12} & \hat{d}_{11} & \hat{d}_{12} \\ \hat{b}_{21} & \hat{b}_{22} & \hat{d}_{21} & \hat{d}_{22} \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix} \right)^{-1} \begin{pmatrix} \hat{c}_1 \\ \hat{c}_2 \\ 0 \\ 0 \end{pmatrix}$$

Step 2 Sample the VAR coefficients from their conditonal distribution. Conditional on μ, equation 4.3 implies that the model is a VAR in the transformed (or de-meaned) variables $Y_t^0 = Z_t - \mu$. The conditional posterior distribution of the VAR coefficients is normal distribution $H\left(\bar{b} \backslash \Sigma, \mu, Z_t^*\right) \tilde{} N\left(M^*, V^*\right)$ where

$$\underset{(N\times(N\times P))\times 1}{M^*} = \left(H^{-1} + \Sigma^{-1} \otimes X_t^{0\prime} X_t^0\right)^{-1} \left(H^{-1}\tilde{b}_0 + \Sigma^{-1} \otimes X_t^{0\prime} X_t^0 \hat{b}\right) \tag{4.5}$$

$$\underset{(N\times(N\times P))\times(N\times(N\times P))}{V^*} = \left(H^{-1} + \Sigma^{-1} \otimes X_t^{0\prime} X_t^0\right)^{-1} \tag{4.6}$$

where $X_t^0 = [Y_{t-1}^0, ..., Y_{t-p}^0]$ and $\hat{b} = vec\left(\left(X_t^{0\prime} X_t^0\right)^{-1} \left(X_t^{0\prime} Y_t^0\right)\right)$. Note that the dimensions of M^* and V^* are different relative to those shown in section 1.1 because X_t^0 does not contain a constant term. Once M^* and V^* are calculated, the VAR coefficients are drawn from the normal distribution as before.

Step 3 Draw Σ from its conditional distribution $H\left(\Sigma \backslash \bar{b}, \mu, Z_t^*\right) \tilde{} IW\left(\bar{\Sigma}, T + \alpha\right)$ where $\bar{\Sigma} = \bar{S} + \left(Y_t^0 - X_t^0 B^1\right)' \left(Y_t^0 - X_t^0 B^1\right)$ where B^1 is the previous draw of the VAR coefficients reshaped into a matrix with dimensions $(N \times P) \times N$ so it is conformable with X_t^*.

Step 4 Draw μ from its conditional distribution. Villani (2009) shows that the conditional distribution of μ is given as $H\left(\mu \backslash \bar{b}, \Sigma, Z_t^*\right) \sim N\left(\mu^*, \Omega^*\right)$ where

$$\Omega^* = \left(\Sigma_\mu^{-1} + U'\left(D'D \otimes \Sigma^{-1}\right) U'\right)^{-1} \tag{4.7}$$

$$\mu^* = \Omega^* \left(U' vec\left(\Sigma^{-1} Y' D\right) + \Sigma_\mu^{-1} \mu_0\right) \tag{4.8}$$

where D is a $T \times (P+1)$ matrix $D = [c_t, -c_{t-1}, ... - c_{t-p}]$ where c_t is the constant term (a $T \times 1$) vector equal to one). U is a matrix with the following structure

$$U = \begin{pmatrix} I_N \\ B_1 \\ . \\ B_P \end{pmatrix}$$

For our two variable VAR U looks as follows

$$U = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ b_{11} & b_{12} \\ b_{21} & b_{22} \\ d_{11} & d_{12} \\ d_{21} & d_{22} \end{pmatrix} \tag{4.9}$$

```
clear
addpath('functions');
% a bi-variate VAR with a Minnesota Prior and Gibbs Sampling
%load data
data=xlsread('\data\datain.xls'); %data for US GDP growth and inflation
1948q1 2010q4
N=size(data,2);
L=2;   %number of lags in the VAR
Y=data;
X=[lag0(data,1) lag0(data,2) ones(rows(data),1) ];
Y=Y(3:end,:);
X=X(3:end,:);
T=rows(X);
%compute standard deviation of each series residual via an ols
regression
%to be used in setting the prior
%first variable
y=Y(:,1);
x=[ones(T,1) X(:,1)];
b0=inv(x'*x)*(x'*y);
s1=sqrt(((y-x*b0)'*(y-x*b0))/(rows(y)-2));  %std of residual standard
error
%second variable
y=Y(:,2);
x=[ones(T,1) X(:,2)];
b0=inv(x'*x)*(x'*y);
s2=sqrt(((y-x*b0)'*(y-x*b0))/(rows(y)-2));
%specify parameters of the minnesota prior
lamda1=1; %controls the prior on own lags
lamda2=1;
lamda3=1;
lamda4=1;
%specify the prior mean of the coefficients of the Two equations of
the VAR
B01=[1;0;0;0];
B02=[0;1;0;0];
B0=[B01;B02];
%Specify the prior variance of vec(B)
H=zeros(8,8);
%for equation 1  of the VAR
H(1,1)=(lamda1)^2;     %own lag
H(2,2)=((s1*lamda1*lamda2)/s2)^2;  %lag of other variable
H(3,3)=(lamda1/(2^lamda3))^2;    %own second lag
H(4,4)=((s1*lamda1*lamda2)/(s2*(2^lamda3)))^2;  %lag of other variable
%for equation 2 of the VAR
H(5,5)=((s2*lamda1*lamda2)/s1)^2;  %lag of other variable
H(6,6)=(lamda1)^2;     %own lag
H(7,7)=((s2*lamda1*lamda2)/(s1*(2^lamda3)))^2;  %lag of other variable
H(8,8)=(lamda1/(2^lamda3))^2;    %own second lag
%prior scale matrix for sigma the VAR covariance
S=eye(N);
%prior degrees of freedom
alpha=N+1;
%set priors for the long run mean which is a N by 1 vector
p(μ) ~ N(μ0, Σμ)
M0=[1 1]; %prior mean
V0=eye(N)*0.001;  %prior variance
%starting values via OLS
betaols=inv(X'*X)*(X'*Y);
F=[betaols(1:N*L,:)';eye(N*(L-1),N*L)]; %companion form
```

FIGURE 9. Matlab code for VAR with steady state priors

Finally $Y = Z_t - B_1 Z_{t-1} - ...B_p Z_{t-p}$ where B_i denotes the VAR coefficients on the ith lag from the previous Gibbs iteration.

Step 5 Repeat steps 2 to 4 M times to obtain $B^1...B^M$ and $(\Sigma)^1 .... (\Sigma)^M$ and $\mu^1...\mu^M$ The last H values of B, μ and Σ from these iterations is used to form the empirical distribution of these parameters.

4.2. Gibbs sampling algorithm for the VAR with steady state priors. The matlab code. We estimate the VAR with steady state priors using the same data used in the first example (quarterly data on annual GDP growth and CPI inflation for the US from 1948Q2 to 2010Q4) and consider a long term forecast of these variables. The code

```
56 C=zeros(rows(F),1);
57 C(1:N)=betaols(N*L+1,:)';
```

$$\begin{pmatrix} \mu_{OLS,1} \\ \mu_{OLS,2} \end{pmatrix} = \left(\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} - \begin{pmatrix} \hat{b}_{11} & \hat{b}_{12} & \hat{d}_{11} & \hat{d}_{12} \\ \hat{b}_{21} & \hat{b}_{22} & \hat{d}_{21} & \hat{d}_{22} \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix} \right)^{-1} \begin{pmatrix} \hat{c}_1 \\ \hat{c}_2 \\ 0 \\ 0 \end{pmatrix}$$

```
58 MU=inv(eye(rows(F))-F)*C;  %ols estimate of the mean inv(I-B)C
59 e=Y-X*betaols;
60 Sigma=(e'*e)/T;
61 Reps=10000;
62 burn=5000;
63 out1=[]; %will store forecast of GDP growth
64 out2=[]; %will store forecast of inflation
65 i=1;
66 for j=1:Reps
67 %demean the data
```

$$Y_t^0 = Z_t - \mu$$

```
68   Y0=data-repmat(MU(1:N)',rows(data),1);
69   X0=[];
70   for jj=1:L
71   X0=[X0 lag0(Y0,jj) ];
72   end
73   Y0=Y0(L+1:end,:);
74   X0=X0(L+1:end,:);
75 %step 1 draw the VAR coefficients
76 bols=vec(inv(X0'*X0)*(X0'*Y0));
```

$$\underset{(N\times(N\times P))\times 1}{M^*} = (H^{-1} + \Sigma^{-1} \otimes X_t^{0\prime} X_t^0)^{-1} \left(H^{-1} \tilde{b}_0 + \Sigma^{-1} \otimes X_t^{0\prime} X_t^0 \hat{b} \right)$$

```
77
M=inv(inv(H)+kron(inv(Sigma),X0'*X0))*(inv(H)*B0+kron(inv(Sigma),X0'*X0)*
bols);
```

$$\underset{(N\times(N\times P))\times(N\times(N\times P))}{V^*} = (H^{-1} + \Sigma^{-1} \otimes X_t^{0\prime} X_t^0)^{-1}$$

```
78 V=inv(inv(H)+kron(inv(Sigma),X0'*X0));
79 beta=M+(randn(1,N*(N*L))*chol(V))';
80 beta1=reshape(beta,N*L,N);
81 %draw sigma from the IW distribution
82 e=Y0-X0*beta1;
83 %scale matrix
```

$$\bar{\Sigma} = \bar{S} + (Y_t^0 - X_t^0 B^1)'(Y_t^0 - X_t^0 B^1)$$

```
84 scale=e'*e+S;
85 Sigma=IWPQ(T+alpha,inv(scale));
86 %step 3 draw MU the long run mean conditional on beta and sigma (see
87   %Appendix A in Villani.
88
```

$$Y = Z_t - B_1 Z_{t-1} - \ldots B_p Z_{t-p}$$

```
89   Y1=Y-X(:,1:end-1)*beta1;
90   U=eye(N);
```

FIGURE 10. Matlab code for steady state VAR (continued)

for the model (example3.m) is presented in figures 9, 10, 11 and 12. The code is identical to the first matlab example until line 50 where we set the prior for the long run means of the two variables. As an example we set the prior mean equal to 1 for both μ_1 and μ_2 and a tight prior variance. Lines 54 to 58 estimate the VAR coefficients via OLS and estimate a starting value for μ_1 and μ_2 as described in Step 1 of the Gibbs sampling algorithm above. Line 68 is the first step of the Gibbs algorithm and computes the demeaned data $Y_t^0 = Z_t - \mu$ and uses this on line 77 and 78 to compute the mean and the variance of the conditional posterior distribution $H\left(\bar{b} \backslash \Sigma, \mu, Z_t^*\right)$ and samples the VAR coefficients from the normal distribution. Lines 81 to 85 draw Σ from the inverse Wishart distribution. Lines 89 to 100 draw μ_1 and μ_2 from the normal distribution. On line 89 the code creates the matrix $Y = Z_t - B_1 Z_{t-1} - ...B_p Z_{t-p}$.

$$U = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ b_{11} & b_{12} \\ b_{21} & b_{22} \\ d_{11} & d_{12} \\ d_{21} & d_{22} \end{pmatrix}$$

```
91   jj=1;
92   for jx=1:L
93       betai=beta1(jj:jj+N-1,:);
94       U=[U;betai'];
95       jj=jj+N;
96   end
```

$D = [c_t, -c_{t-1}, \ldots -c_{t-p}]$

```
97   D=[ones(T,1) -ones(T,L)];
98   vstar1=inv(U'*kron(D'*D,inv(Sigma))*U+inv(V0));
```

$\Omega^* = (\Sigma_\mu^{-1} + U'(D'D \otimes \Sigma^{-1})U')^{-1}$ %posterior variance

```
99   mstar1=vstar1*(U'*vec(inv(Sigma)*Y1'*D)+inv(V0)*M0');
```

$\mu^* = \Omega^*(U'vec(\Sigma^{-1}Y'D) + \Sigma_\mu^{-1}\mu_0)$ %posterior mean

```
100   MU=mstar1+(randn(1,N)*chol(vstar1))';  %draw MU
101
102 if j>burn
103
104     %forecast GDP growth and inflation for 3 years
105    F=[beta1(1:N*L,:)';eye(N*(L-1),N*L)]; %companion form
106    mu=[];
107    for i=1:L
108    mu=[mu;MU];
109    end
```

$$\left(\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} - \begin{pmatrix} b_{11} & b_{12} & d_{11} & d_{12} \\ b_{21} & b_{22} & d_{21} & d_{22} \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}\right)\begin{pmatrix} \mu_1 \\ \mu_2 \\ \mu_1 \\ \mu_2 \end{pmatrix} = \begin{pmatrix} c_1 \\ c_2 \\ 0 \\ 0 \end{pmatrix}$$

```
110    C=(eye(rows(F))-F)*mu;   %implied constant
111
112
113     yhat=zeros(44,2);
114     yhat(1:2,:)=Y(end-1:end,:);
115     for i=3:44
116      yhat(i,:)=C(1:N)'+[yhat(i-1,:) yhat(i-
2,:)]*reshape(beta,N*L,N)+randn(1,N)*chol(Sigma);
117 end
118 out1=[out1 [Y(:,1);yhat(3:end,1)]];
119 out2=[out2 [Y(:,2);yhat(3:end,2)]];
120 end
121 end
122 TT=1948.75:0.25:2021.5;
123 subplot(1,2,1)
124 plot(TT,[mean(out1,2) prctile(out1,[50 10 20 30 70 80 90],2)])
125 xlim([1995 2022])
126 title('GDP Growth');
127 subplot(1,2,2)
```

FIGURE 11. Matlab code for VAR with steady state priors (continued)

Lines 90 to 96 create the matrix $U = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ b_{11} & b_{12} \\ b_{21} & b_{22} \\ d_{11} & d_{12} \\ d_{21} & d_{22} \end{pmatrix}$. Line 97 creates the matrix $D = [c_t, -c_{t-1}, \ldots - c_{t-p}]$. Lines 98 and 99 compute the variance and mean of the conditional posterior distribution of μ (see equation 4.7 and 4.8)while line 100 draws μ from the normal distribution. After the burn-in stage the VAR is used to do a forecast for

```
128 plot(TT,[ mean(out2,2) prctile(out2,[50 10 20 30 70 80 90],2)])
129 xlim([1995 2022])
130 legend('Mean Forecast','Median Forecast','10th percentile','20th
percentile','30th percentile','70th percentile','80th percentile','90th
percentile');
131 title('Inflation');
```

Published with MATLAB® 7.9

FIGURE 12. Matlab code for VAR with Steady State priors.

40 quarters. It is convenient to parameterise the VAR in the usual form i.e as in equation 4.1. On line 110 the code calculates the implied constants in the VAR using the fact that

$$\left(\left(\begin{array}{cccc} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{array}\right) - \left(\begin{array}{cccc} b_{11} & b_{12} & d_{11} & d_{12} \\ b_{21} & b_{22} & d_{21} & d_{22} \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{array}\right)\right)\left(\begin{array}{c} \mu_1 \\ \mu_2 \\ \mu_1 \\ \mu_2 \end{array}\right) = \left(\begin{array}{c} c_1 \\ c_2 \\ 0 \\ 0 \end{array}\right)$$

and lines 113 to 117 calculate the forecast for each retained draw. The resulting forecast distribution in figure 13 is centered around the long run mean close to 1 for both variables.

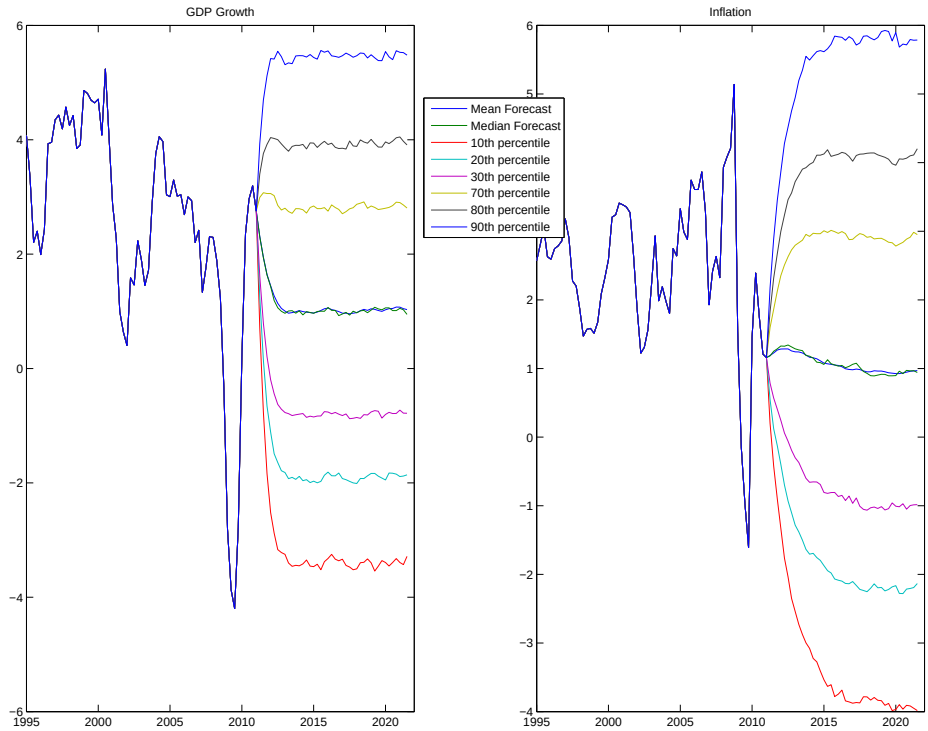

FIGURE 13. Forecast distribution for the VAR with steady state priors.

5. Implementing priors using dummy observations

The computation of the mean of the conditional posterior distribution (see equation 1.5) requires the inversion of $(N \times (N \times (P+1))) \times (N \times (N \times (P+1)))$ matrix $\left(H^{-1} + \Sigma^{-1} \otimes X_t' X_t\right)^{-1}$. For large VARs ($N \geq 20$) this matrix has very large dimensions (e.g for $N = 20$ and $P = 2$ this is a 820×820 matrix). This can slow down the Gibbs sampling algorithm considerably. This computational constraint can be thought of as one potential disadvantage of the way we have incorporated the prior, i.e. via the covariance matrix H which has the dimensions $(N \times (N \times (P+1))) \times (N \times (N \times (P+1)))$.

Note also that our method of implementing the prior makes it difficult to incorporate priors about combination of coefficients in each equation or across equations. For instance, if one is interested in a prior that incorporates the belief that the sum of the coefficients on lags of the dependent variable in each equation sum to 1 (i.e. each variable has a unit root) this is very difficult to implement using a prior covariance matrix. Priors on combinations of coefficients across equations may arise from the implications of DSGE models (see Negro and Schorfheide (2004)). Again these are difficult to implement using the standard approach.

An alternative approach to incorporating prior information into the VAR is via dummy observations or artificial data. Informally speaking this involves generating artificial data from the model assumed under the prior and mixing this with the actual data. The weight placed on the artificial data determines how tightly the prior is imposed.

5.1. The Normal Wishart (Natural Conjugate) prior using dummy observations. Consider artificial data denoted Y_D and X_D (we consider in detail below how to generate this data) such that

$$\begin{aligned} b_0 &= \left(X_D' X_D\right)^{-1} \left(X_D' Y_D\right) \\ S &= \left(Y_D - X_D b_0\right)' \left(Y_D - X_D b_0\right) \end{aligned} \tag{5.1}$$

where $\tilde{b}_0 = vec(b_0)$. In other words a regression of Y_D on X_D gives the prior mean for the VAR coefficients and sum of squared residuals give the prior scale matrix for the error covariance matrix. The prior is of the normal inverse Wishart form

$$\begin{aligned} &p(B \backslash \Sigma) \tilde{} N\left(\tilde{b}_0, \Sigma \otimes \left(X_D' X_D\right)^{-1}\right) \\ &p\left(\Sigma\right) \tilde{} IW(S, T_D - K) \end{aligned} \tag{5.2}$$

where T_D is the length of the artificial data and K denotes the number of regressors in each equation.

Given this artificial data, the conditional posterior distributions for the VAR parameters are given by

$$H\left(b\backslash\Sigma,Y_t\right)\tilde{}N(vec(B^*),\Sigma\otimes\left(X^{*\prime}X^*\right)^{-1}) \tag{5.3}$$
$$H\left(\Sigma\backslash b,Y_t\right)\tilde{}IW(S^*,T^*)$$

where $Y^*=[Y;Y_D],X^*=[X;X_D]$ i.e. the actual VAR left and right hand side variables appended by the artificial data and T^* denotes the number of rows in Y^* and

$$B^*=\left(X^{*\prime}X^*\right)^{-1}\left(X^{*\prime}Y^*\right)$$
$$S^*=\left(Y^*-X^*b\right)'\left(Y^*-X^*b\right)$$

Note that the conditional posterior distribution has a simple form and the variance of $H\left(b\backslash\Sigma,Y_t\right)$ only involves the inversion of $N\times P+1$ matrix making a Gibbs sampler based on this formulation much more computationally efficient in large models.

5.1.1. *Creating the dummy observations for the Normal Wishart prior.* The key question however is, where do Y_D and X_D come from? The artificial observations are formed by the researcher and are created using the following hyper-parameters:

- τ controls the overall tightness of the prior
- d controls the tightness of the prior on higher lags
- c controls the tightness of the prior on constants
- σ_i are standard deviation of error terms from OLS estimates of AR regression for each variable in the model

To discuss the creation of the dummy observations we are going to use the bi-variate VAR given below as an example:

$$\begin{pmatrix} y_t \\ x_t \end{pmatrix}=\begin{pmatrix} c_1 \\ c_2 \end{pmatrix}+\begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{pmatrix}\begin{pmatrix} y_{t-1} \\ x_{t-1} \end{pmatrix}+\begin{pmatrix} d_{11} & d_{12} \\ d_{21} & d_{22} \end{pmatrix}\begin{pmatrix} y_{t-2} \\ x_{t-2} \end{pmatrix}+\begin{pmatrix} v_1 \\ v_2 \end{pmatrix},VAR\begin{pmatrix} v_1 \\ v_2 \end{pmatrix}=\Sigma \tag{5.4}$$

Consider dummy observations that implement the prior on the coefficients on the first lag of y_t and x_t. The artificial data (denoted by $Y_{D,1}$ and $X_{D,1}$) is given by

$$\begin{aligned} Y_{D,1} &= \begin{pmatrix} (1/\tau)\,\sigma_1 & 0 \\ 0 & (1/\tau)\,\sigma_2 \end{pmatrix} \\ X_{D,1} &= \begin{pmatrix} 0 & (1/\tau)\,\sigma_1 & 0 & 0 & 0 \\ 0 & 0 & (1/\tau)\,\sigma_2 & 0 & 0 \end{pmatrix} \end{aligned} \tag{5.5}$$

To see the intuition behind this formulation consider the VAR model using the artificial data

$$\underbrace{\begin{pmatrix} (1/\tau)\,\sigma_1 & 0 \\ 0 & (1/\tau)\,\sigma_2 \end{pmatrix}}_{Y_{D,1}}=\underbrace{\begin{pmatrix} 0 & (1/\tau)\,\sigma_1 & 0 & 0 & 0 \\ 0 & 0 & (1/\tau)\,\sigma_2 & 0 & 0 \end{pmatrix}}_{X_{D,1}}\underbrace{\begin{pmatrix} c_1 & c_2 \\ b_{11} & b_{21} \\ b_{12} & b_{22} \\ d_{11} & d_{21} \\ d_{12} & d_{22} \end{pmatrix}}_{B}+\begin{pmatrix} v_1 \\ v_2 \end{pmatrix} \tag{5.6}$$

Expanding the equation above gives the following

$$\begin{pmatrix} (1/\tau)\,\sigma_1 & 0 \\ 0 & (1/\tau)\,\sigma_2 \end{pmatrix}=\begin{pmatrix} (1/\tau)\,\sigma_1 b_{11} & (1/\tau)\,\sigma_1 b_{21} \\ (1/\tau)\,\sigma_2 b_{12} & (1/\tau)\,\sigma_2 b_{22} \end{pmatrix}+\begin{pmatrix} v_1 \\ v_2 \end{pmatrix} \tag{5.7}$$

Consider the first equation in the expression above $(1/\tau)\,\sigma_1=(1/\tau)\,\sigma_1 b_{11}+v_1$ or $b_{11}=1-\frac{\tau v_1}{\sigma_1}$. Taking the expected value of this gives $E\left(b_{11}\right)=1-E\left(\frac{\tau v_1}{\sigma_1}\right)$ which equals 1 as $E\left(v_1\right)=0$. In other words, the dummy variables imply a prior mean of 1 for b_{11}. Similarly, the variance of b_{11} is $\frac{\tau^2 var(v_1)}{\sigma_1^2}$. Note that the implied prior mean and variance for b_{11} is identical to the Natural conjugate prior discussed above. That is under the prior $b_{11}\tilde{}N\left(1,\frac{\tau^2 var(v_1)}{\sigma_1^2}\right)$. As $\tau\to 0$ the prior is imlemented more tightly.

Consider the second equation implied by expression 5.7: $0=(1/\tau)\,\sigma_1 b_{21}+v_2$ or $b_{21}=-\frac{\tau v_2}{\sigma_1}$. This implies that $E\left(b_{21}\right)=0$ and $var\left(b_{21}\right)=\frac{\tau^2 var(v_2)}{\sigma_1^2}$. Thus $b_{21}\tilde{}N\left(0,\frac{\tau^2 var(v_2)}{\sigma_1^2}\right)$ where the variance is of the same form as the corresponding element in equation 3.9.

Thus, the artificial observations in 5.5 implement the Normal inverse Wishart prior for the coefficients on the first lags of the two variables. We need to create artificial observations to implement the prior on the second lags. These are given by the following matrices

$$\begin{aligned} Y_{D,2} &= \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix} \\ X_{D,2} &= \begin{pmatrix} 0 & 0 & 0 & (1/\tau)\,\sigma_1 2^d & 0 \\ 0 & 0 & 0 & 0 & (1/\tau)\,\sigma_2 2^d \end{pmatrix} \end{aligned} \tag{5.8}$$

Proceeding as in equation 5.7 one can show that these dummy variables imply a prior mean of 0 for the second lag with the prior variance of the same form as in equation 3.9. For example, the prior variance associated with d_{11} is $\frac{\tau^2 var(v_1)}{\sigma_1^2 2^d}$.

The artificial observations that control the prior on the constants in the model are given by:

$$\begin{aligned} Y_{D,3} &= \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix} \\ X_{D,3} &= \begin{pmatrix} 1/c & 0 & 0 & 0 & 0 \\ 1/c & 0 & 0 & 0 & 0 \end{pmatrix} \end{aligned} \tag{5.9}$$

As $c \to 0$ the prior is imlemented more tightly. The dummy observations to implement the prior on the error covariance matrix are given by

$$\begin{aligned} Y_{D,4} &= \begin{pmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{pmatrix} \\ X_{D,4} &= \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix} \end{aligned} \tag{5.10}$$

with the magnitude of the diagonal elements of Σ controlled by the scale of the diagonal elements of $Y_{D,4}$ (i.e. larger diagonal elements implement the prior belief that the variance of v_1 and v_2 is larger).

The prior is implemented by adding all these dummy observations to the actual data. That is

$$Y^* = [Y; Y_{D,1}; Y_{D,2}; Y_{D,3}; Y_{D,4}], X^* = [X; X_{D,1}; X_{D,2}; X_{D,3}; X_{D,4}]$$

With this appended data in hand, the conditional distributions in equation 5.3 can be used to implement the Gibbs sampling algorithm. Note that, as discussed in Banbura *et al.* (2007), these dummy observations for a general N variable VAR with P lags are given as

$$Y_D = \begin{pmatrix} \frac{diag(\chi_1\sigma_1 ... \chi_N\sigma_N)}{\tau} \\ 0_{N\times(P-1)\times N} \\ ..... \\ diag\left(\sigma_1 ... \sigma_N\right) \\ ..... \\ 0_{1\times N} \end{pmatrix}, X_D = \begin{pmatrix} \frac{J_P \otimes diag(\sigma_1 ... \sigma_N)}{\tau} & 0_{NP\times 1} \\ 0_{N\times NP} & 0_{N\times 1} \\ ..... & \\ 0_{1\times NP} & c \end{pmatrix} \tag{5.11}$$

where χ_i are the prior means for the coefficients on the first lags of the dependent variables (these can be different from 1) and $J_P = diag(1..P)$

5.1.2. *Creating dummy variables for the sum of coefficients prior.* If the variables in the VAR have a unit root, this information can be reflected via a prior that incorporates the belief that coefficients on lags of the dependent variable sum to 1 (see Robertson and Tallman (1999)). This prior can be implemented in our example VAR via the following dummy observations

$$Y_{D,5} = \begin{pmatrix} \gamma\mu_1 & 0 \\ 0 & \gamma\mu_2 \end{pmatrix}, X_{D,5} = \begin{pmatrix} 0 & \gamma\mu_1 & 0 & \gamma\mu_1 & 0 \\ 0 & 0 & \gamma\mu_2 & 0 & \gamma\mu_2 \end{pmatrix} \tag{5.12}$$

where μ_1 is the sample mean of y_t and μ_2 is the sample mean of x_t possibly calculated using an intial sample of data. Note that these dummy observations imply prior means of the form $b_{ii} + d_{ii} = 1$ where $i = 1, 2$ and γ controls the tightness of the prior. As $\gamma \to \infty$ the prior is implemented more tightly. Banbura *et al.* (2007) show that these dummy observations for a N variable VAR with P lags are given as

$$Y_D = \frac{diag\left(\chi_1\mu_1 ... \chi_N\mu_N\right)}{\lambda}, X_D = \begin{pmatrix} \frac{(1,2..P)\otimes diag(\chi_1\mu_1 ... \chi_N\mu_N)}{\lambda} & 0_{N\times 1} \end{pmatrix} \tag{5.13}$$

where $\lambda = 1/\gamma$ and $\mu_i, i = 1, .., N$ are sample means of each variable included in the VAR.

5.1.3. *Creating dummy variables for the common stochastic trends prior.* One can express the prior belief that the variables in the VAR have a common stochastic trend via the following dummy observations

$$Y_{D,6} = \begin{pmatrix} \delta\mu_1 & \delta\mu_2 \end{pmatrix}, X_{D,6} = \begin{pmatrix} \delta & \delta\mu_1 & \delta\mu_2 & \delta\mu_1 & \delta\mu_2 \end{pmatrix} \tag{5.14}$$

These dummy observations imply, for example, that $\mu_1 = c_1 + \mu_1 b_{11} + \mu_2 b_{12} + \mu_1 d_{11} + \mu_2 d_{12}$ i.e. the mean of the first variable is a combination of μ_1 and μ_2. Note as $\delta \to \infty$ the prior is implemented more tightly and the series in the VAR share a common stochastic trend.

5.1.4. *Matlab code for implementing priors using dummy observations.* Figures 14, 15 and 16 show the matlab code for the bi-variate VAR(2) model using quarterly data on annual GDP growth and CPI inflation for the US from 1948Q2 to 2010Q4 (example4.m). Line 26 of the code calculates the sample means of the data to be used in setting the dummy observations. Some researchers use a pre-sample to calculate these means and the standard deviations σ_i. Lines 28 to 32 specify the parameters that control the prior. Lines 33 to 37 set the dummy observations for the VAR coefficients on the first lags. Lines 38 to 42 set the dummy observations for the VAR coefficients on the second lag. Lines 43 to 47 specify the dummy observations for the prior on the constant. Lines 48 to 53 specify

```
1 clear
2 addpath('functions');
3 % a bi-variate VAR with dummy variable implementation of priors
4 %load data
5 data=xlsread('\data\datain.xls'); %data for US GDP growth and
inflation 1948q1 2010q4
6 N=size(data,2);
7 L=2;   %number of lags in the VAR
8 Y=data;
9 X=[ones(size(Y,1),1) lag0(data,1) lag0(data,2) ];
10 Y=Y(3:end,:);
11 X=X(3:end,:);
12 T=rows(X);
13 %compute standard deviation of each series residual via an ols
regression
14 %to be used in setting the prior
15 %first variable
16 y=Y(:,1);
17 x=[ones(T,1) X(:,2)];
18 b0=inv(x'*x)*(x'*y);
19 s1=sqrt(((y-x*b0)'*(y-x*b0))/(rows(y)-2));  %std of residual standard
error
20 %second variable
21 y=Y(:,2);
22 x=[ones(T,1) X(:,3)];
23 b0=inv(x'*x)*(x'*y);
24 s2=sqrt(((y-x*b0)'*(y-x*b0))/(rows(y)-2));
25 %mean of the data
26 mu=mean(Y);  μ_i, i = 1,..,N are sample means of each variable included in the VAR
27 %specify parameters of the minnesota prior
28 tau=0.1;  %controls prior on own first lags
τ controls the overall tightness of the prior
29 d=1;    %decay for higher lags   d controls the tightness of the prior on higher lags
30 lamdac=1;  %prior for the constant
controls the tightness of the prior on constants
31 lamda=1;   %sum of coefficients unit roots
As γ → ∞ the prior is implemented more tightly
32 delta=1;  %cointegration prior   as δ → ∞ the prior is implemented more tightly
```

$$Y_{D,1} = \begin{pmatrix} (1/\tau)\sigma_1 & 0 \\ 0 & (1/\tau)\sigma_2 \end{pmatrix}$$

$$X_{D,1} = \begin{pmatrix} 0 & (1/\tau)\sigma_1 & 0 & 0 & 0 \\ 0 & 0 & (1/\tau)\sigma_2 & 0 & 0 \end{pmatrix}$$

```
33 %specify dummy observations for first lag
34 yd1=[(1/tau)*s1 0;
35     0        (1/tau)*s2];
36 xd1=[0 (1/tau)*s1 0 0 0;
37     0 0 (1/tau)*s2 0 0];
```

FIGURE 14. Normal Wishart prior using dummy observations

the dummy observations for the unit root prior. Lines 56 to 57 set out the dummy observations for the common stochastic trends prior. Lines 59 to 64 specify the dummy observations for the prior on the covariance matrix. Lines 68 and 69 mix the actual observations with the dummy data creating $Y^* = [Y; Y_D], X^* = [X; X_D]$. Line 72 computes the mean of the conditional posterior distribution of the VAR coefficients $B^* = (X^{*\prime}X^*)^{-1}(X^{*\prime}Y^*)$. Line 83 calculates the variance of this posterior distribution $\Sigma \otimes (X^{*\prime}X^*)^{-1}$ and line 85 draws the VAR coefficients from the normal distribution with this mean and variance. Line 88 calculates the scale matrix for the inverse Wishart density $S^* = (Y^* - X^*B^*)'(Y^* - X^*B^*)$ and line 89 draws the covariance matrix from the inverse Wishart distribution. Once past the burn-in stage the code forecasts the two variables in the VAR and builds up the predictive density.

$$Y_{D,2} = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}$$

$$X_{D,2} = \begin{pmatrix} 0 & 0 & 0 & (1/\tau)\sigma_1 2^d & 0 \\ 0 & 0 & 0 & 0 & (1/\tau)\sigma_2 2^d \end{pmatrix}$$

```
%specify dummies for second lag
yd2=[0  0;
     0   0];
xd2=[0 0 0 (1/tau)*s1*2^d  0;
     0 0  0  0 (1/tau)*s2*2^d ];
```

$$Y_{D,3} = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}$$

$$X_{D,3} = \begin{pmatrix} 1/c & 0 & 0 & 0 & 0 \\ 1/c & 0 & 0 & 0 & 0 \end{pmatrix}$$

```
%specify priors for the constants
yd3=[0  0;
     0   0];
xd3=[1/lamdac 0 0 0 0;
     1/lamdac 0  0  0 0];
```

$$Y_{D,5} = \begin{pmatrix} \gamma\mu_1 & 0 \\ 0 & \gamma\mu_2 \end{pmatrix}, X_{D,5} = \begin{pmatrix} 0 & \gamma\mu_1 & 0 & \gamma\mu_1 & 0 \\ 0 & 0 & \gamma\mu_2 & 0 & \gamma\mu_2 \end{pmatrix}$$

```
%specify sum of coefficient dummies
yd4=[lamda*mu(1) 0;
      0   lamda*mu(2)];

 xd4=[0 lamda*mu(1) 0 lamda*mu(1) 0;
      0   0   lamda*mu(2) 0 lamda*mu(2)];
```

$$Y_{D,6} = \begin{pmatrix} \delta\mu_1 & \delta\mu_2 \end{pmatrix}, X_{D,6} = \begin{pmatrix} \delta & \delta\mu_1 & \delta\mu_2 & \delta\mu_1 & \delta\mu_2 \end{pmatrix}$$

```
 %specify common stochastic trend dummies
 yd5=[delta*mu(1) delta*mu(2)];
 xd5=[delta delta*mu(1) delta*mu(2) delta*mu(1) delta*mu(2)];
```

$$Y_{D,4} = \begin{pmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{pmatrix}$$

$$X_{D,4} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

```
 %specify dummy variables for covariance matrix
 yd6=[s1 0;
      0   s2];
```

FIGURE 15. Normal Wishart prior using dummy observations continued

6. Application1: Structural VARs and sign restrictions

Structural VAR models offer a simple and flexible framework for analysing several questions of interest. Once structural shocks are identified using an appropriate identification scheme, impulse response analysis, variance decomposition and historical decomposition offer powerful tools. For a detailed explanation of structural VARs see Hamilton (1994) or Canova (2007). In this section we focus on how structural analysis fits in the Gibbs sampling framework established in the chapter.

As shown in the matlab example in section 3.2.1, one can estimate structural VARs by calculating the structural impact matrix A_0 (where $\Sigma = A_0' A_0$) for each retained Gibbs draw and use this to compute impulse response

```
63  xd6=[0 0 0 0 0;
64       0   0  0 0 0];
65 %all dummy observations
66 yd=[yd1;yd2;yd3;yd4;yd5;yd6];
67 xd=[xd1;xd2;xd3;xd4;xd5;xd6];
68 Ystar=[Y;yd];   % Y* = [Y;Y_D], X* = [X;X_D]
69 Xstar=[X;xd];
70 Tstar=rows(Xstar);
71 %compute posterior mean
72 betahat=inv(Xstar'*Xstar)*(Xstar'*Ystar);
73 %compute inital value of sigma
74 e=Ystar-Xstar*betahat;
75 sigma=(e'*e)/Tstar;
76 REPS=2000;
77 BURN=1000;
78 %gibbs algorithm
79 out1=[];
80 out2=[];
81 for i=1:REPS
82     M=vec(betahat);   % vec(B*)
83     V=kron(sigma,inv(Xstar'*Xstar));   % Σ ⊗ (X*'X*)^-1
84     %draw beta
85     beta=M+(randn(1,N*(N*L+1))*chol(V))';
86     %draw sigma
87     e=Ystar-Xstar*reshape(beta,N*L+1,N);
88     scale=e'*e;   % S* = (Y* - X*B*)'(Y* - X*B*)
89     sigma=iwpq(Tstar,inv(scale));
90
91     if i>BURN;
92 %forecast
93 yhat=zeros(14,2);
94    yhat(1:2,:)=Y(end-1:end,:);
95    for i=3:14
96     yhat(i,:)=[1 yhat(i-1,:) yhat(i-
2,:)]*reshape(beta,N*L+1,N)+randn(1,N)*chol(sigma);
97 end
98 out1=[out1 [Y(:,1);yhat(3:end,1)]];
99 out2=[out2 [Y(:,2);yhat(3:end,2)]];
100 end
101 end
102
103 TT=1948.75:0.25:2014;
104 subplot(1,2,1)
105 plot(TT,prctile(out1,[50 10 20 30 70 80 90],2))
106 xlim([1995 2015])
107 title('GDP Growth');
108 subplot(1,2,2)
109 plot(TT,prctile(out2,[50 10 20 30 70 80 90],2))
110 xlim([1995 2015])
111 legend('Median Forecast','10th percentile','20th percentile','30th
percentile','70th percentile','80th percentile','90th percentile');
112 title('Inflation');
```

Published with MATLAB® 7.9

FIGURE 16. Normal Wishart prior using dummy observations continued

functions, variance decompositions and historical decompositions. The Gibbs sampling framework is convenient because it allows one to build a distribution for these objects (i.e. impulse response functions, variance decompositions and historical decompositions) and thus characterise uncertainty about these estimates.

Strictly speaking, this indirect method of estimating structural VARs–i.e. calculating A_0 using a Gibbs draw of Σ (and not sampling A_0 directly) provides the posterior distribution of A_0 only if the structural VAR is exactly identified (for e.g. when A_0 is calculated using a Cholesky decomposition as in section 3.2.1) . In the case of over identification one needs to estimate the posterior of A_0 directly (see Sims and Zha (1998)). We will consider such an example in Chapter 4.

Recent applications of structural VARs have used sign restrictions to identify structural shocks (for a critical survey see Fry and Pagan (2007)). Despite the issues raised in Sims and Zha (1998), sign restrictions are implemented using an indirect algorithm. In other words for each retained draw of Σ one calculates an A_0 matrix which results in impulse responses to a shock of interest with signs that are consistent with theory. For example to identify a monetary policy shock one may want an A_0 matrix that leads to a response of output and inflation that is negative and a response of the policy interest rate that is positive for a few periods after the shock.

Ramirez *et al.* (2010) provide an efficient algorithm to find an A_0 matrix consistent with impulse responses of a certain sign consistent with theory. We review this algorithm by considering the following VAR(1) model as an example

$$\begin{pmatrix} Y_t \\ P_t \\ R_t \end{pmatrix} = \begin{pmatrix} c_1 \\ c_2 \\ c_3 \end{pmatrix} + \begin{pmatrix} b_{11} & b_{12} & b_{13} \\ b_{21} & b_{22} & b_{23} \\ b_{31} & b_{32} & b_{33} \end{pmatrix} \begin{pmatrix} Y_{t-1} \\ P_{t-1} \\ R_{t-1} \end{pmatrix} + \begin{pmatrix} v_{1t} \\ v_{2t} \\ v_{3t} \end{pmatrix} \tag{6.1}$$

where $var \begin{pmatrix} v_{1t} \\ v_{2t} \\ v_{3t} \end{pmatrix} = \Sigma, Y_t$ is output growth, P_t is inflation and R_t is the interest rate. The aim is to calculate the impulse response to a monetary policy shock. The monetary policy shock is assumed to be one that decreases Y_t and P_t and increases R_t in the period of the shock. As described above, the Gibbs sampling algorithm to estimate the parameters of the VAR model cycles through two steps, sampling successively from $H\left(b \backslash \Sigma, Y_t\right)$ and $H\left(\Sigma \backslash b, Y_t\right)$. Once past the burn-in stage the following steps are used calculate the required A_0 matrix:

Step 1 Draw a $N \times N$ matrix K from the standard normal distribution

Step 2 Calculate the matrix Q from the QR decomposition of K. Note that Q is orthonormal i.e. $Q'Q = I$.

Step 3 Calculate the Cholesky decomposition of the current draw of $\Sigma = \tilde{A}_0' \tilde{A}_0$

Step 4 Calculate the candidate A_0 matrix as $A_0 = Q\tilde{A}_0$. Note that because $Q'Q = I$ this implies that $A_0'A_0$ will still equal Σ. By calculating the product $Q\tilde{A}_0$ we alter the elements of $\tilde{A}_0$ but not the property that $\Sigma = \tilde{A}_0'\tilde{A}_0$. The candidate A_0 matrix in our 3 variable VAR example will have the following form

$$A_0 = \begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{pmatrix}$$

The third row of this matrix corresponds with the interest rate shock. We need to check if $a_{31} < 0$ and $a_{32} < 0$ and $a_{33} > 0$. If this is the case a contemporaneous increase in R_t will lead to a fall in Y_t and P_t as the elements$\begin{pmatrix} a_{31} & a_{32} & a_{33} \end{pmatrix}$ correspond to the current period impulse response of Y_t, P_t and R_t respectively. If $a_{31} < 0$ and $a_{32} < 0$ and $a_{33} > 0$ we stop and use this A_0 matrix to compute impulse responses and other objects of interest. If the restriction is not satisfied we go to step 1 and try with an new K matrix.

Step 5 Repeat steps 1 to 4 for every retained Gibbs draw.

6.1. A Structural VAR with sign restrictions in matlab. We estimate a large scale VAR model for the US using quarterly data from 1971Q1 to 2010Q4 (example5.m). The VAR model includes the following variables (in this order): (1) Federal Funds Rate (2) Annual GDP growth (3) Annual CPI Inflation (4) Annual real consumption growth (5) Unemployment rate (6) change in private investment (7) net exports (8) annual growth in M2 (9) 10 year government bond yield (10) annual growth in stock prices (11) annual growth in the yen dollar exchange rate. We identify a monetary policy shock by assuming that a monetary contraction has the following contemporaneous effects

Variable	Sign restriction
Federal Funds Rate	+
Annual GDP growth	-
Annual CPI Inflation	-
Annual Real Consumption Growth	-
Unemployment rate	+
Annual Investment Growth	-
Annual Money Growth	-

The Matlab code for this example is shown in figures 17, 18 and 19. Line 37 of the code builds the dummy observations for the normal wishart prior for the VAR using equations 5.11 and the sum of coefficients prior using equation 5.13 via the function create_dummies.m in function folder. Lines 49 to 55 sample from the conditional posterior distributions as in the previous example. Once past the burn-in stage, on line 61 we draw a $N \times N$ matrix from the standard normal distribution. Line 62 takes the QR decomposition of K and obtaines the matrix Q. Line 63 calculates the Cholesky decomposition of Σ while line 64 calculates the candidate A_0 matrix as $A_0 = Q\tilde{A}_0$. Lines 66 to 72 check if the sign restrictions are satisfied by checking the elements of the first row of the A_0 matrix (the row that corresponds the interest rate). Lines 77 to 83 check if the sign restrictions are satisfied with the sign reversed. If they are, we multiply the entire first row of the A_0 matrix by -1. The code keeps on drawing K and calculating candidate $A_0 = Q\tilde{A}_0$ matrices until an A_0 matrix is found that satisfies the sign restrictions. Once an A_0 matrix is

```
clear
addpath('functions');
REPS=5000;
BURN=3000;
[data,names]=xlsread('\data\usdata1.xls'); %load US data
N=cols(data);
L=2; %lag length of the VAR
Y=data;
%take lags
X=[];
for j=1:L
X=[X lag0(data,j) ];
end
X=[X ones(rows(X),1)];
Y=Y(L+1:end,:);
X=X(L+1:end,:);
T=rows(X);
%priors for VAR coefficients
lamdaP=1;  %This controls the tightness of the priors on the first
lag
tauP=10*lamdaP;  % this controls the tightness of the priors on sum
of coefficients
epsilonP=1;  % this controls tightness of the prior on the constant
muP=mean(Y)';
sigmaP=[];
deltaP=[];
for i=1:N
    ytemp=Y(:,i);
    xtemp=[lag0(ytemp,1) ones(rows(ytemp),1)];
    ytemp=ytemp(2:end,:);
    xtemp=xtemp(2:end,:);
    btemp=xtemp\ytemp;
    etemp=ytemp-xtemp*btemp;
    stemp=etemp'*etemp/rows(ytemp);
    deltaP=[deltaP;btemp(1)];
    sigmaP=[sigmaP;stemp];
end
%dummy data to implement priors see
http://ideas.repec.org/p/ecb/ecbwps/20080966.html
This function (create dummies) builds the dummy observations using:
```

$$Y_D = \begin{pmatrix} \frac{diag(\chi_1\sigma_1 \ldots \chi_N\sigma_N)}{\tau} \\ 0_{N\times(P-1)\times N} \\ \ldots\ldots\ldots\ldots \\ diag(\sigma_1 \ldots \sigma_N) \\ \ldots\ldots\ldots\ldots \\ 0_{1\times N} \end{pmatrix}, X_D = \begin{pmatrix} \frac{J_P \otimes diag(\sigma_1 \ldots \sigma_N)}{\tau} & 0_{NP\times 1} \\ 0_{N\times NP} & 0_{N\times 1} \\ \ldots\ldots\ldots\ldots & \\ 0_{1\times NP} & c \end{pmatrix}$$

$$Y_D = \frac{diag(\chi_1\mu_1 \ldots \chi_N\mu_N)}{\lambda}, X_D = \begin{pmatrix} \frac{(1,2..P)\otimes diag(\chi_1\mu_1 \ldots \chi_N\mu_N)}{\lambda} & 0_{N\times 1} \end{pmatrix}$$

```
[yd,xd] = create dummies(lamdaP,tauP,deltaP,epsilonP,L,muP,sigmaP,N);
%yd and xd are the dummy data. Append this to actual data
  Y0=[Y;yd];
  X0=[X;xd];
  %conditional mean of the VAR coefficients
  mstar=vec(X0\Y0);  %ols on the appended data
  xx=X0'*X0;
```

FIGURE 17. A Structural VAR with sign restrictions: Matlab Code

found that satisfies the sign restrictions, this is used to calculate the impulse response to a monetary policy shock and the impulse response functions for each retained Gibbs draw are saved. The file example6.m has exactly the same code but makes the algorithm to find the A_0 matrix more efficient by searching all rows of candidate $A_0 = Q\tilde{A}_0$ matrix for the one consistent with the policy shock–i.e. with the signs as in the table above. Once this is found we insert this row into the first row of the candidate A_0 matrix. Note that that this re-shuffling of rows does not alter the property that $\Sigma = A_0'A_0$. Note also that the A_0 matrix is not unique. That is, one could find A_0 matrices that satisfy the sign restrictions but have elements of different magnitude. Some researchers deal with this issue by generating M A_0 matrices that satisfy the sign restrictions for each Gibbs draw and then retaining the A_0 matrix that is closest to

```
44   ixx=xx\eye(cols(xx));  %inv(X0'X0) to be used later in the Gibbs
sampling algorithm
45   sigma=eye(N); %starting value for sigma
46   out=zeros(REPS-BURN,36,N);
47   jj=1;
48   for i=1:REPS
49        vstar=kron(sigma,ixx);
50        beta=mstar+(randn(1,N*(N*L+1))*chol(vstar))';
51
52        %draw covariance
53        e=Y0-X0*reshape(beta,N*L+1,N);
54     scale=e'*e;
55     sigma=iwpq(T+rows(yd),inv(scale));
56
57     if i>=BURN
58         %impose sign restrictions
59         chck=-1;
60         while chck<0
61         K=randn(N,N);
62         Q=getQR(K);
```

Draw a $N \times N$ matrix K from the standard normal distribution

Calculate the matrix Q from the QR decomposition of K. Note that Q is orthonormal i.e. $Q'Q = I$.

```
63         A0hat=chol(sigma);
```

Calculate the Cholesky decomposition of the current draw of $\Sigma = \tilde{A}_0'\tilde{A}_0$

```
64         A0hat1=(Q*A0hat);  %candidate draw
```

Calculate the candidate A_0 matrix as $A_0 = Q\tilde{A}_0$. Note that because $Q'Q = I$ this implies that $A_0'A_0$ will still equal Σ

```
65         %check signs
66         e1=A0hat1(1,1)>0;  %Response of R
67         e2=A0hat1(1,2)<0;  %Response of Y
68         e3=A0hat1(1,3)<0;  %Response of Inflation
69         e4=A0hat1(1,4)<0;  %Response of consumption
70         e5=A0hat1(1,5)>0;  %Response of U
71         e6=A0hat1(1,6)<0;  %Response of investment
72         e7=A0hat1(1,8)<0; %response of money
73         if e1+e2+e3+e4+e5+e6+e7==7
74             chck=10;
75         else
76             %check signs but reverse them
77         e1=-A0hat1(1,1)>0;  %Response of R
78         e2=-A0hat1(1,2)<0;  %Response of Y
79         e3=-A0hat1(1,3)<0;  %Response of Inflation
80         e4=-A0hat1(1,4)<0;  %Response of consumption
81         e5=-A0hat1(1,5)>0;  %Response of U
82         e6=-A0hat1(1,6)<0;  %Response of investment
83         e7=-A0hat1(1,8)<0; %response of money
84         if e1+e2+e3+e4+e5+e6+e7==7
85             A0hat1(1,1:N)=-A0hat1(1,1:N);
86             chck=10;
87         end
88         end
89         end
90
91 yhat=zeros(36,N);
92 vhat=zeros(36,N);
93 vhat(3,1)=1; %shock to the Federal Funds rate
94 for j=3:36
95  yhat(j,:)=[ yhat(j-1,:) yhat(j-2,:)
0]*reshape(beta,N*L+1,N)+vhat(j,:)*A0hat1;
96 end
```

FIGURE 18. A structural VAR with sign restrictions (continued)

the mean or median of these M matrices. This imples that one restricts the distribution of the selected A_0 matrices via a (arbitrary) rule. The file example7.m does this for our example by generating 100 A_0 matrices for each retained Gibbs draw and using the A_0 matrix closest to the median to compute the impulse response functions. Figure 20 shows the estimated impulse response functions computed using example7.m

7. Application 2: Conditional forecasting using VARs and Gibbs sampling

In many cases (relevant to central bank applications) forecasts of macroeconomic variables that are conditioned on fixed paths for other variables is required. For example, one may wish to forecast credit and asset prices assuming

```
97  out(jj,:,:)=yhat;
98  jj=jj+1;
99     end
100
101   end
```

Published with MATLAB® 7.9

FIGURE 19. A structural VAR with sign restrictions: Matlab code (continued)

that inflation and GDP growth follow future paths fixed at the official central bank forecast. Waggoner and Zha (1999) provide a convenient framework to calculate not only the conditional forecasts but also the forecast distribution using a Gibbs sampling algorithm.

To see their approach consider a simple VAR(1) model

$$Y_t = c + BY_{t-1} + A_0\varepsilon_t \tag{7.1}$$

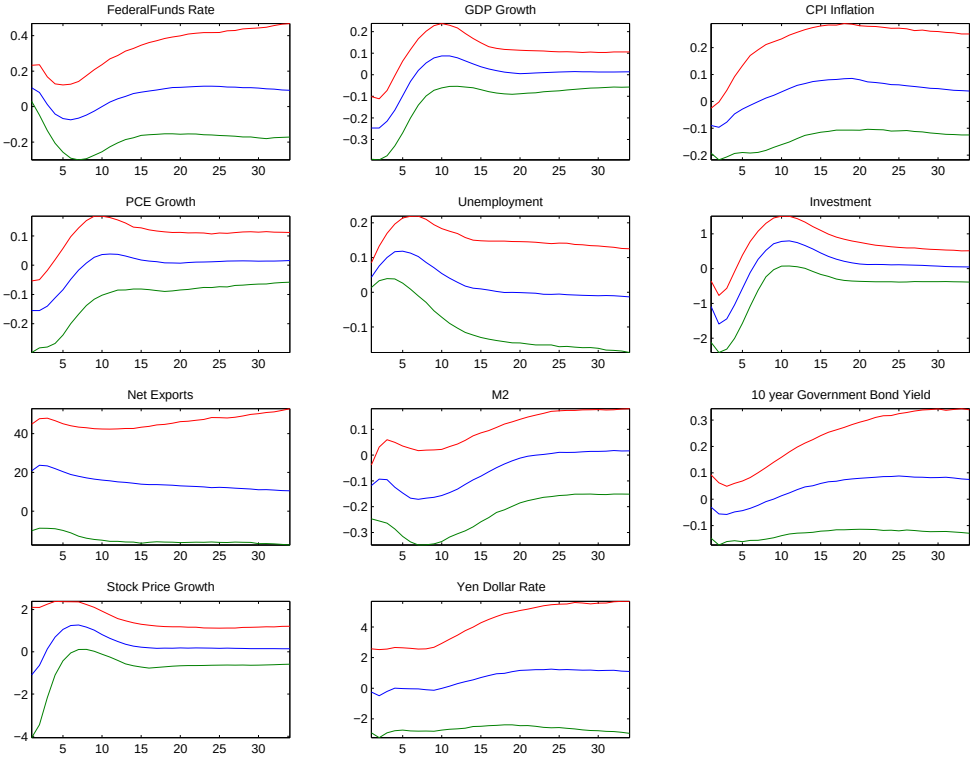

FIGURE 20. Impulse response to a monetary policy shock using sign restrictions

where Y_t denotes a $T \times N$ matrix of endogenoeus variables, ε_t are the uncorrelated structural shocks and $A_0 A_0' = \Sigma$ where Σ denotes the variance of the reduced form VAR esiduals. Iterating equation 7.1 forward K times we obtain

$$Y_{t+K} = c\sum_{j=0}^{K} B^j + B^j Y_{t-1} + A_0 \sum_{j=0}^{K} B^j \varepsilon_{t+K-j} \tag{7.2}$$

Equation 7.2 shows that the K period ahead forecast Y_{t+K} can be decomposed into components with and without structural shocks. The key point to note is that if a restriction is placed on the future path of the J^{th} variable in Y_t, this implies restrictions on the future shocks to the other variables in the system. This can easily be seen by re-arranging equation 7.2

$$Y_{t+K} - c\sum_{j=0}^{K} B^j - B^j Y_{t-1} = A_0 \sum_{j=0}^{K} B^j \varepsilon_{t+K-j} \tag{7.3}$$

If some of the variables in Y_{t+K} are constrained to follow a fixed path, this implies restrictions on the future innovations on the RHS of equation 7.3. Waggoner and Zha (1999) express these constraints on future innovations as

$$R\varepsilon = r \tag{7.4}$$

where r is a $(M \times k) \times 1$ vector where M are the number of constrained variables and k denotes the number of periods the constraint is applied. The elements of the vector r are the path for the constrained variables minus the unconditional forecast of the constrained variables. R is a matrix with dimensions $(M \times k) \times (N \times k)$. The elements of this matrix are the impulse responses of the constrained variables to the structural shocks ε at horizon $1, 2...k$. The $(N \times k) \times 1$ vector ε contains the constrained future shocks. We give a detailed example showing the structure of these matrices below.

Doan *et al.* (1983) show that a least square solution for the constrained innovations in equation 7.4 is given as

$$\hat{\varepsilon} = R'(R'R)^{-1} r \tag{7.5}$$

With these constrained shocks $\hat{\varepsilon}$ in hand, the conditional forecasts can be calculated by substituting these in equation7.2.

7.1. Calculating conditional forecasts. To see the details of this calculation, consider the following VAR model with two endogenous variables

$$\begin{pmatrix} Y_t \\ X_t \end{pmatrix} = \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} + \begin{pmatrix} B_1 & B_2 \\ B_3 & B_4 \end{pmatrix} \begin{pmatrix} Y_{t-1} \\ X_{t-1} \end{pmatrix} + \begin{pmatrix} A_{11} & \\ A_{12} & A_{22} \end{pmatrix} \begin{pmatrix} \varepsilon_{1t} \\ \varepsilon_{2t} \end{pmatrix} \tag{7.6}$$

In addition denote $z^k_{i,j}$ as the impulse response of the j^{th} variable at horizon i to the k^{th} structural shock where $k = 1, 2$. Consider forecasting Y_t three periods in the future using the estimated VAR in equation 7.6. However we impose the condition that $\begin{pmatrix} \hat{X}_{t+1} \\ \hat{X}_{t+2} \\ \hat{X}_{t+3} \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$, i.e. variable X is fixed at 1 over the forecast horizon. In order to calculate the forecast for Y_t under this condition, the first step involves using equation 7.5 to calculate the restricted structural shocks. Using equation 7.5 requires building the matrices R and r. We now describe the structure of these matrices for our example. First note that the restricted structural shocks (to be calculated) are stacked as

$$\hat{\varepsilon} = \begin{pmatrix} \hat{\varepsilon}_{1t+1} \\ \hat{\varepsilon}_{2t+1} \\ \hat{\varepsilon}_{1t+2} \\ \hat{\varepsilon}_{2t+2} \\ \hat{\varepsilon}_{1t+3} \\ \hat{\varepsilon}_{2t+3} \end{pmatrix} \tag{7.7}$$

The matrix of impulse responses R is built to be compatible with $\hat{\varepsilon}$ (see equation 7.4). In this example, it has the following structure

$$R = \begin{pmatrix} z^1_{1,2} & z^2_{1,2} & 0 & 0 & 0 & 0 \\ z^1_{2,2} & z^2_{2,2} & z^1_{1,2} & z^2_{1,2} & 0 & 0 \\ z^1_{3,2} & z^2_{3,2} & z^1_{2,2} & z^2_{2,2} & z^1_{1,2} & z^2_{1,2} \end{pmatrix} \tag{7.8}$$

The matrix R is made of the response of the constrained variable 2 (i.e. X) to the two structural shocks. The first row of the matrix has the response of X to ε_1 and ε_2 at horizon 1. Note that this row corresponds to the first two elements in $\hat{\varepsilon}$– it links the constrained shocks 1 period ahead to their responses. The second row of R has this impulse response at horizon 2 (first two elements) and then at horizon 1 (third and fourth element). This row corresponds to the forecast two periods ahead and links the structural shocks at horizon 1 and 2 to their respective impulse responses. A similar interpretation applies to the subsequent rows of this matrix.

The matrix r is given as

$$r = \begin{pmatrix} 1 - \tilde{X}_{t+1} \\ 1 - \tilde{X}_{t+2} \\ 1 - \tilde{X}_{t+3} \end{pmatrix} \tag{7.9}$$

where $\tilde{X}_{t+j}$ denotes the unconditional forecast of X. Once these matrices are constructed, the restricted structural shocks are calculated as $\hat{\varepsilon} = R'(R'R)^{-1}r$. These are then used calculate the conditional forecast by substituting them in equation 7.6 and iterating forward.

In figures 21 and 22 we show the matlab code for this simple example of calculating a conditional forecast (the matlab file is example8.m). We estimate a VAR(2) model for US GDP growth and inflation and use the estimated VAR to forecast GDP growth 3 periods ahead assuming inflation remains fixed at 1% over the forecast horizon.Lines 18 to 21 of the code estimate the VAR coefficients and error covariance via OLS and calculate A_0 as the Choleski decomposition of the error covariance matrix. As shown in Waggoner and Zha (1999) the choice of identifying restrictions (i.e. the structure of A_0) does not affect the conditional forecast which depends on the reduced form VAR. Therefore it is convenient to use the Choleski decomposition to calculate A_0 for this application. Lines 25 to 28 estimate the impulse response functions $z^k_{i,j}$. Line 33 to 39 constructs the unconditional forecast by simulating the estimated VAR model for three periods. Lines 41 to 43 construct the R matrix as specified in equation 7.8. Line 45 constructs the r matrix. With these in hand, the restricted future shocks are calculated on line 50. The conditional forecast is calculated by simulating the VAR using these restricted shocks 50 to 59 of the code with the matlab variable yhat2 holding the conditional forecast.

7.2. Calculating the distribution of the conditional forecast. The main contribution of Waggoner and Zha (1999) is to provide a Gibbs sampling algorithm to construct the distribution of the conditional forecast and thus allow a straigth forward construction of fan charts whn some forecasts are subject to constraints. In particular, Waggoner and Zha (1999) show that the distribution of the restricted future shocks ε is normal with mean $\bar{M}$ and variance $\bar{V}$ where

$$\begin{aligned} \bar{M} &= R'(R'R)^{-1}r \\ \bar{V} &= I - R'(RR')^{-1}R \end{aligned} \tag{7.10}$$

The Gibbs sampling algorithm to generate the forecast distribution proceeds in the following steps

(1) Initialise the VAR coefficients and the A_0 matrix.

```
clear
addpath('functions');
data=xlsread('\data\datain.xls'); %data for US GDP growth and inflation
1948q1 2010q4
N=cols(data);
horizon=3;
path=[1;1;1]; %constrained values for X
L=2; %lag length of the VAR
Y=data;
%take lags
X=[];
for j=1:L
X=[X lag0(data,j) ];
end
X=[X ones(rows(X),1)];
Y=Y(L+1:end,:);
X=X(L+1:end,:);
T=rows(X);
B=X\Y;  %ols estimate
res=Y-X*B;
sigma=(res'*res)/T;
A0=chol(sigma);
%calculate impulse responses to be used to construct R
S=zeros(1,N);
S(1)=1; %shock to first eq
Z1=irfsim(B,N,L,A0,S,horizon+L); % z_{ij}^k
S=zeros(1,N);
S(2)=1; %shock to 2nd eq
Z2=irfsim(B,N,L,A0,S,horizon+L); % z_{ij}^k
%calculate unconditional forecast to be used to construct r
yhat1=zeros(horizon+L,N);
yhat1(1:L,:)=Y(end-L+1:end,:);
for i=L+1:horizon+L
    x=[];
    for j=1:L
    x=[x yhat1(i-j,:)];
    end
    yhat1(i,:)=[x 1]*B;
end
yhat1=yhat1(L+1:end,:);
```

$$R=\begin{pmatrix} z_{1,2}^1 & z_{1,2}^2 & 0 & 0 & 0 & 0 \\ z_{2,2}^1 & z_{2,2}^2 & z_{1,2}^1 & z_{1,2}^2 & 0 & 0 \\ z_{3,2}^1 & z_{3,2}^2 & z_{2,2}^1 & z_{2,2}^2 & z_{1,2}^1 & z_{1,2}^2 \end{pmatrix}$$

```
%construct the R matrix
R=[Z1(1,2) Z2(1,2) 0  0  0  0;
    Z1(2,2) Z2(2,2) Z1(1,2) Z2(1,2) 0  0;
    Z1(3,2) Z2(3,2) Z1(2,2) Z2(2,2) Z1(1,2) Z2(1,2)];
%construct the r matrix
```

$$r=\begin{pmatrix} 1-\tilde{X}_{t+1} \\ 1-\tilde{X}_{t+2} \\ 1-\tilde{X}_{t+3} \end{pmatrix}$$

```
r=path-yhat1(:,2);
%compute the restricted structural shocks
```

FIGURE 21. Matlab code for computing the conditional forecast

(2) Form the matrices R and r. Draw the restricted structural shocks from the $N(\bar{M},\bar{V})$ distribution where $\bar{M}$ and $\bar{V}$ are calculated as in equation 7.10. This draw of structural shocks is used to calculate the conditional forecast $\hat{Y}_{t+k}$.

(3) Construct the appended dataset $Y_t^* = [Y_t; \hat{Y}_{t+k}]$. This the actual data for the VAR model with the forecasts added to it. The conditional posterior of the VAR coefficients and covariance matrix is construced using Y_t^* and new values of the coefficients and covariance matrix are drawn. The A_0 matrix can be updated as the Cholesky decomposition of the new draw of the covariance matrix. Note that by using Y_t^* we ensure

$$\hat{\varepsilon} = R'(R'R)^{-1}r$$

```
47 ehat=R'*pinv(R*R')*r;
48 ehat=reshape(ehat,N,horizon)';
49 %compute the conditional forecast
50 yhat2=zeros(horizon+L,N);
51 yhat2(1:L,:)=Y(end-L+1:end,:);
52 for i=L+1:horizon+L
53     x=[];
54     for j=1:L
55     x=[x yhat2(i-j,:)];
56     end
57     yhat2(i,:)=[x 1]*B+ehat(i-L,:)*A0;
58 end
59 yhat2=yhat2(L+1:end,:);
60
61
```

Published with MATLAB® 7.9

FIGURE 22. Matlab code for computing the conditional forecast (continued)

that the draws of the VAR parameters take into account the restrictions $R\varepsilon = r$. This procedure therefore accounts for parameter uncertainty and the restrictions imposed on the forecasts by the researcher.

(4) Goto step 2 and repeat M times. The last R draws of $\hat{Y}_{t+k}$ can be used to construct the distribution of the forecast.

In order to demonstrate this algorithm we continue our Matlab example above and calculate the distribution of the GDP growth forecast, leaving the inflation forecast restricted at 1%.

The Matlab code is shown in figures 23 and 24. Note that this is a continuation of the code in the previous example from line 60. We use 5000 Gibbs iterations and discard the first 3000 as burn in. Line 70 of the code constructs the

```
60 %Gibbs sampling algorithm
61 REPS=5000;
62 BURN=3000;
63 out1=[]; %will hold forecast for GDP
64 out2=[]; %will hold forecast for inflation
65 yhatg=yhat2; %initialise conditional forecast
66 sig=sigma; %initialise error covariance
67 for igibbs=1:REPS
68
69     %step 1 DRAW VAR parameters
70  datag=[data;yhatg]; %appended data  Y_t^* = [Y_t; Ŷ_{t+k}]
71  YSTAR=datag;
72     %take lags
73 XSTAR=[];
74 for j=1:L
75 XSTAR=[XSTAR lag0(datag,j) ];
76 end
77 XSTAR=[XSTAR ones(rows(XSTAR),1)];
78 YSTAR=YSTAR(L+1:end,:);
79 XSTAR=XSTAR(L+1:end,:);
80 T=rows(XSTAR);
81 %conditional mean
82 M=vec(XSTAR\YSTAR);
83 %conditional variance
84 V=kron(sig,inv(XSTAR'*XSTAR));
85 bg=M+(randn(1,N*(N*L+1))*chol(V))';
86 bg1=reshape(bg,N*L+1,N);
87 %draw sigma from the IW distribution
88 e=YSTAR-XSTAR*bg1;
89 scale=e'*e;
90 sig=IWPQ(T,inv(scale));
91 %A0 matrix
92 A0g=chol(sig);
93 %step 2 Construct conditional forecast
94 %impulse responses
95 S=zeros(1,N);
96 S(1)=1; %shock to first eq
97 Z1=irfsim(bg1,N,L,A0g,S,horizon+L);
98 S=zeros(1,N);
99 S(2)=1; %shock to 2nd eq
100 Z2=irfsim(bg1,N,L,A0g,S,horizon+L);
101 %calculate unconditional forecast to be used to construct r
102 yhat1=zeros(horizon+L,N);
103 yhat1(1:L,:)=Y(end-L+1:end,:);
104 for i=L+1:horizon+L
105     x=[];
106     for j=1:L
107     x=[x yhat1(i-j,:)];
108     end
109     yhat1(i,:)=[x 1]*bg1;
110 end
111 yhat1=yhat1(L+1:end,:);
112 %construct the R matrix
113 R=[Z1(1,2) Z2(1,2) 0  0  0  0;
114     Z1(2,2) Z2(2,2) Z1(1,2) Z2(1,2) 0  0;
115     Z1(3,2) Z2(3,2) Z1(2,2) Z2(2,2) Z1(1,2) Z2(1,2)];
116 %construct the r matrix
117 r=path-yhat1(:,2);
118 %compute the mean of the distribution of restricted structural
shocks
```

FIGURE 23. Calculating the distribution of the conditional forecast via Gibbs sampling

appended dataset and lines 82 to 90 use this appended data to draw the VAR coefficients and covariance matrix from their conditional distributions. The impulse responses and unconditional forecasts based on this draw of the VAR coefficients and the new A_0 matrix are used to construct the R matrix and the r vector on lines 113 to 117. Lines 119 to 122 construct the mean and variance of the restricted structural shocks $\bar{M}$ and $\bar{V}$. On line 124 we draw the structural shocks from the $N(\bar{M}, \bar{V})$ distrbution and lines 126 to 136 use these to construct the conditional forecast. Once past the burn-in stage the conditional forecasts are saved in the matrices out1 and out2.

Running the code produces figure 25. The left panel of the figure shows the forecast distribution for GDP growth. The right panel shows the forecast for inflation which is restricted at 1% over the forecast horizon.

```
119 MBAR=R'*pinv(R*R')*r;  % \bar{M} = R'(R'R)^{-1}r
120 %compute the variance of the distribution of restricted structural
shocks
121 VBAR=R'*pinv(R*R')*R;
122 VBAR=eye(cols(VBAR))-VBAR;  % \bar{V} = I - R'(RR')^{-1}R
123 %draw structural shocks from the N(MBAR,VBAR) distribution
124  edraw=MBAR+(randn(1,rows(MBAR))*real(sqrtm(VBAR)))';
125  edraw=reshape(edraw,N,horizon)';
126 %conditional forecast using new draw of shocks
127 yhatg=zeros(horizon+L,N);
128 yhatg(1:L,:)=Y(end-L+1:end,:);
129 for i=L+1:horizon+L
130     x=[];
131     for j=1:L
132     x=[x yhatg(i-j,:)];
133     end
134     yhatg(i,:)=[x 1]*bg1+edraw(i-L,:)*A0g;
135 end
136 yhatg=yhatg(L+1:end,:);
137 if igibbs>BURN
138     out1=[out1;[Y(:,1);yhatg(:,1)]'];
139     out2=[out2;[Y(:,2);yhatg(:,2)]'];
140 end
141 end
142
143 TT=1948.75:0.25:2011+(.75);
144 subplot(1,2,1)
145 plot(TT,prctile(out1,[50 10 20 30 70 80 90],1))
146 xlim([1995 max(TT)+0.25])
147 title('GDP Growth');
148 subplot(1,2,2)
149 plot(TT,prctile(out2,[50 10 20 30 70 80 90],1))
150 xlim([1995 max(TT)+0.25])
151 legend('Median Forecast','10th percentile','20th percentile','30th
percentile','70th percentile','80th percentile','90th percentile');
152 title('Inflation');
```

Published with MATLAB® 7.9

FIGURE 24. Calculating the distribution of the conditional forecast via Gibbs sampling (continued)

7.3. Extensions and other issues. The example above places restrictions on both structural shocks ε_1 and ε_2 to produce the conditional forecast. In some applications it may be preferable to produce the conditional forecast by placing restrictions only on a subset of shocks. For instances one may wish to restrict ε_1 only in our application. This can be done easily by modifying the R matrix as follows:

$$R = \begin{pmatrix} z^1_{1,2} & 0 & 0 & 0 & 0 & 0 \\ z^1_{2,2} & 0 & z^1_{1,2} & 0 & 0 & 0 \\ z^1_{3,2} & 0 & z^1_{2,2} & 0 & z^1_{1,2} & 0 \end{pmatrix} \tag{7.11}$$

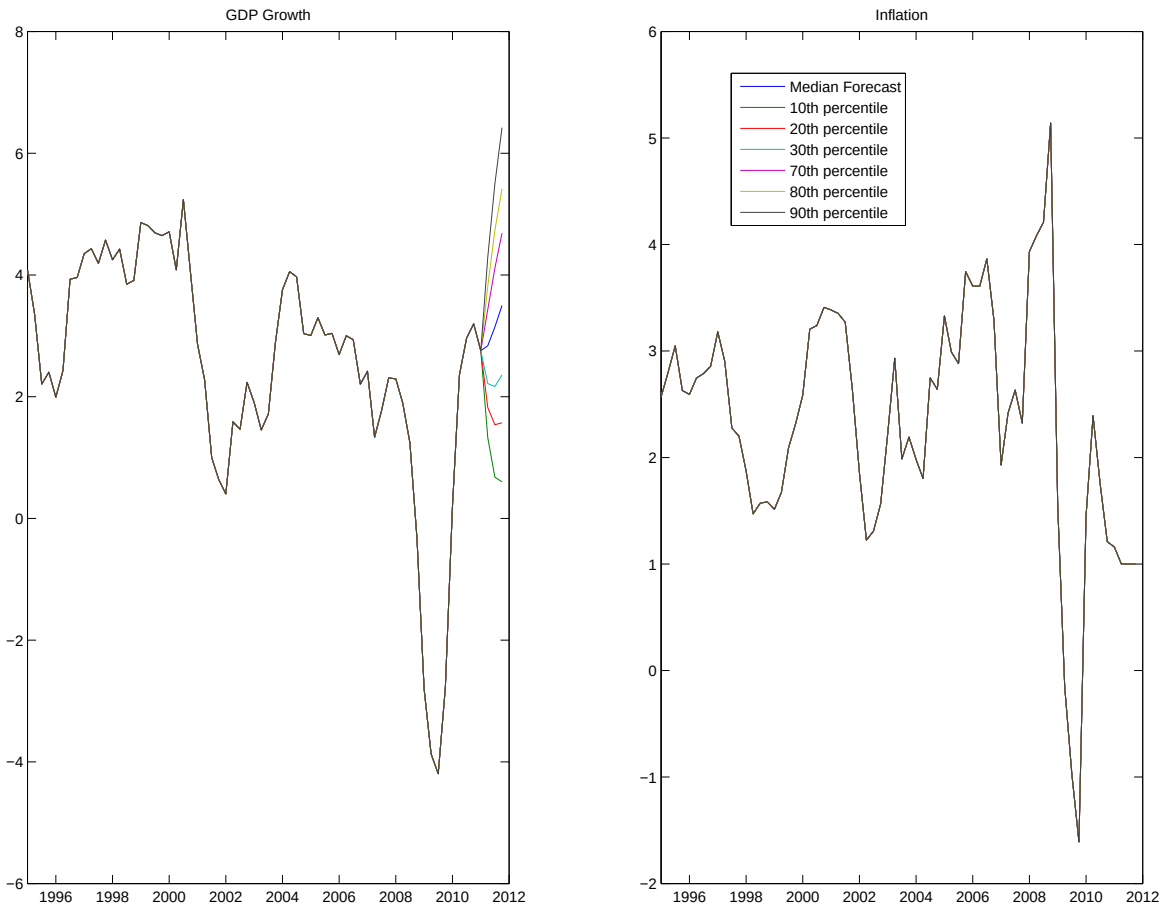

FIGURE 25. Conditional forecast for US GDP growth

Waggoner and Zha (1999) also discuss a simple method for imposing 'soft conditions' on forecasts– i.e. restricting the forecasts for some variables to lie within a range rather than the 'hard condition' we examine in the example above. Robertson *et al.* (2005) introduce an alternative method to impose 'soft conditions'.

8. Further Reading

- A comprehensive general treatment of Bayesian VARs can be found in Canova (2007) Chapter 10.
- An excellent intuitive explanation of priors and conditional forecasting can be found in Robertson and Tallman (1999).
- A heavily cited article discussing different prior distributions for VARs and methods for calculating posterior distributions is Kadiyala and Karlsson (1997).
- Banbura *et al.* (2007) is an illuminating example of implementing the natural conjugate prior via dummy observations.
- The appendix of Zellner (1971) provides an excellent description of the Inverse Wishart density.

9. Appendix: The marginal likelihood for a VAR model via Gibbs sampling

We can easily apply the method in Chib (1995) to calculate the marginal likelihood for a VAR model. This can then be used to select prior tightness (see for example Carriero *et al.* (2010)) or to choose the lag length and compare different models.

Consider the following VAR model

$$Y_t = c + \sum_{j=1}^{J} b_j Y_{t-j} + v_t, VAR(v_t) = \Sigma \tag{9.1}$$

The prior for the VAR coefficients $B = \{c, b_j\}$ is $P(B)\tilde{}N(\tilde{b}, H)$ and for the covariance matrix $P(\Sigma)\tilde{}IW(\bar{S}, \alpha)$. The posterior distribution of the model parameters $\Phi = B, \Sigma$ is defined via the Bayes rule

$$H\left(\Phi \backslash Y\right) = \frac{F\left(Y \backslash \Phi\right) \times P\left(\Phi\right)}{F\left(Y\right)} \tag{9.2}$$

where $\ln F\left(Y \backslash \Phi\right) = \frac{-TN}{2} \ln 2\pi + \frac{T}{2} \ln \left|\Sigma^{-1}\right| - 0.5 \sum_{i=1}^{T} \left(v_i \Sigma^{-1} v_i'\right)$ is the likelihood function with N representing the number of endogenous variables, $P\left(\Phi\right)$ is the joint prior distribution while $F\left(Y\right)$ is the marginal likelihood that we want to compute. Chib (1995) suggests computing the marginal likelihood by re-arranging equation 9.2. Note that in logs we can re-write equation 9.2 as

$$\ln F\left(Y\right) = \ln F\left(Y \backslash \Phi\right) + \ln P\left(\Phi\right) - \ln H\left(\Phi \backslash Y\right) \tag{9.3}$$

Note that equation 9.3 can be evaluated at any value of the parameters Φ to calculate $\ln F\left(Y\right)$. In practice a high density point Φ^* such as the posterior mean or posterior mode is used.

The likelihood function is easy to evaluate. In order to evaluate the priors, the pdf of the normal density and the inverse Wishart is needed. The latter is given in definition 3 above.

Following Chib (1995) the posterior density $H\left(\Phi^* \backslash Y\right) = H\left(B^*, \Sigma^* \backslash Y\right)$can be factored as

$$H\left(B^*, \Sigma^* \backslash Y\right) = H(B^* \backslash \Sigma^*, Y) \times H\left(\Sigma^* \backslash Y\right) \tag{9.4}$$

The first term on the RHS of equation 9.4 can be evaluated easily as this is simply the conditional posterior distribution of the VAR coefficients–i.e. a normal distribution with a known mean and covariance matrix.

$$\begin{aligned} & H(B^* \backslash \Sigma^*, Y)\tilde{}N(M, V) \\ M &= \left(H^{-1} + \Sigma^{*-1} \otimes X_t' X_t\right)^{-1} \left(H^{-1} \tilde{b}_0 + \Sigma^{*-1} \otimes X_t' X_t \hat{b}\right) \\ V &= \left(H^{-1} + \Sigma^{*-1} \otimes X_t' X_t\right)^{-1} \end{aligned}$$

The second term on the on the RHS of equation 9.4 can be evaluated by noting that

$$H\left(\Sigma^* \backslash Y\right) \approx \frac{1}{J} \sum_{j=1}^{J} H\left(\Sigma^* \backslash B_j, Y\right) \tag{9.5}$$

where B_j represent $j = 1...J$ draws of the VAR coefficients from the Gibbs sampler used to estimate the VAR model. Note that $H\left(\Sigma^* \backslash B_j, Y\right)$ is the inverse Wishart distribution with scale matrix $\bar{\Sigma} = \bar{S} + v_t' v_t$ and degrees of freedom $T + \alpha$ where the residuals v_t are calculated using the draws B_j.

Figures 26 and 27 show the matlab code for estimating the marginal likelihood in a simple BVAR with a natural conjugate prior implemented via dummy observations. On line 44 we calculate the marginal likelihood analytically for comparison with Chib's estimate. Analytical computation is possible with the natural conjugate prior (see Bauwens *et al.* (1999)), while Chib's estimator can be used more generally. Lines 46 to 67 estimate the VAR model using Gibbs sampling with the posterior means calculated on lines 69 and 70. Lines 73 to 76 calculate the prior moments which are used to evaluate the prior densities on lines 79 and 81. Line 83 evaluates the log likelihood function for the VAR model. Line 86 evaluates the term $H(B^* \backslash \Sigma^*, Y)$. Lines 88 to 99 evaluate the term $H\left(\Sigma^* \backslash Y\right) \approx \frac{1}{J} \sum_{j=1}^{J} H\left(\Sigma^* \backslash B_j, Y\right)$. These components are used to calculate the marginal likelihood on line 102.

```
clear
clc
addpath('functions');
REPS=15000;
BURN=10000;
[data,names]=xlsread('\data\usdata1.xls'); %load US data
N=cols(data);
L=2; %lag length of the VAR
Y=data;
%take lags
X=[];
for j=1:L
X=[X lag0(data,j) ];
end
X=[X ones(rows(X),1)];
Y=Y(L+1:end,:);
X=X(L+1:end,:);
T=rows(X);
%priors for VAR coefficients
lamdaP=1;  %This controls the tightness of the priors on the first
lag
tauP=0;  % this controls the tightness of the priors on sum of
coefficients 0 means not applied
epsilonP=1;  % this controls tightness of the prior on the constant
muP=mean(Y)';
sigmaP=[];
deltaP=[];
for i=1:N
    ytemp=Y(:,i);
    xtemp=[lag0(ytemp,1) ones(rows(ytemp),1)];
    ytemp=ytemp(2:end,:);
    xtemp=xtemp(2:end,:);
    btemp=xtemp\ytemp;
    etemp=ytemp-xtemp*btemp;
    stemp=etemp'*etemp/rows(ytemp);
    deltaP=[deltaP;btemp(1)];
    sigmaP=[sigmaP;stemp];
end
%dummy data to implement priors see
http://ideas.repec.org/p/ecb/ecbwps/20080966.html
[yd,xd] = create_dummies(lamdaP,tauP,deltaP,epsilonP,L,muP,sigmaP,N);
%yd and xd are the dummy data. Append this to actual data
  Y0=[Y;yd];
  X0=[X;xd];

 %compute the marginal likelihood analytically for comparison
 temp1=mlikvar1(Y,X,yd,xd);
  %conditional mean of the VAR coefficients
  mstar=vec(X0\Y0);  %ols on the appended data
  xx=X0'*X0;
  ixx=xx\eye(cols(xx));  %inv(X0'X0) to be used later in the Gibbs
sampling algorithm
  sigma=eye(N); %starting value for sigma
  out1=zeros(REPS-BURN,N*(N*L+1),1);
  out2=zeros(REPS-BURN,N,N);
  jj=1;
  for i=1:REPS
       vstar=kron(sigma,ixx);
       beta=mstar+(randn(1,N*(N*L+1))*chol(vstar))';

       %draw covariance
```

FIGURE 26. Marginal Likelihood for a VAR model

```
58        e=Y0-X0*reshape(beta,N*L+1,N);
59      scale=e'*e;
60      sigma=iwpq(T+rows(yd),inv(scale));
61
62      if i>=BURN
63          out1(jj,:,:)=beta;
64          out2(jj,:,:)=sigma;
65          jj=jj+1;
66      end
67    end
68
69    betam=squeeze(mean(out1,1));
70    sigmam=squeeze(mean(out2,1));
71
72    %evaluate priors
73    b0=vec(xd\yd);
74    b01=reshape(b0,N*L+1,N);
75    e0=yd-xd*b01;
76    S=e0'*e0;
77
78    %evaluate log prior distribution for VAR coefficients
79    bp=multivariatenormal(betam',b0,kron(S,pinv(xd'*xd)));
80    %evaluate log prior for VAR covariance
81    sp= invwishpdf(sigmam,S,size(yd,1)-size(xd,2));
82    %evaluate log likelihood
83    lik=loglik(reshape(betam,N*L+1,N),sigmam,Y,X);
84    %evaluate H(Bstar\sigmastar);
85    vstar1=kron(sigmam,ixx);
86    H1=multivariatenormal(betam',mstar,vstar1);
87    %evaluate H(sigmastar\beta[j])
88    H2i=[];
89    for j=1:size(out1,1)
90        betaj=out1(j,:);
91        e=Y0-X0*reshape(betaj,N*L+1,N);
92        scale=e'*e;
93        H2i= [H2i;invwishpdf(sigmam,scale,size(Y0,1))];
94    end
95    %take mean taking care of possible underflow/overflow with exp
96    factor=max(H2i);
97    H2=exp(H2i-factor);
98    H2m=mean(H2);
99    H2m=log(H2m)+factor;
100
101    %marginal lik
102    mlik=lik+bp+sp-H1-H2m;
103
104    disp('Analytical log Marginal Likelihood')
105    disp(temp1);
106
107    disp('Chib log Marginal Likelihood')
108    disp(mlik);
```

Published with MATLAB® 7.9

FIGURE 27. Marginal Likelihood for a VAR model (continued)

CHAPTER 3

Gibbs Sampling for state space models

1. Introduction

State space models have become a key tool for research and analysis in central banks. In particular, they can be used to detect structural changes in time series relationships and to extract unobserved components from data (such as the trend in a time series). The state space formulation is also used when calculating the likelihood function for DSGE models.

The classic approach to state space modelling can be computationally inefficient in large scale models as it is based on maximising the likelihood function with respect to all parameters. In contrast, Gibbs sampling proceeds by drawing from conditional distributions which implies dealing with smaller components of the model. In addition, Gibbs sampling provides an approximation to the marginal posterior distribution of the state variable and therefore directly provides a measure of uncertainty associated with the estimate of the state variable. The use of prior information also helps along the dimensions of the model where the data is less informative.

This chapter discusses the Gibbs sampling algorithm for state space models and provides examples of implementing the algorithm in Matlab.

2. Examples of state space models

In general, a state space model consists of the following two equations

$$Y_t = H\beta_t + Az_t + e_t \text{ Observation Equation} \tag{2.1}$$

$$\beta_t = \mu + F\beta_{t-1} + v_t \text{ Transition Equation} \tag{2.2}$$

Consider first the components of the observation equation 2.1. Here Y_t is observed data, H denotes either the right hand side variables or a coefficient matrix depending on the model as discussed below. β_t is the unobserved component or the state variable. z_t denotes exogenous variables with coefficient A. The observation equation, therefore, connects observed data to the unobserved state variable.

Consider the transition equation 2.2. This equation describes the dynamics of the state variable. Note that the order of the AR process in equation 2.2 is restricted to be 1. This condition is not restrictive in a practical sense as any AR(p) process can always be re-written in first order companion form. This is described in the examples below.

Finally, note that we make the following assumptions about the error terms e_t and v_t :

$$VAR(e_t) = R, VAR(v_t) = Q, COV(e_t, v_t) = 0 \tag{2.3}$$

As an example of a state space model consider a time-varying parameter regression: $Y_t = c_t + B_t X_t + e_t$, where the coefficients c_t and B_t are assumed to evolve as random walks. In state-space form this model can be expressed as:

$$Y_t = \overset{H}{\begin{pmatrix} 1 & X_t \end{pmatrix}} \overset{\beta_t}{\begin{pmatrix} c_t \\ B_t \end{pmatrix}} + e_t \text{ Observation Equation} \tag{2.4}$$

$$\overset{\beta_t}{\begin{pmatrix} c_t \\ B_t \end{pmatrix}} = \overset{\beta_{t-1}}{\begin{pmatrix} c_{t-1} \\ B_{t-1} \end{pmatrix}} + \overset{v_t}{\begin{pmatrix} v_{1t} \\ v_{2t} \end{pmatrix}} \text{ Transition Equation} \tag{2.5}$$

where $VAR(e_t) = R, VAR(v_t) = Q, COV(e_t, v_t) = 0$. Note that: (a) In this model $\mu = 0$ and $F = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ by assumption and (b) the matrix H in the observation equation represents the right hand side variables of the time-varying regression.

As a second example of a state space model, consider decomposing a series Y_t into two unobserved components, i.e. $Y_t = C_t + \tau_t$. We assume that: (1) the trend component τ_t follows random walk: $\tau_t = \tau_{t-1} + v_{2t}$ and (2) the cyclical component C_t follows an AR(2) process with a constant: i.e. $C_t = c + \rho_1 C_{t-1} + \rho_2 C_{t-2} + v_{1t}$. In state space form this model can be expressed as:

$$Y_t = \overset{H}{\begin{pmatrix} 1 & 1 & 0 \end{pmatrix}} \overset{\beta_t}{\begin{pmatrix} C_t \\ \tau_t \\ C_{t-1} \end{pmatrix}} \text{ Observation Equation} \tag{2.6}$$

$$\overset{\beta_t}{\begin{pmatrix} C_t \\ \tau_t \\ C_{t-1} \end{pmatrix}} = \overset{\mu}{\begin{pmatrix} c \\ 0 \\ 0 \end{pmatrix}} + \overset{F}{\begin{pmatrix} \rho_1 & 0 & \rho_2 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix}} \overset{\beta_{t-1}}{\begin{pmatrix} C_{t-1} \\ \tau_{t-1} \\ C_{t-2} \end{pmatrix}} + \overset{v_t}{\begin{pmatrix} v_{1t} \\ v_{2t} \\ 0 \end{pmatrix}} \text{ Transition Equation} \tag{2.7}$$

where $var \begin{pmatrix} v_{1t} \\ v_{2t} \\ 0 \end{pmatrix} = Q = \begin{pmatrix} Q_{1,1} & Q_{1,2} & 0 \\ Q_{1,2} & Q_{2,2} & 0 \\ 0 & 0 & 0 \end{pmatrix}$. Consider the observation equation for this model. Here the matrix H is a coefficient matrix which links the state variables $\begin{pmatrix} C_t \\ \tau_t \\ C_{t-1} \end{pmatrix}$ to Y_t. Note that the observation equation has no error term as we assume that Y_t decomposes exactly into the two components.

The left hand side of the transition equation has the state vector at time t i.e. β_t. The right hand side contains the state vector lagged one period i.e. $\beta_{t-1} = \begin{pmatrix} C_{t-1} \\ \tau_{t-1} \\ C_{t-2} \end{pmatrix}$. The fact that the state vector contains C_{t-1} implies that β_{t-1} contains C_{t-2}. This gives us a way to incorporate the AR(2) process for C_t into the transition equation. In general, if the state variable follows an AR(p) process, this implies adding $p-1$ lags of that state-variable into the state vector β_t.

The first row of the matrix F contains the AR coefficients for C_t with the constant in the corresponding row of μ. The second row forms the random walk process for τ_t. Note that the last row of F contains a 1 (element (1,1)) to link C_{t-1} on the left hand side and C_{t-1} on the right hand side and represents an identity. As a consequence, the last row of v_t equals zero with corresponding zeros in the Q matrix.

As a final example of a state space model, consider a dynamic factor model for a panel of series Y_{it} where $t = 1, 2..T$ represents time and $i = 1, 2...N$ represents the cross-section. Each series in the panel is assumed to depend on a common component F_t i.e. $Y_{it} = B_i F_{t-1} + e_{it}$.We assume that the common unobserved component F_t follows an $AR(2)$ process: $F_t = c + \rho_1 F_{t-1} + \rho_2 F_{t-2} + v_t$. This model has the following state-space representation:

$$\begin{pmatrix} Y_{1t} \\ Y_{2t} \\ . \\ Y_{Nt} \end{pmatrix} = \overset{H}{\begin{pmatrix} B_1 & 0 \\ B_1 & 0 \\ . & 0 \\ B_N & 0 \end{pmatrix}} \overset{\beta_t}{\begin{pmatrix} F_t \\ F_{t-1} \end{pmatrix}} + \overset{e_t}{\begin{pmatrix} e_{1t} \\ e_{2t} \\ . \\ e_{Nt} \end{pmatrix}} \text{ Observation Equation} \tag{2.8}$$

$$\overset{\beta_t}{\begin{pmatrix} F_t \\ F_{t-1} \end{pmatrix}} = \overset{\mu}{\begin{pmatrix} c \\ 0 \end{pmatrix}} + \overset{F}{\begin{pmatrix} \rho_1 & \rho_2 \\ 1 & 0 \end{pmatrix}} \overset{\beta_{t-1}}{\begin{pmatrix} F_{t-1} \\ F_{t-2} \end{pmatrix}} + \overset{v_t}{\begin{pmatrix} v_{1t} \\ 0 \end{pmatrix}} \text{ Transition Equation} \tag{2.9}$$

where $VAR(e_t) = R = \begin{pmatrix} R_{1,1} & 0 & 0 & 0 \\ 0 & R_{2,2} & 0 & 0 \\ 0 & 0 & . & 0 \\ 0 & 0 & 0 & R_{N,N} \end{pmatrix}$ and $VAR(v_t) = Q = \begin{pmatrix} Q_{1,1} & 0 \\ 0 & 0 \end{pmatrix}$. As in the unobserved component model, the matrix H contains the coefficients linking the data Y_{it} to the state variables β_t. The first lag of F_t appears in the state vector because of our assumption that F_t follows an AR(2) process. The transition equation of the system incorporates the AR(2) dynamics for the state variable in companion form with appropriate structures for the μ, F and Q matrices.

See Kim and Nelson (1999) Chapter 2 for further examples of state space models.

3. The Gibbs sampling algorithm for state space models

It is instructive to consider the unknown parameters of our state space system:

$$y_t = H\beta_t + Az_t + e_t,\ VAR(e_t) = R \tag{3.1}$$

$$\beta_t = \mu + F\beta_{t-1} + v_t, VAR(v_t) = Q \tag{3.2}$$

In the observation equation the unknown parameters consist of the elements of H that are not fixed or given as data (for e.g. the coefficients B_i in equation 2.8), the elements of A and the non-zero elements of the covariance matrix R. In the transition equation, the parameters to be estimated are the non-zero and free elements of μ, F and Q.In addition, the state variable β_t is unknown and needs to be estimated.

A Gibbs sampling algorithm for this problem can be discerned by considering the hypothetical case where the state variable β_t is known and observed. If this is the case, then the observation and the transition equations collapse to linear regressions with the conditional posterior distribution of coefficients and variances exactly as in Chapter 1. For example if the common factor F_t in equations 2.8 and 2.9 is known, these equations become a series of linear regressions. Equation 2.8 is then simply N linear regressions while equation 2.9 is an AR(2) model. The conditional

distributions of the parameters in this case are known from Chapter 1. This observation indicates the following general Gibbs algorithm for the state space model in equations 3.1 and 3.2.

Step 1 Conditional on β_t, sample H and R from their posterior distributions.
Step 2 Conditional on β_t, sample μ, F and Q from their posterior distributions.
Step 3 Conditional on the parameters of the state space: H, R, μ, F and Q sample the state variable β_t from its conditional posterior distribution.
Step 4 Repeat steps 1 to 3 until convergence is detected.

As emphasised above, steps 1 to 3 are standard and involve linear regressions and/or VARs with known conditional posteriors. *The new step required for the state space model is step 3 where we sample β_t from its conditional posterior distribution.* We turn to a description of the conditional posterior distribution for β_t next.

3.1. The conditional distribution of the state variable. We follow Kim and Nelson (1999) chapter 8 closely in this description. Let $\tilde{\beta}_T = [\beta_1, \beta_2, ...\beta_T]$ i.e. the time series of β from time $1, 2..T$. Similarly, let $\tilde{Y}_T = [Y_1, ...Y_T]$. Recall that we are interested in deriving the conditional posterior distribution $H\left(\tilde{\beta}_T \backslash H, Q, R, F, \mu, \tilde{Y}_T\right)$ i.e. the joint posterior for $\beta_1, \beta_2, ...\beta_T$. As shown by Carter and Kohn (1994), it is convenient to consider a factorisation of the joint density $H\left(\tilde{\beta}_T \backslash \tilde{Y}_T\right)$. Note, we drop the conditioning arguments for simplicity in what follows below.

We can factor $H\left(\tilde{\beta}_T \backslash \tilde{Y}_T\right)$ into the following conditional distributions

$$H\left(\tilde{\beta}_T \backslash \tilde{Y}_T\right) = H\left(\beta_T \backslash \tilde{Y}_T\right) \times H\left(\tilde{\beta}_{T-1} \backslash \beta_T, \tilde{Y}_T\right) \tag{3.3}$$

Note that the right hand side of 3.3 splits $H\left(\tilde{\beta}_T \backslash \tilde{Y}_T\right)$ into the product of the marginal distribution of the state variable at time T and the distribution of the vector $\tilde{\beta}_{T-1} = [\beta_1, \beta_2, ...\beta_{T-1}]$ conditioned on β_T. We can expand the term $H\left(\tilde{\beta}_{T-1} \backslash \beta_T, \tilde{Y}_T\right)$ as $H\left(\tilde{\beta}_{T-1} \backslash \beta_T, \tilde{Y}_T\right) = H\left(\beta_{T-1} \backslash, B_T, \tilde{Y}_T\right) \times H\left(\tilde{\beta}_{T-2} \backslash \beta_T, \beta_{T-1}, \tilde{Y}_T\right)$ where $\tilde{\beta}_{T-2} = [\beta_1, \beta_2, ...\beta_{T-2}]$. Thus:

$$H\left(\tilde{\beta}_T \backslash \tilde{Y}_T\right) = H\left(\beta_T \backslash \tilde{Y}_T\right) \times H\left(\beta_{T-1} \backslash B_T, \tilde{Y}_T\right) \times H\left(\tilde{\beta}_{T-2} \backslash \beta_T, \beta_{T-1}, \tilde{Y}_T\right) \tag{3.4}$$

Continuing in this vein and expanding $H\left(\tilde{\beta}_{T-2} \backslash \beta_T, \beta_{T-1}, \tilde{Y}_T\right) = H\left(\beta_{T-2} \backslash, \beta_T, \beta_{T-1}, \tilde{Y}_T\right) \times$
$H\left(\tilde{\beta}_{T-3} \backslash \beta_T, \beta_{T-1}, \beta_{T-2}, \tilde{Y}_T\right)$

$$H\left(\tilde{\beta}_T \backslash \tilde{Y}_T\right) = H\left(\beta_T \backslash \tilde{Y}_T\right) \times H\left(\beta_{T-1} \backslash B_T, \tilde{Y}_T\right) \times H\left(\beta_{T-2} \backslash \beta_T, \beta_{T-1}, \tilde{Y}_T\right) \times H\left(\tilde{\beta}_{T-3} \backslash \beta_T, \beta_{T-1}, \beta_{T-2}, \tilde{Y}_T\right)$$

Expanding further $\rightarrow$

$$\begin{aligned} H\left(\tilde{\beta}_T \backslash \tilde{Y}_T\right) &= H\left(\beta_T \backslash \tilde{Y}_T\right) \times H\left(\beta_{T-1} \backslash B_T, \tilde{Y}_T\right) \times H\left(\beta_{T-2} \backslash \beta_T, \beta_{T-1}, \tilde{Y}_T\right) \\ \times H\left(\beta_{T-3} \backslash \beta_T, \beta_{T-1}, \beta_{T-2}, \tilde{Y}_T\right) &\quad \times ....H\left(\beta_1 \backslash \beta_T, \beta_{T-1}, \beta_{T-2}, ...\beta_2, \tilde{Y}_T\right) \end{aligned} \tag{3.5}$$

As shown in Kim and Nelson (1999) (page 191) expression 3.5 can be simplified by considering the fact that β_T follows a first order AR or Markov process. Because of this Markov property, given $\tilde{Y}_T$ and β_{T-1}, in the term $H\left(\beta_{T-2} \backslash, \beta_T, \beta_{T-1}, \tilde{Y}_T\right)$, β_T contains no additional information for β_{T-2}. This term can therefore be re-written as $H\left(\beta_{T-2} \backslash \beta_{T-1}, \tilde{Y}_T\right)$. Similarly $H\left(\beta_{T-3} \backslash \beta_T, \beta_{T-1}, \beta_{T-2}, \tilde{Y}_T\right)$ can be re-written as $H\left(\beta_{T-3} \backslash \beta_{T-2}, \tilde{Y}_T\right)$.

A similar argument applies to the data vector $\tilde{Y}_T$. For example, in the term $H\left(\beta_{T-2} \backslash \beta_{T-1}, \tilde{Y}_T\right)$, $\tilde{Y}_{T-2} = [Y_1, ...Y_{T-2}]$ contains all the required information for β_{T-2} (given β_{T-1}). Therefore, this term can be re-written as $H\left(\beta_{T-2} \backslash \beta_{T-1}, \tilde{Y}_{T-2}\right)$. Similarly, the term $H\left(\beta_{T-3} \backslash \beta_{T-2}, \tilde{Y}_T\right)$ can be re-written as $H\left(\beta_{T-3} \backslash \beta_{T-2}, \tilde{Y}_{T-3}\right)$.

Given these simplifications we can re-write expression 3.5 as

$$\begin{aligned} H\left(\tilde{\beta}_T \backslash \tilde{Y}_T\right) &= H\left(\beta_T \backslash \tilde{Y}_T\right) \times H\left(\beta_{T-1} \backslash B_T, \tilde{Y}_{T-1}\right) \times H\left(\beta_{T-2} \backslash \beta_{T-1}, \tilde{Y}_{T-2}\right) \times H\left(\beta_{T-3} \backslash \beta_{T-2}, \tilde{Y}_{T-3}\right) \\ &\quad \times ....H\left(\beta_1 \backslash \beta_2, \tilde{Y}_1\right) \end{aligned} \tag{3.6}$$

or more compactly

$$H\left(\tilde{\beta}_T \backslash \tilde{Y}_T\right) = H\left(\beta_T \backslash \tilde{Y}_T\right) \prod_{t=1}^{T-1} H\left(\beta_t \backslash B_{t+1}, \tilde{Y}_t\right) \tag{3.7}$$

The conditional distribution of the state variable is given by expression 3.7.

Assuming that the disturbances of the state space model e_t and v_t are normally distributed:

$$\begin{aligned} H\left(\beta_T \backslash \tilde{Y}_T\right) &\tilde{} N(\beta_{T\backslash T}, P_{T\backslash T}) \\ H\left(\beta_t \backslash B_{t+1}, \tilde{Y}_t\right) &\tilde{} N(\beta_{t\backslash t,\beta_{t+1}}, P_{t\backslash t,\beta_{t+1}}) \end{aligned} \tag{3.8}$$

where the notation $\beta_{i\backslash j}$ denotes an estimate of β at time i given information upto time j. The two components on the right hand side of expression 3.7 are normal distributions. However, to draw from these distributions, we need to calculate their respective means and variances. To see this calculation we consider each component in turn.

3.1.1. *The mean and variance of* $H\left(\beta_T \backslash \tilde{Y}_T\right)$. We can compute the mean $\beta_{T\backslash T}$ and the variance $P_{T\backslash T}$ using the Kalman filter. The Kalman filter is a recursive algorithm which provides with an estimate of the state variable at each time period, given information up to that time period–i.e. it provides an estimate of $\beta_{t\backslash t}$ and its variance $P_{t\backslash t}$. To estimate the state variable, the Kalman filter requires knowledge of the parameters of the state space H, R, μ, F and Q. These are available in our Gibbs sampling framework from the previous draw of the Gibbs sampler.

The Kalman filter consists of the following equations which are evaluated recursively through time starting from an initial value $\beta_{0\backslash 0}$ and $P_{0\backslash 0}$.

$$\begin{aligned} \beta_{t\backslash t-1} &= \mu + F\beta_{t-1\backslash t-1} \\ P_{t\backslash t-1} &= FP_{t-1\backslash t-1}F' + Q \\ \eta_{t\backslash t-1} &= Y_t - H\beta_{t\backslash t-1} - Az_t \\ f_{t\backslash t-1} &= HP_{t\backslash t-1}H' + R \\ \beta_{t\backslash t} &= \beta_{t\backslash t-1} + K\eta_{t\backslash t-1} \\ P_{t\backslash t} &= P_{t\backslash t-1} - KHP_{t\backslash t-1} \end{aligned} \tag{3.9}$$

where $K = P_{t\backslash t-1}H'f_{t\backslash t-1}^{-1}$. Running these equations from $t = 1, 2...T$ delivers $\beta_{T\backslash T}$ and $P_{T\backslash T}$ at the end of the recursion.

Consider the intuition behind each equation of the Kalman filter. The first and the second equation are referred to as the prediction equations. The first equation $\beta_{t\backslash t-1} = \mu + F\beta_{t-1\backslash t-1}$ simply predicts the value of the state variable one period ahead using the transition equation of the model. This equation can be easily derived by taking the expected value of the transition equation i.e. $E\left(\mu + F\beta_{t-1} + v_t \backslash \bar{Y}_{t-1}\right) = \mu + F\beta_{t-1\backslash t-1}$ where $\bar{Y}_t = \{Y_t, z_t\}$. This follows by noting that $E(v_t \backslash \bar{Y}_{t-1}) = 0$ and $E\left(\beta_{t-1} \backslash \bar{Y}_{t-1}\right) = \beta_{t-1\backslash t-1}$. The second equation is simply the estimated variance of the state variable given information at time $t-1$ and can be derived by taking the variance of β_t (i.e. calculating $E\left[\beta_t - E\left(\beta_{t-1} \backslash \bar{Y}_{t-1}\right)\right]$). The prediction equations of the Kalman filter threfore produce an estimate of the state variable simply based on the parameters of the transition equation. Note that the observed data $\bar{Y}_t$ is not involved upto this point. The third equation of the Kalman filter calculates the prediction error $\eta_{t\backslash t-1} = Y_t - H\beta_{t\backslash t-1} - Az_t$. The fourth equation calculates the variance of the prediction error $f_{t\backslash t-1} = HP_{t\backslash t-1}H' + R$. This equation can be derived by calculating $E\left(Y_t - H\beta_{t\backslash t-1} - Az_t\right)^2$.

The final two equations of the Kalman filter are referred to as the updating equations. These equations update the initial estimates $\beta_{t\backslash t-1}$ and $P_{t\backslash t-1}$ using the information contained in the prediction error $\eta_{t\backslash t-1}$. Note that $K = P_{t\backslash t-1}H'f_{t\backslash t-1}^{-1}$ (referred to as the Kalman gain) can be thought of as the weight attached to prediction error. The updating equations can be derived by considering the formula for updating a linear projection.

As shown in Hamilton (1994) page 99 this formula is given as

$$\hat{P}\left(Y_3 \backslash Y_2, Y_1\right) = \hat{P}\left(Y_3 \backslash Y_1\right) + H_{32}H_{22}^{-1}\left[Y_2 - \hat{P}\left(Y_2 \backslash Y_1\right)\right] \tag{3.10}$$

In equation 3.10 we consider the hypothetical case where we have three variables Y_1, Y_2 and Y_3. Originally we were forecasting Y_3 based on Y_1 i.e. the term $\hat{P}\left(Y_3 \backslash Y_1\right)$ and we want to update this projection using the variable Y_2. According to equation 3.10 the updated projection is the sum of $\hat{P}\left(Y_3 \backslash Y_1\right)$ and the error in predicting Y_2 where the projection of Y_2 is based on Y_1. The weight attached to this prediction error is $H_{32}H_{22}^{-1}$ where H_{ij} is the covariance between Y_i, Y_j. Consider first the intuition behind the prediction error $Y_2 - \hat{P}\left(Y_2 \backslash Y_1\right)$. If the information contained in Y_1 and Y_2 is very similar, it is likely that $\hat{P}\left(Y_2 \backslash Y_1\right)$ and Y_2 will be similar and hence the extra unanticipated information contained in Y_2 will be limited. The weight attached to this extra information $H_{32}H_{22}^{-1}$ can be interpreted as the regression coefficient between Y_3 and Y_2. A larger value of this coefficient implies that the information contained in Y_2 receives a larger weight when updating the forecast $\hat{P}\left(Y_3 \backslash Y_1\right)$.

In our application, if we let $Y_3 = \beta_t$, $Y_2 = Y_t$ and $Y_1 = z_t, Y_{t-1} \rightarrow$

$$\begin{aligned} \beta_{t\backslash t} = \beta_{t\backslash t-1} &+ E\left[\left(\beta_t - \beta_{t\backslash t-1}\right)\left(Y_t - Y_{t\backslash t-1}\right)'\right] \times \\ &E\left[\left(Y_t - Y_{t\backslash t-1}\right)\left(Y_t - Y_{t\backslash t-1}\right)'\right]^{-1} \times \eta_{t\backslash t-1} \end{aligned} \tag{3.11}$$

where $Y_{t\backslash t-1} = H\beta_{t\backslash t-1} + Az_t$. Note that the term $E\left[\left(Y_t - Y_{t\backslash t-1}\right)\left(Y_t - Y_{t\backslash t-1}\right)'\right]$ is simply the forecast error variance $f_{t\backslash t-1}$. Also note that $Y_t - Y_{t\backslash t-1} = (H\beta_t + Az_t + e_t) - \left(H\beta_{t\backslash t-1} + Az_t\right) = H\left(\beta_t - \beta_{t\backslash t-1}\right) + e_t$. Thus

$$E\left[\left(\beta_t - \beta_{t\backslash t-1}\right)\left(Y_t - Y_{t\backslash t-1}\right)'\right] = E\left[\left(\beta_t - \beta_{t\backslash t-1}\right)\left(H\left(\beta_t - \beta_{t\backslash t-1}\right) + e_t\right)'\right]$$

$\rightarrow$

$$E\left[\left(\beta_t - \beta_{t\backslash t-1}\right)\left(H\left(\beta_t - \beta_{t\backslash t-1}\right)\right)'\right] = P_{t\backslash t-1}H'$$

Substituting these in equation 3.11 produces the updating equation $\beta_{t\backslash t} = \beta_{t\backslash t-1} + K\eta_{t\backslash t-1}$. A similar derivation can be used to obtain the final updating equation $P_{t\backslash t}$ as shown in Hamilton (1994) page 380. Finally note that the likelihood function is given as a by product of the Kalman filter recursions as $-\frac{1}{2}\sum_{t=1}^{T}\ln 2\pi^n\left|f_{t\backslash t-1}\right| - \frac{1}{2}\sum_{t=1}^{T}\eta'_{t\backslash t-1}f^{-1}_{t\backslash t-1}\eta_{t\backslash t-1}$.

For a stationary transition equation, the initial values for the Kalman filter recursions $\beta_{0\backslash 0}$ and $P_{0\backslash 0}$ are given as the unconditional mean and variance. That is $\beta_{0\backslash 0} = (I_k - F)^{-1}\mu$ and $vec(P_{0\backslash 0}) = (I - F\otimes F)^{-1}vec(Q)$. If the transition equation of the system is non-stationary (for e.g. if the state variable follows a random walk) the unconditional moments do not exist. In this case $\beta_{0\backslash 0}$ can be set arbitrarily. $P_{0\backslash 0}$ is then set as a diagonal matrix with large diagonal entries reflecting uncertainty around this initial guess.

To recap, we evaluate the equations of the Kalman filter given in 3.9 for periods $t = 1...T$. The final recursion delivers $\beta_{T\backslash T}$ and $P_{T\backslash T}$ the mean and variance of $H\left(\beta_T\backslash\tilde{Y}_T\right)$.

3.1.2. *The mean and variance of* $H\left(\beta_t\backslash\beta_{t+1},\tilde{Y}_t\right)$. The mean and variance of the conditional distribution $H\left(\beta_t\backslash\beta_{t+1},\tilde{Y}_t\right)$ can also be derived using the Kalman filter updating equations. As discussed in Kim and Nelson (1999) page 192, deriving the mean $\beta_{t\backslash t,\beta_{t+1}}$ can be thought of as updating $\beta_{t\backslash t}$ (the kalman filter estimate of the state variable) for information contained in β_{t+1}which we treat as observed (for e.g. at time $T-1$, β_{t+1} is given using a draw from $H\left(\beta_T\backslash\tilde{Y}_T\right)$ which we discussed above) Note that this task fits into the framework of the updating equations discussed in the previous section as we are updating an estimate using new information. In other words, the updating equations of the Kalman filter apply with parameters and the prediction error chosen to match our problem.

For the purpose of this derivation we can consider a state space system with the observation equation:

$$\beta_{t+1} = \mu + F\beta_t + v_{t+1} \tag{3.12}$$

This implies that the prediction error is given by $\eta^*_{t+1\backslash t} = \beta_{t+1} - \mu + F\beta_{t\backslash t}$. The forecast error variance is given by $f^*_{t+1\backslash t} = FP_{t\backslash t}F' + Q$. Note also that for this observation equation, the matrix that relates the state variable β_t to the observed data β_{t+1} is $H^* = F$. With these definitions in hand we can simply use the updating equations of the Kalman filter. That is

$$\beta_{t\backslash t,\beta_{t+1}} = \beta_{t\backslash t} + K^*\left(\beta_{t+1} - \mu + F\beta_{t\backslash t}\right) \tag{3.13}$$

$$P_{t\backslash t,\beta_{t+1}} = P_{t\backslash t} - K^*H^*P_{t\backslash t} \tag{3.14}$$

where the gain matrix is $K^* = P_{t\backslash t}H^{*\prime}f^{*-1}_{t+1\backslash t}$.

Equations 3.13 and 3.14 are evaluated backwards in time starting from period $T-1$ and iterating backwards to period 1. This recursion consists of the following steps:

Step 1 Run the Kalman filter from $t = 1...T$ to obtain the mean $\beta_{T\backslash T}$ and the variance $P_{T\backslash T}$ of the distribution $H\left(\beta_T\backslash\tilde{Y}_T\right)$. Also save $\beta_{t\backslash t}$ and $P_{t\backslash t}$ for $t = 1...T$. Draw β_T from the normal distribution with mean $\beta_{T\backslash T}$ and the variance $P_{T\backslash T}$. Denote this draw by $\hat{\beta}_T$

Step 2 At time $T-1$, use 3.13 to calculate $\beta_{T-1\backslash T-1,\beta_T} = \beta_{T-1\backslash T-1} + K^*\left(\hat{\beta}_T - \mu + F\beta_{T-1\backslash T-1}\right)$ where $\beta_{T-1\backslash T-1}$ is the Kalman filter estimate of the state variable (from step 1) at time $T-1$. Use equation 3.14 to calculate $P_{t\backslash t,\beta_{t+1}}$. Draw $\hat{\beta}_{T-1}$ from the normal distribution with mean $\beta_{T-1\backslash T-1,\beta_T}$ and variance $P_{t\backslash t,\beta_{t+1}}$.

Step 3 Repeat step 2 for $t = T-2, T-3, ...1$.

This backward recursion (The Carter and Kohn algorithm) delivers a draw of $\tilde{\beta}_T = [\beta_1, \beta_2, ...\beta_T]$ from its conditional posterior distribution.

A minor modification to this algorithm is required if the matrix Q is singular (see the example of the state space model given in equation 2.6). In this case we evaluate equations 3.13 and 3.14 using $\bar{F}$ instead of F, $\bar{Q}$ instead of Q and $\bar{\mu}$ instead of μ where $\bar{F}, \bar{Q}, \bar{\mu}$ correspond to the non-singular block of Q. In the example given in equation 2.6 above $\bar{\mu} = \begin{pmatrix} c \\ 0 \end{pmatrix}, \bar{F} = \begin{pmatrix} \rho_1 & 0 & \rho_2 \\ 0 & 1 & 0 \end{pmatrix}$ and $\bar{Q} = \begin{pmatrix} Q_{1,1} & Q_{1,2} \\ Q_{1,2} & Q_{2,2} \end{pmatrix}$.

3.2. The Gibbs sampling algorithm. We can now re-state the Gibbs alogrithm for the state space model in equations 3.1 and 3.2.

Step 1 Conditional on β_t, sample H and R from their posterior distributions.
Step 2 Conditional on β_t, sample μ, F and Q from their posterior distributions.
Step 3 Conditional on the parameters of the state space: H, R, μ, F and Q sample the state variable β_t from its conditional posterior distribution. That is, run the Kalman filter to obtain $\beta_{t\backslash t}$ and $P_{t\backslash t}$ for $t = 1...T$ and draw β_T. Use equations 3.13 and 3.14 to draw $\beta_1, \beta_2, ...\beta_{T-1}$.
Step 4 Repeat steps 1 to 3 until convergence is detected.

Implementing this Gibbs sampling algorithm therefore requires programming the Kalman filter and equations 3.13 and 3.14 in matlab. The remainder of this chapter describes this task with the help of several examples.

4. The Kalman filter in Matlab

To discuss the implementation of the Kalman filter in Matlab we will consider the following time varying parameter model as an example

$$\begin{aligned} Y_t &= X_t\beta_t + v_t \\ \beta_t &= \mu + F\beta_{t-1} + e_t \\ VAR(v_t) &= R \\ VAR(e_t) &= Q \end{aligned} \tag{4.1}$$

where Y_t is a $T \times 1$ matrix containing the dependent variable, X_t is a $T \times 1$ matrix containg the regressor with time-varying coefficient β_t. For the moment we assume that the parameters of this state space model μ, F, R and Q are known and we are interested in estimating the time-varying coefficient β_t the state variable.

Figures 1 and 2 show the matlab code for the Kalman filter equations (Example1.m). Lines 7 to 20 of the file generate artificial data for Y_t (see equation 4.1) assuming that $\mu = 0, F = 1, Q = 0.001, R = 0.01$. Line 21 starts the Kalman filter by setting up the initial conditions for the state variable β_t. Line 22 assumes that $\beta_{0\backslash 0} = 0$ and line 23 sets $p_{0\backslash 0}$ the variance of the initial state equal to 1. The Kalman filter starts with $\beta_{t-1\backslash t-1} = \beta_{0\backslash 0}$ and $P_{t-1\backslash t-1} = P_{0\backslash 0}$ (lines 27 and 28) and then iterates through the sample (loop starts on line 29). Line 32 is the first equation of the prediction step of the Kalman filter $\beta_{t|t-1} = \mu + F\beta_{t-1|t-1}$. Line 33 calculates the variance of $\beta_{t|t-1}$ using the equation $P_{t\backslash t-1} = FP_{t-1\backslash t-1}F' + Q$. Line 34 calculates the fitted value of Y_t for that time period as $X_t\beta_{t|t-1}$ and the next line calculates the prediction error for that time period $\eta_{t\backslash t-1} = Y_t - X_t\beta_{t|t-1}$. Line 36 calculates the variance of the prediction error $f_{t\backslash t-1} = X_tP_{t\backslash t-1}X_t' + R$. Line 38 starts the updating step of the Kalman filter by calculating the Kalman gain $K_t = P_{t\backslash t-1}X_t'f_{t\backslash t-1}^{-1}$. Line 39 updates the the estimate of the state variable based on information contained in the prediction error $\beta_{t\backslash t} = \beta_{t|t-1} + K_t\eta_{t\backslash t-1}$ where this information is weighted by the Kalman gain. The final equation of the Kalman filter (line 40) updates the variance of the state variable $P_{t\backslash t} = P_{t\backslash t-1} - K_tX_tp_{t\backslash t-1}$

Figure 3 shows the estimates of β_t obtained using the Kalman filter. These closely match the assumed true value of β_t.

5. The Carter and Kohn algorithm in Matlab

Recall that that the conditional distribution of the state variable $\tilde{\beta}_T = [\beta_1, \beta_2, ...\beta_T]$ is

$$H\left(\tilde{\beta}_T\right) = H\left(\beta_T, \tilde{Y}_T\right)\prod_{t=1}^{T-1} H\left(\beta_t\backslash\beta_{t+1}, \tilde{Y}_t\right) \tag{5.1}$$

As discussed above, this implies that

$$\begin{aligned} &\beta_T \tilde{} N(\beta_{T\backslash T}, p_{T\backslash T}) \\ &\beta_t\backslash\beta_{t+1} \tilde{} N(\beta_{t\backslash t,\beta_{t+1}}, P_{t\backslash t,\beta_{t+1}}) \end{aligned} \tag{5.2}$$

As described above, the mean and variance in $\beta_T \tilde{} N(\beta_{T\backslash T}, P_{T\backslash T})$ is delivered by the Kalman filter at time $t = T$. The computation of the mean and variance in $N(\beta_{t\backslash t,\beta_{t+1}}, P_{t\backslash t,\beta_{t+1}})$ requires the updating equations 3.13 and 3.14. Written in full these are:

$$\beta_{t\backslash t,\beta_{t+1}} = \beta_{t\backslash t} + P_{t\backslash t}F'\left(Fp_{t\backslash t}F' + Q\right)^{-1}\left(\beta_{t+1} - \mu - F\beta_{t\backslash t}\right) \tag{5.3}$$

$$P_{t\backslash t,\beta_{t+1}} = P_{t\backslash t} - P_{t\backslash t}F'\left(FP_{t\backslash t}F' + Q\right)^{-1}FP_{t\backslash t} \tag{5.4}$$

These are computed going backwards in time from period $t-1$ to 1. We now turn to the implementation of the algorithm in Matlab

Figures 4 and 5 show the matlab code for the Carter and Kohn algorithm for artificial data on the state space model shown in equation 4.1) assuming that $\mu = 0, F = 1, Q = 0.001, R = 0.01$ (See example2.m). As alluded to above, the algorithm works in two steps. As a first step we run the Kalman filter to obtain $\beta_{T\backslash T}, p_{T\backslash T}$. Lines 21 to

```
1 clear
2 %generate data for a state space model
3 %Y=Beta[t]*X+e1
4 %Beta[t]=mu+F*Beta[t-1]+e2
5 %var(e1)=R
6 %var(e2)=Q
```

$$Y_t = X_t\beta_t + v_t$$

$$\beta_t = \mu + F\beta_{t-1} + e_t$$

$$VAR(v_t) = R$$

$$VAR(e_t) = Q$$

```
7 t=500;
8 Q=0.001;
9 R=0.01;
10 F=1;  %these are fixed
11 mu=0;  %these are fixed
12 e1=randn(t,1)*sqrt(R);
13 e2=randn(t,1)*sqrt(Q);
14 Beta=zeros(t,1);
15 Y=zeros(t,1);
16 X=randn(t,1);
17 for j=2:t
18     Beta(j,:)=Beta(j-1,:)+e2(j,:);
19     Y(j)=X(j,:)*Beta(j,:)'+e1(j);
20 end
Start of the Kalman filter
21 %%Step 1 Set up matrices for the Kalman Filter
22 beta0=zeros(1,1);   %state variable  b[0/0]   β_{0\0}
23 p00=1;          %variance of state variable p[0/0]   p_{0\0}
24 beta tt=[];          %will hold the filtered state variable
25 ptt=zeros(t,1,1);    % will hold its variance
26 %initialise the state variable
27 beta11=beta0;   β_{t-1|t-1}
28 p11=p00;   p_{t-1\t-1}
29 for i=1:t Loop from period 1 to end of sample
30     x=X(i);
31     %Prediction
32 beta10=mu+beta11*F';   β_{t|t-1} = μ + Fβ_{t-1|t-1}
33 p10=F*p11*F'+Q;   p_{t\t-1} = Fp_{t-1\t-1}F' + Q
34 yhat=(x*(beta10)')';   X_tβ_{t|t-1}
35 eta=Y(i,:)-yhat;   η_{t\t-1} = Y_t − X_tβ_{t|t-1}
36 feta=(x*p10*x')+R;   f_{t\t-1} = X_t p_{t\t-1} X_t' + R
37 %updating
38 K=(p10*x')*inv(feta);   K_t = p_{t\t-1}X_t' f_{t\t-1}^{-1}   Kalman gain
39 beta11=(beta10'+K*eta')';   β_{t\t} = β_{t|t-1} + K_t η_{t\t-1}
```

FIGURE 1. The Kalman filter in Matlab

44 of the code are the Kalman filter equations and are identical to the example above. Note that the matrix ptt saves $p_{t\backslash t}$ for each time period.[1]The matrix beta_tt saves $\beta_{t\backslash t}$ for each time period. Line 47 specifies an empty matrix to hold the draw of β_t. Line 48 specifies a $T \times 1$ vector from the $N(0,1)$ distribution to be used below. Line 51 draws from $H\left(\beta_T, \tilde{Y}_T\right)$ where the mean of this distribution is $\beta_{T\backslash T}$ and the variance is $P_{T\backslash T}$ where both these quantities are delivered by the kalman filter and saved as the last row of beta_tt and ptt respectively. Line 53 starts the second step of the Carter and Kohn algorithm and begins a loop going backwards from period T-1 to1. Line 55 computes the

[1]This is set up as a three dimensional matrix where the first dimension is time and the second two dimensions are the rows and columns of the covariance matrix $p_{t\backslash t}$. In this example this matrix has dimension $500 \times 1 \times 1$.

$p_{t\backslash t} = p_{t\backslash t-1} - K_t X_t p_{t\backslash t-1}$

```
40 p11=p10-K*(x*p10);
41 ptt(i,:,:)=p11;
42 beta tt=[beta tt;beta11];
43 end
44 %%%%%%%%%%%end of Kalman
Filter%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
45 plot([beta tt Beta])
46 axis tight
47 legend('estimated \beta_{t}','true \beta_{t}');
```

Published with MATLAB® 7.9

FIGURE 2. The Kalman filter in Matlab continued

mean of $H\left(\beta_t \backslash \beta_{t+1}, \tilde{Y}_t\right)$ using the expression $\beta_{t\backslash t,\beta_{t+1}} = \beta_{t\backslash t} + P_{t\backslash t}F'\left(FP_{t\backslash t}F' + Q\right)^{-1}\left(\beta_{t+1} - \mu - F\beta_{t\backslash t}\right)$. Note that the term β_{t+1} is the draw of β_t one period in the future. Line 56 calculates the variance of $H\left(\beta_t \backslash \beta_{t+1}, \tilde{Y}_t\right)$ using the expression $P_{t\backslash t,\beta_{t+1}} = P_{t\backslash t} - P_{t\backslash t}F'\left(Fp_{t\backslash t}F' + Q\right)^{-1} FP_{t\backslash t}$.Line 57 draws the state variable from a normal distribution using this mean and variance.

Figure 6 plots the result of running this code and shows the draw for β_t.

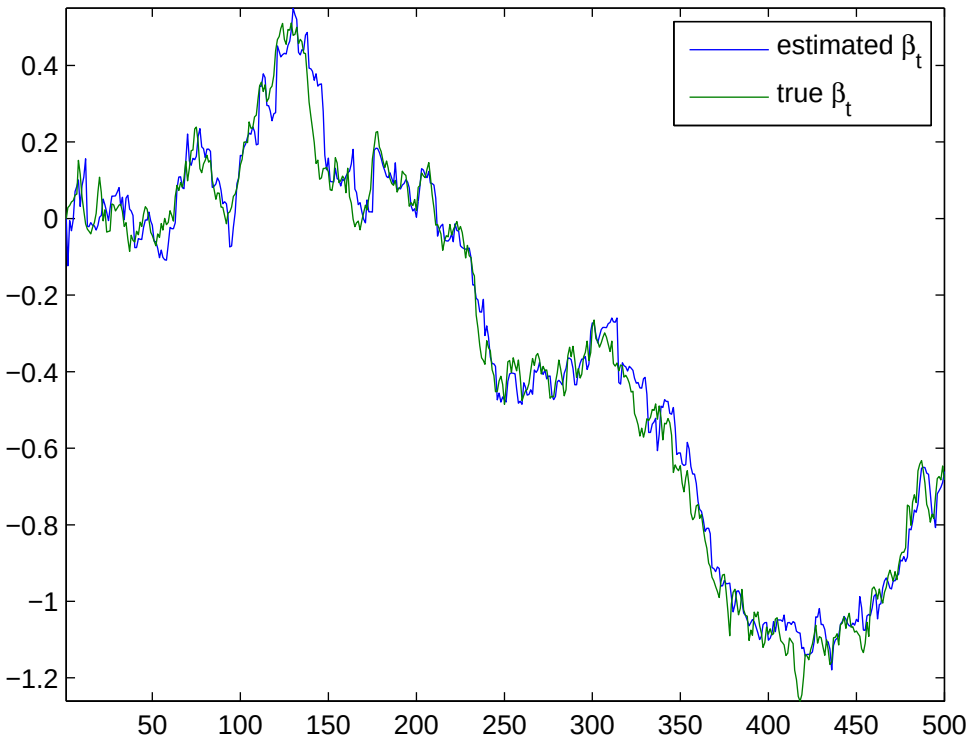

FIGURE 3. Estimates of β_t from the Kalman filter

6. The Gibbs sampling algorithm for a VAR with time-varying parameters

We now consider our first example that illustrates the Carter and Kohn algorithm. Following Cogley and Sargent (2002), we consider the following VAR model with time-varying coefficients

$$\begin{aligned} Y_t &= c_t + \sum_{j=1}^{P} B_{j,t} Y_{t-j} + v_t, VAR(v_t) = R \\ \beta_t &= \{c_t, B_{1,t} .... B_{P,t}\} \\ \beta_t &= \mu + F\beta_{t-1} + e_t, VAR(e_t) = Q \end{aligned} \tag{6.1}$$

Note that most empirical applications of this model assume that $\mu = 0$ and $F = 1$ and we are going to implement this assumption in our code below. The Gibbs sampling algorithm for this model can be discerned by noticing that if the time-varying coefficients β_t are known, the conditional posterior distribution of R is inverse Wishart. Similarly, conditional on β_t the distribution of Q is inverse Wishart. Conditional on R and Q and with $\mu = 0$ and $F = 1$ the model in 6.1 is a linear Gaussian state space model. The conditional posterior of β_t is normal and the mean and the variance can be derived via the Carter Kohn algorithm. Therefore the Gibbs sampling algorithm consists of the following steps

Step 1 Set a prior for R and Q and starting values for the Kalman filter. The prior for Q is inverse Wishart $p(Q) \sim IW(Q_0, T_0)$. Note that this prior is quite crucial as it influences the amount of time-variation allowed for in the VAR model. In other words, a large value for the scale matrix Q_0 would imply more fluctuation in β_t. This prior is typically set using a training sample. The first T_0 observations of the sample are used to estimate a standard fixed coefficient VAR via OLS such that $\beta_0 = (X_{0t}' X_{0t})^{-1} (X_{0t}' Y_{0t})$ with a coefficient covariance matrix given by $p_{0\backslash 0} = \Sigma_0 \otimes (X_{0t}' X_{0t})^{-1}$ where $X_{0t} = \{Y_{0t-1}, ...Y_{0t-p}, 1\}$, $\Sigma_0 = \frac{(Y_{0t} - X_{0t}\beta_0)'(Y_{0t} - X_{0t}\beta_0)}{T_0 - K}$ and the subscript 0 denotes the fact that this is the training sample. The scale matrix Q_0 is set equal to $p_{0\backslash 0} \times T_0 \times \tau$ where τ is a scaling factor chosen by the researcher. Some studies set $\tau = 3.510^{-4}$ i.e. a small number to reflect the fact that the training sample in typically short and the resulting estimates of $p_{0\backslash 0}$ maybe imprecise. Note that one can control the apriori amount of time-variation in the model by varying τ. The prior degrees of freedom are set equal to T_0. The prior for R is inverse Wishart with scale parameter R_0 and degrees of freedom v_{R_0}. The initial state is set equal to $\beta_{0\backslash 0} = vec(\beta_0)'$ and the intial state covariance is given by $P_{0\backslash 0}$. We set a starting value for R and Q.

```
clear
%generate data for a state space model
%Y=Beta[t]*X+e1
%Beta[t]=mu+F*Beta[t-1]+e2
%var(e1)=R
%var(e2)=Q
t=500;
Q=0.001;
R=0.01;
F=1;  %these are fixed
mu=0;  %these are fixed
e1=randn(t,1)*sqrt(R);
e2=randn(t,1)*sqrt(Q);
Beta=zeros(t,1);
Y=zeros(t,1);
X=randn(t,1);
for j=2:t
    Beta(j,:)=Beta(j-1,:)+e2(j,:);
    Y(j)=X(j,:)*Beta(j,:)'+e1(j);
end
%%Step 1 Set up matrices for the Kalman Filter
beta0=zeros(1,1);   %state variable  b[0/0]
p00=1;           %variance of state variable p[0/0]
beta_tt=[];          %will hold the filtered state variable
ptt=zeros(t,1,1);    % will hold its variance
%initialise the state variable
beta11=beta0;
p11=p00;
for i=1:t
    x=X(i);
    %Prediction
beta10=mu+beta11*F';
p10=F*p11*F'+Q;
yhat=(x*(beta10)')';
eta=Y(i,:)-yhat;
feta=(x*p10*x')+R;
%updating
K=(p10*x')*inv(feta);
beta11=(beta10'+K*eta')';
p11=p10-K*(x*p10);
ptt(i,:,:)=p11;
beta_tt=[beta_tt;beta11];
end
%%%%%%%%%%%end of Kalman Filter%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% Carter and Kohn Backward recursion to calculate the mean and variance of the distribution of the state
%vector
beta2 = zeros(t,1);   %this will hold the draw of the state variable
wa=randn(t,1);
i=t;  %period T
p00=squeeze(ptt(i,:,:));     p_T\T
β_T\T ──→
beta2(i,:)=beta_tt(i:i,:)+(wa(i:i,:)*chol(p00));     β_T~N(β_T\T,p_T\T)
%periods T-1..to 1
for i=t-1:-1:1
pt=squeeze(ptt(i,:,:));
```

Figure 4. The Carter and Kohn algorithm in Matlab

Step 2 Sample $\tilde{\beta}_t$ conditional on R and Q from its conditional posterior distribution $H\left(\tilde{\beta}_T \backslash R, Q, \tilde{Y}_t\right)$ where $\tilde{\beta}_T = [vec(\beta_1)', vec(\beta_2)', ...vec(\beta_T)']$ and $\tilde{Y}_t = [Y_1, ...Y_T]$. This is done via the Carter and Kohn algorithm as described in the example above. We describe the Matlab implementation in the next section.

Step 3 Sample Q from its conditional posterior distribution. Conditional on $\tilde{\beta}_t$ the posterior of Q is inverse Wishart with scale matrix $\left(\tilde{\beta}_t^1 - \tilde{\beta}_{t-1}^1\right)' \left(\tilde{\beta}_t^1 - \tilde{\beta}_{t-1}^1\right) + Q_0$ and degrees of freedom $T + T_0$ where T denotes the length of the estimation sample and $\tilde{\beta}_t^1$ is the previous draw of the state variable $\tilde{\beta}_t$. Notice that once the state variable is drawn from its distribution we treat it like data. It is therefore easy to extend this step to sample

```
55 bm=beta tt(i:i,:)+(pt*F'*inv(F*pt*F'+Q)*(beta2(i+1:i+1,:)-mu-
beta tt(i,:)*F')')';
56 pm=pt-pt*F'*inv(F*pt*F'+Q)*F*pt;
57 beta2(i:i,:)=bm+(wa(i:i,:)*chol(pm));
58 end
59 plot([beta_tt beta2 Beta])
60 axis tight
61 legend('Kalman filter estimated \beta_{t}','Draw from
H(\beta_{t})','true \beta_{t}');
```

$\beta_{t|t,B_{t+1}} = \beta_{t|t} + p_{t|t}F'(Fp_{t|t}F' + Q)^{-1}(\beta_{t+1} - \mu - F\beta_{t|t})$

$p_{t|t,B_{t+1}} = p_{t|t} - p_{t|t}F'(Fp_{t|t}F' + Q)^{-1}Fp_{t|t}$

$\beta_t \backslash \beta_{t+1} \sim N(\beta_{t|t,B_{t+1}}, p_{t|t,B_{t+1}})$

Published with MATLAB® 7.9

FIGURE 5. The Carter and Kohn algorithm in Matlab (continued)

μ, F which are just the intercept and AR coefficients in an AR regression for each individual coefficient in $\tilde{\beta}_t$ (the conditional distributions for linear regression models are described in chapter 1)

Step 4. Sample R from its conditional posterior distribution. Conditional on the draw $\tilde{\beta}_t^1$ the posterior of R is inverse Wishart with scale matrix $\left(Y_t - \left(c_t^1 + \sum_{j=1}^{P} B_{j,t}^1 Y_{t-j}\right)\right)' \left(Y_t - \left(c_t^1 + \sum_{j=1}^{P} B_{j,t}^1 Y_{t-j}\right)\right) + R_0$ and degrees of freedom $T + v_{R_0}$.

Step 5. Repeat steps 2 to 4 M times and use the last L draws for inference. Note that unlike fixed coefficient VAR models, this state space model requires a large number of draws for convergence (e.g. $M \geq 100,000$).

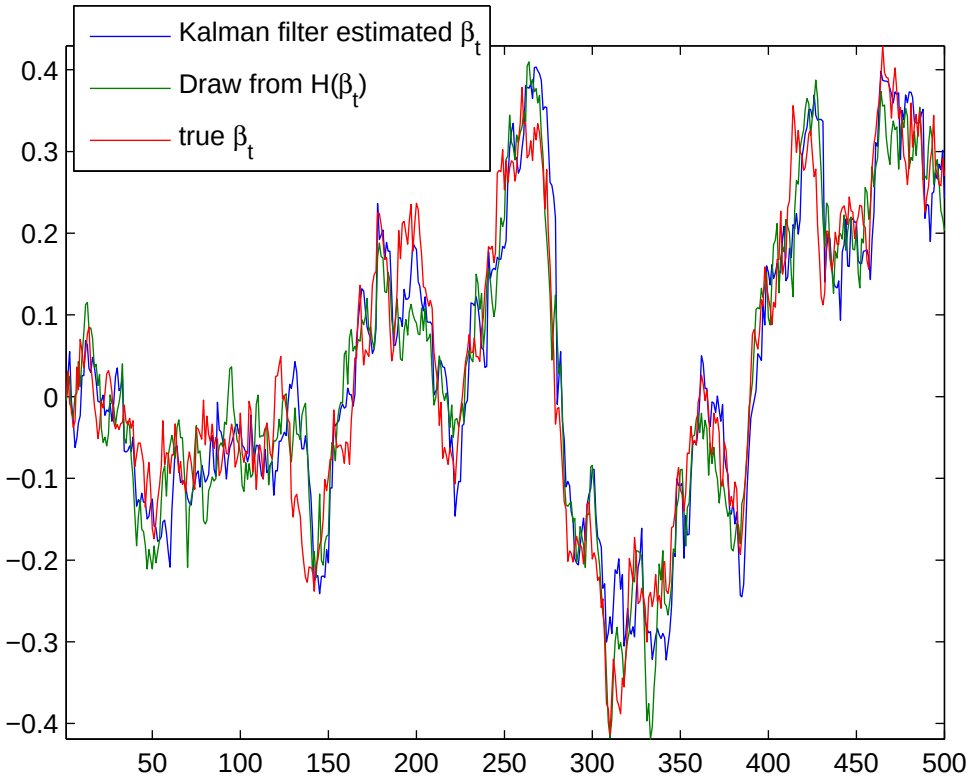

FIGURE 6. A draw from the conditional posterior distribution of β_t using the Carter and Kohn algorithm

6.1. Matlab code for the time-varying parameter VAR. We consider a time-varying VAR model with two lags using US data on GDP growth, CPI inflation and the Federal Funds rate over the period 1954Q3 to 2010Q2 (Example3.m). We use the time-varying VAR model to compute the impulse response to a monetary policy shock at each point in time and see if the response of the US economy to this shock has changed over this period. The code for this model can be seen in figures 7, 8, 9 and 10. Line 13 of the code sets the training sample equal to the first 40 observations of the sample and line 16 calculates a fixed coefficient VAR using this training sample to obtain β_0 and $p_{0\backslash 0}$. In calculating Q_0 on line 21 we set $\tau = 3.510^{-4}$. Lines 25 and 26 set a starting value for R and Q. Lines 29 and 30 remove the training sample from the data–the model is estimated on the remaining sample. Lines 38 to 88 sample the time-varying coefficients conditional on R and Q using the Carter and Kohn algorithm. The code for this is exactly the same as in the previous example with some monor differences. First, note that the VAR is expressed as $Y_t = (I_N \otimes X_t)\, vec(\beta_t) + v_t$ for each time period $t = 1....T$. This is convenient as it allows us to write the transition equation in terms of $vec(\beta_t)$ i.e. the VAR coefficients in vectorised form at each point in time. Therefore, on line 47 x is set equal to $(I_N \otimes X_t)$. The second practical differences arises in the backward recursion on lines 64 to 87. In particular (following earlier papers) we draw $\tilde{\beta}_T = [vec(\beta_1)', vec(\beta_2)', ...vec(\beta_T)']$ for $H\left(\tilde{\beta}_T \backslash R, Q, \tilde{Y}_t\right)$ but ensure that the VAR is stable at each point in time. If the stability condition fails for one time period, the entire matrix $\tilde{\beta}_T = [vec(\beta_1)', vec(\beta_2)', ...vec(\beta_T)']$ is discarded and the algorithm tries again. With the draw of $\tilde{\beta}_T$ in hand line 89 calculates the residuals of the transition equation e_t. Line 90 calculates the scale matrix $\left(\tilde{\beta}_t^1 - \tilde{\beta}_{t-1}^1\right)' \left(\tilde{\beta}_t^1 - \tilde{\beta}_{t-1}^1\right) + Q_0$ and line 91 draws Q from the inverse Wishart distribution. Line 94 draws the VAR error covariance R from the inverse Wishart distribution. Note that we use a flat prior for R in this example. One past the burn-in stage we save the draws for $\tilde{\beta}_T$, R and Q. We use the saved draws to compute the impulse response to a monetary policy shock and use sign restrictions to identify a monetary policy shock (lines 106 to 180). We assume that a monetary policy shock is one that increases interest rates, decreases inflation and output growth. The results for the time-varying impulse response are shown in 11. The 3-D surface charts show the impulse response horizon on the Y-axis and the time-series on the X-axis. These results show little evidence of significant variation in the impulse response functions across time for this dataset.

```
clear
addpath('functions');
% a TVP-VAR with dlog(GDP) dlog(CPI) and R for the US 1962 2004
%load data
data=xlsread('\data\usdata.xls')/100;
N=size(data,2);
L=2;   %number of lags in the VAR
Y=data;
X=[ lag0(Y,1) lag0(Y,2) ones(size(Y,1),1) ];
Y=Y(3:end,:);
X=X(3:end,:);
%step 1 set starting values and priors using a pre-sample of 10 years
T0=40;
y0=Y(1:T0,:);
x0=X(1:T0,:);
b0=x0\y0;   % \beta_0 = (X'_{0t} X_{0t})^{-1} (X'_{0t} Y_{0t})
e0=y0-x0*b0;
sigma0=(e0'*e0)/T0;
V0=kron(sigma0,inv(x0'*x0));  % p_{0\0} = \Sigma_0 \otimes (X'_{0t} X_{0t})^{-1}
%priors for the variance of the transition equation
Q0=V0*T0*3.5e-04;  % p_{0\0} \times T_0 \times \tau  %prior for the variance of the
transition equation error
P00=V0;      % p_{0\0}                % variance of the intial state vector
variance of state variable p[t-1/t-1]
beta0=vec(b0)';  % \beta_{0\0} = vec( \beta_0)'      % intial state vector   %state
variable  b[t-1/t-1]
%initialise
Q=Q0;
R=sigma0;
%remove intial Sample
Y=Y(T0+1:end,:);
X=X(T0+1:end,:);
T=rows(X);
%Gibbs sampling algorithm Step 2
reps=110000;
burn=109000;
mm=1;
for m=1:reps
m
%%Step 2a Set up matrices for the Kalman Filter
ns=cols(beta0);
F=eye(ns); fixed
mu=0;  fixed
beta tt=[];          %will hold the filtered state variable
ptt=zeros(T,ns,ns);    % will hold its variance
beta11=beta0;
p11=P00;
% %%%%%%%%%%Step 2b run Kalman Filter
for i=1:T
   x=kron(eye(N),X(i,:));  % (I_N \otimes X_t)
    %Prediction
beta10=mu+beta11*F';  % \beta_{t|t-1} = \mu + F \beta_{t-1|t-1}
p10=F*p11*F'+Q;       % p_{t|t-1} = F p_{t-1\t-1} F' + Q
```

FIGURE 7. Matlab code for a time-varying VAR

7. The Gibbs sampling algorithm for a Factor Augmented VAR

Our second example is based on the Factor augmented VAR model introduced in Bernanke *et al.* (2005). The FAVAR model can be written compactly as

$$\begin{aligned} X_{i,t} &= b_i F_t + \gamma_i FFR_t + v_{i,t} \\ Z_t &= c_t + \sum_{j=1}^{P} B_j Z_{t-j} + e_t \\ Z_t &= \{F_t, FFR_t\} \\ VAR(v_{i,t}) &= R, VAR(e_t) = Q \end{aligned} \tag{7.1}$$

```
yhat=(x*(beta10)')';
```

$$X_t\beta_{t|t-1}$$

```
eta=Y(i,:)-yhat;
```

$$\eta_{t|t-1} = Y_t - X_t\beta_{t|t-1}$$

```
feta=(x*p10*x')+R;
```

$$f_{t|t-1} = X_t p_{t|t-1} X_t' + R$$

```
%updating
K=(p10*x')*inv(feta);
```

$$K_t = p_{t|t-1} X_t' f_{t|t-1}^{-1}$$

```
beta11=(beta10'+K*eta')';
```

$$\beta_{t|t} = \beta_{t|t-1} + K_t \eta_{t|t-1}$$

```
p11=p10-K*(x*p10);
```

$$p_{t|t} = p_{t|t-1} - K_t X_t p_{t|t-1}$$

```
ptt(i,:,:)=p11;
beta tt=[beta tt;beta11];
end
%%%%%%%%%%%end of Kalman
Filter%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%step 2c Backward recursion to calculate the mean and variance of the
distribution of the state
%vector
chck=-1;
while chck<0 while loop to ensure VAR stable at each point in time
beta2 = zeros(T,ns);   %this will hold the draw of the state variable
wa=randn(T,ns);
error=zeros(T,N);
roots=zeros(T,1);
i=T;  %period t
p00=squeeze(ptt(i,:,:));
beta2(i,:)=beta tt(i:i,:)+(wa(i:i,:)*chol(p00));
```

$$\beta_T \sim N(\beta_{T|T}, p_{T|T})$$

```
%draw for beta in period t from N(beta tt,ptt)
error(i,:)=Y(i,:)-X(i,:)*reshape(beta2(i:i,:),N*L+1,N);  %var
residuals calculate var residuals in the same loop for convenience
roots(i)=stability(beta2(i,:)',N,L); checking stability at ith time
period roots(i)=1 if stability violated
%periods t-1..to .1
for i=T-1:-1:1
pt=squeeze(ptt(i,:,:));
bm=beta tt(i:i,:)+(pt*F'*inv(F*pt*F'+Q)*(beta2(i+1:i+1,:)-
beta_tt(i,:)*F')')';
```

$$\beta_{t|t,B_{t+1}} = \beta_{t|t} + p_{t|t}F'(Fp_{t|t}F' + Q)^{-1}(\beta_{t+1} - \mu - F\beta_{t|t})$$

```
pm=pt-pt*F'*inv(F*pt*F'+Q)*F*pt;
```

$$p_{t|t,B_{t+1}} = p_{t|t} - p_{t|t}F'(Fp_{t|t}F' + Q)^{-1}Fp_{t|t}$$

```
beta2(i:i,:)=bm+(wa(i:i,:)*chol(pm));
```

$$\beta_t \backslash \beta_{t+1} \sim N(\beta_{t|t,B_{t+1}}, p_{t|t,B_{t+1}})$$

```
error(i,:)=Y(i,:)-X(i,:)*reshape(beta2(i:i,:),N*L+1,N);
roots(i)=stability(beta2(i,:)',N,L);
end
if sum(roots)==0
    chck=1;
end
end
% step 3 sample Q from the IW distribution
errorq=diff(beta2);
```

$$\tilde{\beta}_t^1 - \tilde{\beta}_{t-1}^1$$

```
scaleQ=(errorq'*errorq)+Q0;
```

$$\left(\tilde{\beta}_t^1 - \tilde{\beta}_{t-1}^1\right)'\left(\tilde{\beta}_t^1 - \tilde{\beta}_{t-1}^1\right) + Q_0$$

Sample Q from its conditional posterior distribution

```
Q=iwpQ(T+T0,inv(scaleQ));
%step4 sample R from the IW distribution
```

FIGURE 8. Matlab code for a time-varying VAR (continued)

where $X_{i,t}$ is a $T \times M$ matrix containing a panel of macroeconomic and financial variables. FFR_t denotes the Federal Funds rate and F_t are the unobserved factors which summarise the information in the data $X_{i,t}$. The first equation is the observation equation of the model, while the second equation is the transition equation. Bernanke *et al.* (2005) consider a shock to the interest rate in the transition equation and calculate the impulse response of each variable in $X_{i,t}$.

It is instructive to consider the state-space representation of the FAVAR model in more detail. We assume in this example that the lag length in the transition equation equals 2 and there are 3 unobserved factors $F_t = \{F_{1t}, F_{2t}, F_{3t}\}$.

```
93 scaleR=(error'*error);
(Y_t − (c_t^1 + Σ_{j=1}^P B_{j,t}^1 Y_{t−j}))' (Y_t − (c_t^1 + Σ_{j=1}^P B_{j,t}^1 Y_{t−j})) + R_0
94 R=iwpQ(T,inv(scaleR));   Sample R from its conditional posterior distribution
95 if m>burn
96     %save output from Gibbs sampler
97     out1(mm,1:T,:)=beta2;
98     out2(mm,1:N,1:N)=R;
99     out3(mm,1:N*(N*L+1),1:N*(N*L+1))=Q;
100     mm=mm+1;
101 end
102 end
103 %save results
104 save tvp.mat out1 out2 out3
105 %compute irf to a policy shock using sign restrictions
106 horz=40;% impulse response horizon
107 irfmat=zeros(size(out1,1),T,horz,N); %empty matrix to save impulse
response to a policy shock
108 for i=1:size(out1,1);
109     sigma=squeeze(out2(i,:,:));
110    %sign restrictions
111         chck=-1;
112         while chck<0
113         K=randn(N,N);
114         QQ=getQR(K);
115         A0hat=chol(sigma);
116         A0hat1=(QQ*A0hat);  %candidate draw
117         for m=1:N
118         %check signs in each row
119         e1=A0hat1(m,1)<0;  %Response of Y
120         e2=A0hat1(m,2)<0;  %Response of P
121         e3=A0hat1(m,3)>0;  %Response of R
122
123         if e1+e2+e3==3
124             MP=A0hat1(m,:);
125             chck=10;
126         else
127             %check signs but reverse them
128          e1=-A0hat1(m,1)<0;  %Response of Y
129         e2=-A0hat1(m,2)<0;  %Response of P
130         e3=-A0hat1(m,3)>0;  %Response of R
131
132         if e1+e2+e3==3
133             MP=-A0hat1(m,:);
134             chck=10;
135         end
136         end
137         end
138         end
139         %re-shuffle rows of A0hat1 and insert MP in the first row
140         A0x=[]; %will hold rows of A0hat1 not equal to MP
141         for m=1:N
142             ee=sum(abs(A0hat1(m,:))==abs(MP));
143             if ee==0
144                 A0x=[A0x;A0hat1(m,:)];
145             end
146         end
147         A0new=[A0x;MP]; %A0 matrix to be used for impulse response
148     shock=[0 0 1];
149     for j=1:size(out1,2)
```

FIGURE 9. Matlab code for a time-varying VAR (continued)

Consider the observation equation of the model

$$\underbrace{\begin{pmatrix} X_{1t} \\ X_{2t} \\ X_{3t} \\ . \\ . \\ . \\ . \\ X_{Mt} \\ FFR_t \end{pmatrix}}_{\tilde{X}_{it}} = \underbrace{\begin{pmatrix} b_{11} & . & b_{13} & \gamma_1 & 0 & 0 & 0 & 0 \\ b_{21} & & . & . & 0 & 0 & 0 & 0 \\ . & & . & . & 0 & 0 & 0 & 0 \\ . & & . & . & 0 & 0 & 0 & 0 \\ . & & . & . & 0 & 0 & 0 & 0 \\ . & & . & . & 0 & 0 & 0 & 0 \\ . & & . & . & 0 & 0 & 0 & 0 \\ b_{M1} & & b_{M3} & \gamma_M & 0 & 0 & 0 & 0 \\ & & & 1 & 0 & 0 & 0 & 0 \end{pmatrix}}_{H} \underbrace{\begin{pmatrix} F_{1t} \\ F_{2t} \\ F_{3t} \\ FFR_t \\ F_{1t-1} \\ F_{2t-1} \\ F_{3t-1} \\ FFR_{t-1} \end{pmatrix}}_{\beta_t} + \underbrace{\begin{pmatrix} v_{1t} \\ v_{2t} \\ v_{3t} \\ . \\ . \\ . \\ . \\ v_{Nt} \\ 0 \end{pmatrix}}_{v_t} \tag{7.2}$$

```
        btemp=squeeze(out1(i,j,:));
        btemp=reshape(btemp,N*L+1,N);
        zz=irfsim(btemp,N,L,A0new,shock,horz+L);
        zz=zz./repmat(zz(1,3),horz,N);
        irfmat(i,j,:,:)=zz;
    end
end
TT=1964.75:0.25:2010.5;
HH=0:horz-1;
irf1=squeeze(median(irfmat(:,:,:,1),1));
irf2=squeeze(median(irfmat(:,:,:,2),1));
irf3=squeeze(median(irfmat(:,:,:,3),1));
figure(1)
subplot(2,2,1);
mesh(TT,HH,irf1')
ylabel('Impulse Horizon');
xlabel('Time');
axis tight
title('GDP growth');
subplot(2,2,2);
mesh(TT,HH,irf2')
ylabel('Impulse Horizon');
xlabel('Time');
axis tight
title('Inflation');
subplot(2,2,3);
mesh(TT,HH,irf3')
ylabel('Impulse Horizon');
xlabel('Time');
axis tight
title('Federal Funds Rate');
```

Published with MATLAB® 7.9

FIGURE 10. Matlab code for a time-varying VAR (continued)

The left hand side of the observation equation 7.2 contains the dataset $X_{i,t}$ with the Funds rate as the last variable (thus $\tilde{X}_{it} = \{X_{i,t}, FFR_t\}$) . $X_{i,t}$ is related to the three factors via the factor loadings b_{ij} where $i = 1...M$ and $j = 1, 2, 3$. $X_{i,t}$ is related to the Federal Funds rate via the coefficients γ_i. Bernanke *et al.* (2005) assume that γ_i are non-zero only for fast moving financial variables. FFR_t appears in the state vector β_t (even though it is observed) as we want it to be part of the transition equation. Therefore the last row of the coefficient matrix H describes the identity $FFR_t = FFR_t$. The state vector contains the first lag of all state variables as we want two lags in the VAR

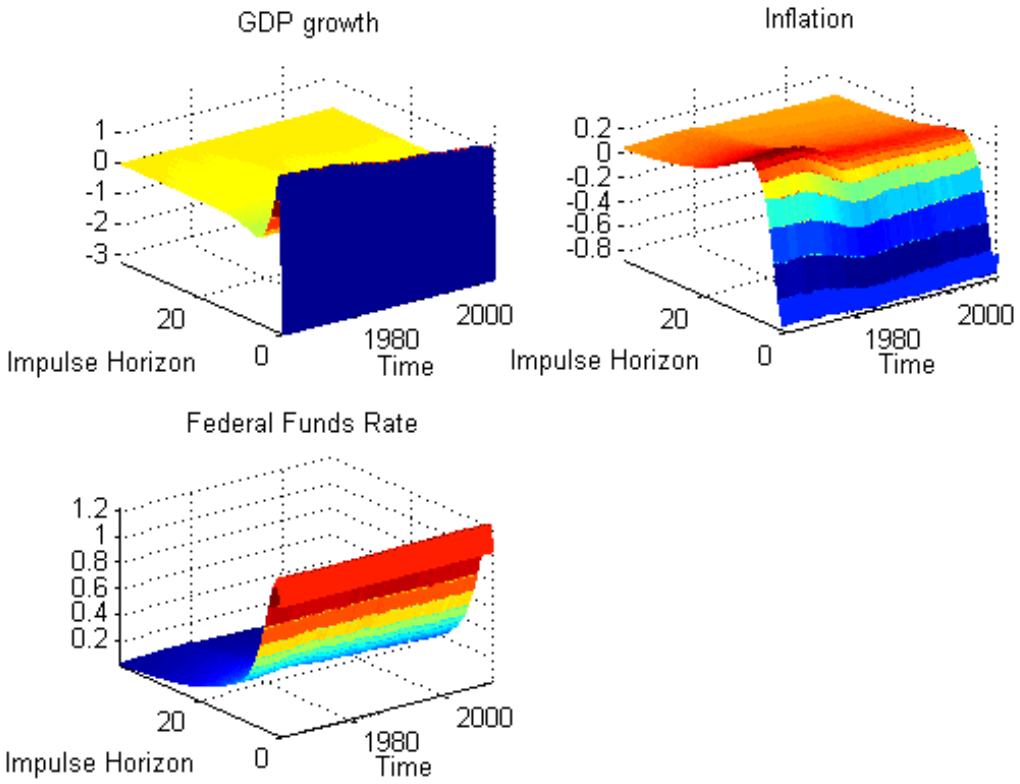

FIGURE 11. Time-varying impulse responses to a monetary policy shock

that forms the transition equation. Note also that

$$VAR\left(v_t\right) = R = \begin{pmatrix} R_1 & 0 & 0 & 0 & 0 \\ 0 & R_2 & 0 & 0 & 0 \\ 0 & 0 & . & 0 & 0 \\ 0 & 0 & 0 & R_M & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix} \tag{7.3}$$

The transition equation of the model is

$$\underbrace{\begin{pmatrix} F_{1t} \\ F_{2t} \\ F_{3t} \\ FFR_t \\ F_{1t-1} \\ F_{2t-1} \\ F_{3t-1} \\ FFR_{t-1} \end{pmatrix}}_{\beta_t} = \underbrace{\begin{pmatrix} u_1 \\ u_2 \\ u_3 \\ u_4 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}}_{\mu} + \underbrace{\begin{pmatrix} A_{11} & A_{12} & A_{13} & A_{14} & A_{15} & A_{16} & A_{17} & A_{18} \\ A_{21} & A_{22} & A_{23} & A_{24} & A_{25} & A_{26} & A_{27} & A_{28} \\ A_{31} & A_{32} & A_{33} & A_{34} & A_{35} & A_{36} & A_{37} & A_{38} \\ A_{41} & A_{42} & A_{43} & A_{44} & A_{45} & A_{46} & A_{47} & A_{48} \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \end{pmatrix}}_{F} \underbrace{\begin{pmatrix} F_{1t-1} \\ F_{2t-1} \\ F_{3t-1} \\ FFR_{t-1} \\ F_{1t-2} \\ F_{2t-2} \\ F_{3t-2} \\ FFR_{t-2} \end{pmatrix}}_{\beta_{t-1}} + \underbrace{\begin{pmatrix} e_{1t} \\ e_{2t} \\ e_{3t} \\ e_{4t} \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}}_{e_t} \tag{7.4}$$

Note that this is simply a VAR(2) in F_{1t}, F_{2t}, F_{3t} and FFR_t written in first order companion form to make consistent with the usual form of a transition equation (i.e. the transition equation needs to be in AR(1) form). Note that:

$$VAR(e_t) = Q = \begin{pmatrix} Q_{11} & Q_{12} & Q_{13} & Q_{14} & 0 & 0 & 0 & 0 \\ Q_{12} & Q_{22} & Q_{23} & Q_{24} & 0 & 0 & 0 & 0 \\ Q_{13} & Q_{23} & Q_{33} & Q_{34} & 0 & 0 & 0 & 0 \\ Q_{14} & Q_{24} & Q_{23} & Q_{44} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \tag{7.5}$$

where the zeros result from the fact that the last 4 equations in the transition equation describe identities. Therefore the matrix Q is singular in this FAVAR model. This implies that the Carter and Kohn recursion has to be generalised slightly to take this singularity into account as discussed above. This modification implies that we use $\mu^*, F^*, Q^*, \beta_{t+1}^*$

in equations 3.13 and 3.14 where $\mu^*, F^*, Q^*, \beta^*_{t+1}$ denote the first jv rows of μ, F, Q, β_{t+1}. In our example $jv = 3$ as we extract three factors (the FFR_t equation in the observation equation is an identity).

The Gibbs sampling algorithm can be discerned by imagining the situation where the factors F_t are observed. Give the factors, the observation equation is just M linear regressions of the form $X_{i,t} = b_{ij}F_{jt} + \gamma_i FFR_t + v_{i,t}$ and the conditional distributions studies in Chapter 1 apply immediately to sample b_{ij} and γ_i (i.e. the elements of H) and R. Similarly, given the factors, the transition equation is simply a VAR model. The conditional distributions in Chapter 2 can be used to sample u, F and Q. Finally, given a draw for H, R, u, F and Q the model can be cast into the state-space form shown in equations 7.2 and 7.4. Then the Carter and Kohn algorithm can be used to draw F_t from its conditional distribution. The Gibbs sampling algorithm consists of the following steps

Step 1 Set priors and starting values. The prior for the factor loadings is normal. Let $H_i = \{b_{ij}, \gamma_i\}$. Then $p(H_i) \sim N(H_{i0}, \Sigma_{H_i})$. The prior for the diagonal elements of R is inverse Gamma and given by $p(R_{ii}) \sim IG(R_{ii0}, V_{R0})$. The prior for the VAR parameters μ, F and Q can be set using any of the priors for VARs considered in the previous chapter. For example, one may consider setting an independent Normal inverse Wishart prior. Collecting the VAR coefficients in the matrix B and the non-zero elements of Q in the matrix Ω this prior can be represented as $p(B) \sim N(B_0, \Sigma_B)$ and $p(\Omega) \sim IW(\Omega_0, V_0)$. The Kalman filter requires the initial value of the state vector $\beta_t = \begin{pmatrix} F_{1t} \\ F_{2t} \\ F_{3t} \\ FFR_t \\ F_{1t-1} \\ F_{2t-1} \\ F_{3t-1} \\ FFR_{t-1} \end{pmatrix}$. One can use principal components to get an initial estimate of F_{1t}, F_{2t} and F_{3t} to set $\beta_{0\backslash 0}$. The principal component estimate also provides a good starting value for the factors $F_t = F_{1t}, F_{2t}$ and F_{3t}. One can arbitrarily set $R_{ii} = 1$ and Ω to an identity matrix to start the algorithm.

Step 2. Conditional on the factors F_t and R_{ii} sample the factor loadings $H_i = \{b_{ij}, \gamma_i\}$ from their conditional distributions. For each variable in X_{it} the factor loadings have a normal conditional posterior (as described in Chapter 1) $H(H_i \backslash F_t, R_{ii}) \sim N(H_i^*, V_i^*)$

$$H_i^* = \left(\Sigma_{H_i}^{-1} + \frac{1}{R_{ii}} Z_t' Z_t\right)^{-1} \left(\Sigma_{H_i}^{-1} H_{i0} + \frac{1}{R_{ii}} Z_t' X_{it}\right)$$

$$V_i^* = \left(\Sigma_{H_i}^{-1} + \frac{1}{R_{ii}} Z_t' Z_t\right)^{-1}$$

where $Z_t = \{F_{1t}, F_{2t}, F_{3t}, FFR_t\}$ if the ith series X_{it} is a fast moving data series which has a contemporaneous relationship with the Federal Funds rate (e.g. stock prices) and $Z_t = \{F_{1t}, F_{2t}, F_{3t}\}$ if the ith series X_{it} is a slow moving data series which has no contemporaneous relationship with the Federal Funds rate (e.g. GDP). Note that as F_{1t}, F_{2t}, F_{3t} and H_i are both estimated the model is unidentified. Bernanke *et al.* (2005) suggest fixing the top 3×3 block of b_{ij} to an identity matrix and the top 3×1 block of γ_i to zero for identification. See Bernanke *et al.* (2005) for more details on this issue.

Step 3. Conditional on the factors F_t and the factor loadings $H_i = \{b_{ij}, \gamma_i\}$ sample the variance of the error terms of the observation equation R_{ii} from the inverse Gamma distribution with scale parameter $(X_{it} - Z_t H_i)'(X_{it} - Z_t H_i) + R_{ii0}$ with degrees of freedom $T + V_{R0}$ where T is the length of the estimation sample.

Step 4. Conditional on the factors F_t and the error covariance matrix Ω, the posterior for the VAR coefficients B (recall $B = \{\mu, F\}$ the coefficients in the transition equation of the model) is normal (see Chapter 2) and given as $H(B \backslash F_t, \Omega) \sim N(B^*, D^*)$ where B_0, Σ_B

$$\begin{aligned} B^* &= \left(\Sigma_B^{-1} + \Omega^{-1} \otimes \bar{X}_t' \bar{X}_t\right)^{-1} \left(\Sigma_B^{-1} vec(B_0) + \Omega^{-1} \otimes \bar{X}_t' \bar{X}_t vec(\hat{B})\right) \\ D^* &= \left(\Sigma_B^{-1} + \Omega^{-1} \otimes \bar{X}_t' \bar{X}_t\right)^{-1} \end{aligned}$$

where $\bar{X}_t = \{F_{t-1}, FFR_{t-1}, F_{t-2}, FFR_{t-2}, 1\}$ and $\hat{B}$ is the OLS estimate of B.

Step 5. Conditional on the factors F_t and the VAR coefficients B the error covariance Ω has a inverse Wishart posterior with scale matrix $\left(Y_t - \bar{X}_t B\right)'\left(Y_t - \bar{X}_t B\right) + \Omega_0$ and degrees of freedom $T + V_0$.

Step 6. Given H_i, R, B and Ω the model can be cast into state-space form and then the factors F_t are sampled via the Carter and Kohn algorithm.

Step 7 Repeat steps 2 to 6 M times and use the last L values for inference

7.1. Matlab code for the FAVAR model. We estimate a FAVAR model using UK data over the period 1970Q1 to 2006Q1. We use 40 Macroeconomic and financial time series along with the Bank of England policy rate to estimate the model and consider the impact of a monetary policy shock.

The Matlab code for this example (example4.m) can be seen in figures 12, 13, 14, 15, 16 and 17.

```
1 clear
2 addpath('functions');
3 [ data0 junk ]=xlsread('\data\datain.xls');
4 [ junk names ]=xlsread('\data\names.xls');
5 %names=names(1,2:end);
6 index=xlsread('\data\index.xls');
7 dindex=index(:,1); %dindex=1 for series that are log differenced
dindex=3 differencing without logs
8 index=index(:,2);  %index=1 for 'fast moving' series
9 %first difference the data where appropriate
10 data=[];
11 for i=1:cols(data0);
12     if dindex(i)==1
13         dat=log(data0(:,i));
14         dat=diff(dat)*100;
15     elseif dindex(i)==3
16         dat=diff(data0(:,i));
17     else
18         dat=data0(2:end,i);
19     end
20     data=[data dat];
21 end
22 %standardise the data
23 data=standardise(data);
24 %load policy rate and standardize it
25 z=xlsread('\data\baserate.xls');
26 z=z(2:end);
27 z=standardise(z);
28 KK=3;  %number of factors
29 L=2;  %number of lags in the VAR
30 N=KK+1; %number of Variables in var K factors plus the interest rate
31 NN=cols(data);% size of the panel
32 T=rows(data)
33 %step 1 of the algorithm set starting values and priors
34 %get an intial guess for the factor via principal components
35 pmat=extract(data,KK);

36 beta0=[pmat(1,:) z(1) zeros(1,N)];  %state vector S[t-1/t-1]
37 ns=cols(beta0);

38 P00=eye(ns);  %P[t-1/t-1]
One can arbitrarily set R_ii = 1 and Ω to an identity matrix to start the algorithm
39 rmat=ones(NN,1); %arbitrary starting value for the variance of the
idiosyncratic component
40 Sigma=eye(N);  %arbitrary starting value for the variance of VAR
errors
41 %flat prior for the factor loadings,variances and VAR
42 reps=5000;
43 burn=4000;
44 mm=1;
45 for m=1:reps;
46 %gibbs sampling
47 %step 2 sample factor loadings
48 fload=[];
49 floadr=[];
50 error=[];
51 for i=1:NN
52     y=data(:,i);
53     if index(i)==0
```

FIGURE 12. Code for the FAVAR model

Lines 3 and 4 load the $T \times M$ panel of UK data and the variable names ($M = 40$). Line 6 reads a variable called index. The first column is a $M \times 1$ vector which equals 1 if the corresponding data series in the panel has to be first differenced. The second column is a $M \times 1$ vector which equals 1 if the corresponding data series is a fast-moving variable (like an asset price) and will have a contemporaneous relationship with the policy interest rate i.e. $\gamma_i \neq 0$ for this variable. Lines 10 to 23 transform the data to stationarity and standardises it. Lines 25 to 27 read the bank rate and standardises it. Line 35 extracts three principal components from the dataset to use as starting values for the three factors in this example. Line 36 defines $\beta_{0\backslash 0}$ =[pmat(1,:) z(1) zeros(1,N)]. Notice that there are 8 state variables: 3 factors, the interest rate and the first lags of these 4 state variables and thus $\beta_{0\backslash 0}$ is 1×8. Line 38 sets

$Z_t = \{F_{1t}, F_{2t}, F_{3t}\}$

```
54          x=pmat;
55      else
56          x=[pmat z];
57      end
58      M=inv(x'*x)*(x'*y);
```

$Z_t = \{F_{1t}, F_{2t}, F_{3t}, FFR_t\}$

$H_i^* = \left(\Sigma_{H_i}^{-1} + \frac{1}{R_{ii}} Z_t' Z_t\right)^{-1} \left(\Sigma_{H_i}^{-1} H_{i0} + \frac{1}{R_{ii}} Z_t' X_{it}\right)$ with a flat prior

$V_i^* = \left(\Sigma_{H_i}^{-1} + \frac{1}{R_{ii}} Z_t' Z_t\right)^{-1}$ with a flat prior

```
59      V=rmat(i)*inv(x'*x);
60      %draw
61      ff=M+(randn(1,cols(x))*cholx(V))';
62
63      %save
64      if index(i)==0;
65          fload=[fload;ff'];
66          floadr=[floadr;0];
67      else
68              fload=[fload;ff(1:end-1)'];
69          floadr=[floadr;ff(end)];
70      end
71      error=[error y-x*ff];
72  end
73  %for identification top K by K block of fload is identity
74  fload(1:KK,1:KK)=eye(KK);
75  %for identification top K by 1 block of Floadr is zero
76  floadr(24:24+KK-1,1)=zeros(KK,1);
77  %step 3 sample variance of the idiosyncratic components from inverse
78  %gamma
```

sample the variance of the error terms of the observation equation R_{ii} from the inverse Gamma distribution with scale parameter $(X_{it} - Z_t H_i)'(X_{it} - Z_t H_i) + R_{i0}$ with degrees of freedom $T + V_{i0}$ where T is the length of the estimation sample.

```
79  rmat=[];
80  for i=1:NN
81      rmati= IG(0,0,error(:,i));
82      rmat=[rmat rmati];
83  end
84  %step 4 sample VAR coefficients
85  Y=[pmat z];
86  X=[lag0(Y,1) lag0(Y,2) ones(rows(Y),1)];
87  Y=Y(2:end,:);
88  X=X(2:end,:);
89  M=vec(inv(X'*X)*(X'*Y));  %conditional mean
```

$B^* = (\Sigma_B^{-1} + \Omega^{-1} \otimes \bar{X}_t' \bar{X}_t)^{-1} \left(\Sigma_B^{-1} vec(B_0) + \Omega^{-1} \otimes \bar{X}_t' \bar{X}_t vec(\hat{B})\right)$ with a flat prior

```
90  V=kron(Sigma,inv(X'*X)); %conditional variance
```

$D^* = (\Sigma_B^{-1} + \Omega^{-1} \otimes \bar{X}_t' \bar{X}_t)^{-1}$ with a flat prior

```
91  chck=-1;                %make sure VAR is stationary
92  while chck<0
93  beta=M+(randn(1,N*(N*L+1))*cholx(V))';  %draw for VAR coefficients
94  S=stability(beta,N,L);
95  if S==0
96      chck=10;
97  end
98  end
99  beta1=reshape(beta,N*L+1,N);
100 errorsv=Y-X*beta1;
```

Conditional on the factors F_t and the VAR coefficients B the error covariance Ω has a inverse Wishart posterior with scale matrix $(Y_t - \bar{X}_t B)'(Y_t - \bar{X}_t B) + \Omega_0$ and degrees of freedom $T+V_0$.

FIGURE 13. Code for the FAVAR model (continued)

$p_{0\backslash 0}$ as a 8×8 identity matrix. We arbitrarily set $R_{ii} = 1$ and $\Omega = I$ to start the algorithm on lines 39 and 40. Note that following Bernanke *et al.* (2005) we will not use prior distributions for the regression or VAR coefficients which will imply that the conditional posteriors collapse to OLS formulae.

Lines 48 to 72 sample the factor loadings. The code loops through the 40 data series and selects each as the dependent variable (line 52) to be regressed on the factors only for slow moving series (line 54) or the factors and the policy interest rate for fast moving series (line 56). Line 58 calculates the mean of the conditional posterior distribution of the factor loadings without the priors

$$H_i^* = (Z_t' Z_t)^{-1} (Z_t' X_{it})$$

```
101 %sample VAR covariance
102 scale=errorsv'*errorsv;
103 Sigma=iwpQ(T,inv(scale));
104 %step 5 prepare matrices for the state space
105 %Y=H*factors+e
106 %factors=MU+F*Factors(-1)+v
107 %e~N(0,R)
108 %v~N(0,Q)
109 %matrix of factor loadings
```

$$H = \begin{pmatrix} b_{11} & . & b_{13} & \gamma_1 & 0 & 0 & 0 & 0 \\ b_{21} & & . & . & 0 & 0 & 0 & 0 \\ . & & . & . & 0 & 0 & 0 & 0 \\ . & & . & . & 0 & 0 & 0 & 0 \\ . & & . & . & 0 & 0 & 0 & 0 \\ . & & . & . & 0 & 0 & 0 & 0 \\ . & & . & . & 0 & 0 & 0 & 0 \\ b_{M1} & & b_{M3} & \gamma_M & 0 & 0 & 0 & 0 \\ & & & 1 & 0 & 0 & 0 & 0 \end{pmatrix}$$

```
110 H=zeros(NN,(KK+1)*L);
111 H(1:rows(fload),1:KK+1)=[fload floadr];
112 H(rows(floadr)+1,KK+1)=1;
113 %matrix R
```

$$VAR(v_t) = R = \begin{pmatrix} R_1 & 0 & 0 & 0 & 0 \\ 0 & R_2 & 0 & 0 & 0 \\ 0 & 0 & . & 0 & 0 \\ 0 & 0 & 0 & R_M & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

```
114 R=diag([rmat 0]);
115 %vector MU
```

$$\mu = \begin{pmatrix} u_1 \\ u_2 \\ u_3 \\ u_4 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}$$

```
116 MU=[beta1(end,:)';zeros(N*(L-1),1)];
117 %matrix F
```

FIGURE 14. Code for the FAVAR model (continued)

and line 59 calculates the variance of this distribution (without the prior information).

$$V_i^* = \left(\frac{1}{R_{ii}} Z_t' Z_t\right)^{-1}$$

The coefficients b_{ij} are stored in the matrix fload and the coefficients γ_i in floadr. Lines 74 and 76 impose the identification conditions and fix the top 3×3 block of fload to an identity matrix and top 3×1 block of floadr to 0. Lines 79 to 83 sample R_{ii} from the inverse Gamma distribution (using the function IG in the functions folder) with prior degrees of freedom and the prior scale matrix set to 0 (hence using information from the data only). Lines 85

$$\begin{pmatrix} A_{11} & A_{12} & A_{13} & A_{14} & A_{15} & A_{16} & A_{17} & A_{18} \\ A_{21} & A_{22} & A_{23} & A_{24} & A_{25} & A_{26} & A_{27} & A_{28} \\ A_{31} & A_{32} & A_{33} & A_{34} & A_{35} & A_{36} & A_{37} & A_{38} \\ A_{41} & A_{42} & A_{43} & A_{44} & A_{45} & A_{46} & A_{47} & A_{48} \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \end{pmatrix}$$

F

```
118 F=[beta1(1:N*L,:)';eye(N*(L-1),N*L)];
119 %matrix Q
```

$$VAR(e_t) = Q = \begin{pmatrix} Q_{11} & Q_{12} & Q_{13} & Q_{14} & 0 & 0 & 0 & 0 \\ Q_{12} & Q_{22} & Q_{23} & Q_{24} & 0 & 0 & 0 & 0 \\ Q_{13} & Q_{23} & Q_{33} & Q_{34} & 0 & 0 & 0 & 0 \\ Q_{14} & Q_{24} & Q_{23} & Q_{44} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

```
120 Q=zeros(rows(F),rows(F));
121 Q(1:N,1:N)=Sigma;
122 %Carter and Kohn algorithm to draw the factor
123 beta tt=[];           %will hold the filtered state variable
124 ptt=zeros(T,ns,ns);    % will hold its variance
125 % %%%%%%%%%%Step 6a run Kalman Filter
126 i=1;
127 x=H; This is no longer data but a matrix of coefficients
128 %Prediction
129 beta10=MU+beta0*F';
130 p10=F*P00*F'+Q;
131 yhat=(x*(beta10)')';
132 eta=[data(i,:) z(i,:)]-yhat;
```

$$\begin{pmatrix} X_{1t} \\ X_{2t} \\ X_{3t} \\ . \\ . \\ . \\ . \\ X_{Mt} \\ FFR_t \end{pmatrix}$$

$\tilde{X}_{it}$

```
133 feta=(x*p10*x')+R;
134 %updating
135 K=(p10*x')*inv(feta);
136 beta11=(beta10'+K*eta')';
```

FIGURE 15. Code for the FAVAR model (continued)

and 86 set up the left hand side and the right hand side variables for the VAR model using the factors (pmat) and the policy rate. Lines 89 and 90 calculate the mean and variance of the conditional distribution of the VAR coefficients (without prior information these are just OLS). Line 93 draws the VAR coefficients ensuring stability. Lines 102 and 103 draw the covariance matrix Ω from the inverse Wishart distribution. We now have a draw for all parameters of the state space representation so we build the matrices necessary to cast the FAVAR into the state space form. Lines 110 to 112 build the matrix H seen in equation 7.2. Line 114 builds the covariance matrix of the error term R. Line 116 builds the matrix μ seen in equation 7.4. Line 118 builds the matrix F, while line 120 builds the matrix Q. With the matrices of the state space representation in hand we start the Carter and Kohn algorithm by running

```
137 p11=p10-K*(x*p10);
138 beta_tt=[beta_tt;beta11];
139 ptt(i,:,:)=p11;
140 for i=2:T
141     %Prediction
142 beta10=MU+beta11*F';
143 p10=F*p11*F'+Q;
144 yhat=(x*(beta10)')';
145 eta=[data(i,:) z(i,:)]-yhat;
146 feta=(x*p10*x')+R;
147 %updating
148 K=(p10*x')*inv(feta);
149 beta11=(beta10'+K*eta')';
150 p11=p10-K*(x*p10);
151 ptt(i,:,:)=p11;
152 beta tt=[beta tt;beta11];
153 end
154 % Backward recursion to calculate the mean and variance of the
distribution of the state
155 %vector
156 beta2 = zeros(T,ns);   %this will hold the draw of the state
variable
157 jv=1:3; %index of state variables to extract 3 factors to extract
158 wa=randn(T,ns);
                F*
159 f=F(jv,:);
                Q*
160 q=Q(jv,jv);
                μ*
161 mu=MU(jv);
162 i=T;  %period t
163 p00=squeeze(ptt(i,jv,jv));
164 beta2(i,jv)=beta tt(i:i,jv)+(wa(i:i,jv)*cholx(p00));   %draw for
beta in period t from N(beta_tt,ptt)
165 %periods t-1..to .1
166 for i=T-1:-1:1
167 pt=squeeze(ptt(i,:,:));
β_{t|t} + P_{t|t}F'(FP_{t|t}F' + Q)^{-1}(β_{t+1} − μ − Fβ_{t|t})
168 bm=beta tt(i:i,:)+(pt*f'*inv(f*pt*f'+q)*(beta2(i+1:i+1,jv)-mu-
beta tt(i,:)*f')')';
P_{t|t} − P_{t|t}F'(FP_{t|t}F' + Q)^{-1}FP_{t|t}
169 pm=pt-pt*f'*inv(f*pt*f'+q)*f*pt;
170 beta2(i:i,jv)=bm(jv)+(wa(i:i,jv)*cholx(pm(jv,jv)));
171 end
172 pmat=beta2(:,1:3);   %update the factors
173 if m>burn
174     %compute impulse response
175     A0=cholx(Sigma);
176     yhat=zeros(36,N);
177 vhat=zeros(36,N);
178 vhat(3,1:N)=[0 0 0 1];
179 for i=3:36
180  yhat(i,:)=[yhat(i-1,:) yhat(i-2,:)
1]*[beta1(1:N*L,:);zeros(1,N)]+vhat(i,:)*A0;
181 end
182 yhat1=yhat*H(:,1:KK+1)';  %impulse response for the panel
183 irfmat(mm,1:36,1:NN+1)=(yhat1);
184 mm=mm+1;
185 end
```

FIGURE 16. Code for the FAVAR model (continued)

the Kalman filter from lines 123 to 153. Note a minor difference to the previous example is that the observation equation now does not have a regressor on the right hand side. Hence on line 127 x is set equal to the matrix H. Line 156 starts the backward recursion. Recall that the last 5 state variables represent identities and Q is singular. Therefore we will only work with the first 3 rows (and columns for covariance matrices) of μ, F, Q and B_{t+1}. Lines 159 to 161 create μ^*, F^*, Q^*. Lines 168 and 169 are the modified Carter and Kohn updating equations. Line 172 sets the factors pmat equal to the last draw using the Carter and Kohn algorithm and we return to the first step of the Gibbs sampler. Once past the burn-in period we calculate an impulse response of the factors to a shock to the bank rate in the transition equation using a Cholesky decomposition of the covariance matrix to form the $A0$ matrix. Line

```
186
187 end
188 irf=prctile(irfmat,[50 16 84],1);
189 figure(1)
190 j=1
191 for i=1:size(irf,3)
192 subplot(4,10,j)
193 plotx1(squeeze(irf(:,:,i))');
194 title(strcat('\fontsize{8}', names(i)))
195 j=j+1
196 axis tight
197 end
```

Published with MATLAB® 7.9

FIGURE 17. Code for the FAVAR model (continued)

182 uses the observation equation of the model to calculate the impulse response of all the underlying data series. The estimated impulse responses are shown in 18.

8. Further reading

- Kim and Nelson (1999) chapter 3 is an excellent intuitive intrduction to state space models.
- Hamilton (1994) chapter provides a more formal derivation of the Kalman filter.
- Kim and Nelson (1999) chapter 8 provides a detailed description of Gibbs sampling for state space models.
- Code and a monograph by Gary Koop:

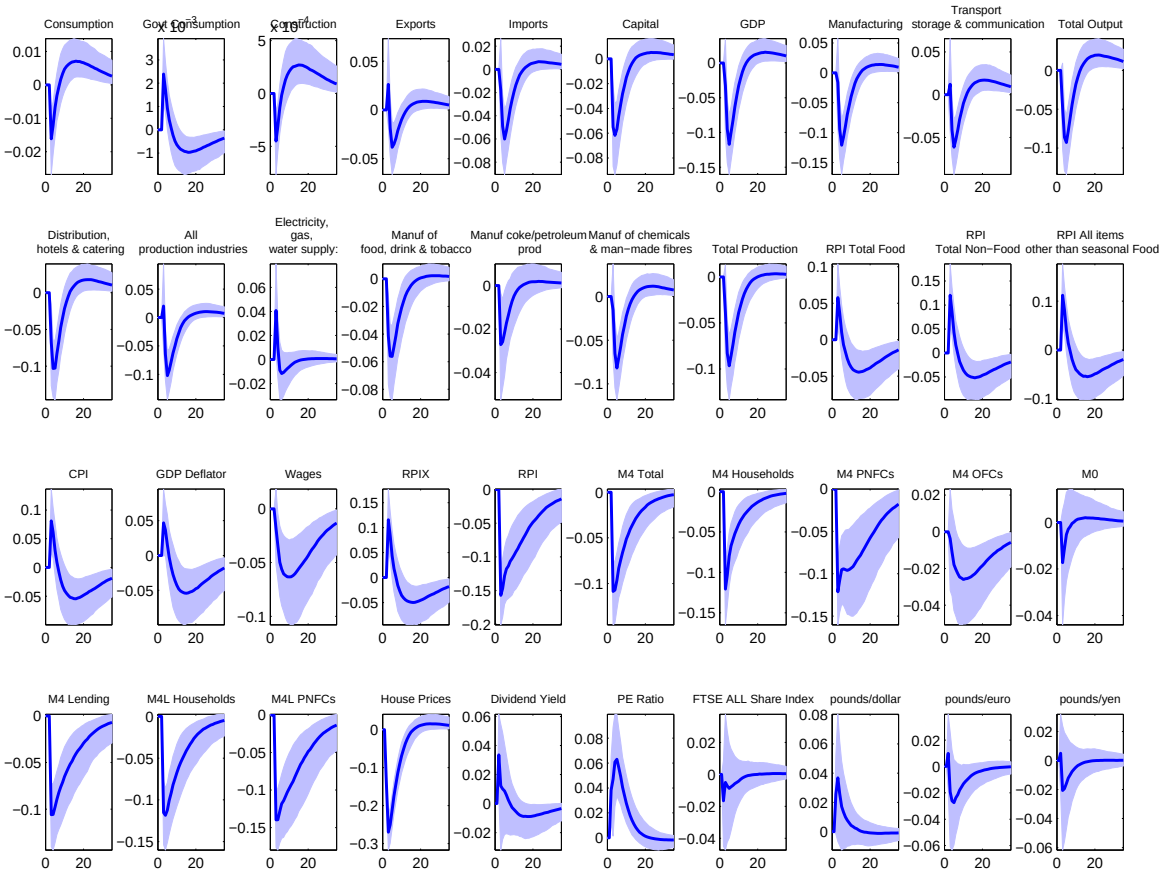

FIGURE 18. Impulse response of UK Macroeconomic series to a monetary policy shock using a FAVAR model.

http://personal.strath.ac.uk/gary.koop/bayes_matlab_code_by_koop_and_korobilis.html.

Part 2

The Metropolis Hastings algorithm

CHAPTER 4

An introduction to the the Metropolis Hastings Algorithm

1. Introduction

The Gibbs sampling algorithm relies on the availability of conditional distributions to be operational. In many cases (of practical relevance) conditional distributions are not available in closed form. An important example of such a situation is the estimation of Dynamic Stochastic General Equilibrium (DSGE) models where the conditional distribution of different parameter blocks is unavailable. In such cases an algorithm more general than the Gibbs sampler is required to approximate the posterior distribution. The Metropolis Hastings algorithm offers such an alternative. This chapter introduces this algorithm and discusses its implementation in Matlab for a number of important cases. The algorithm is applied to DSGE models in the next chapter.

2. The Metropolis Hastings algorithm

In this section we describe the Metropolis Hastings (MH) algorithm in a general setting. We follow that with a number of specific examples and Matlab code in the subsequent sections.

Suppose that we are interested in drawing samples from the following distribution (this is referred to as the target density below)

$$\pi(\Phi) \tag{2.1}$$

where Φ is a $K \times 1$ vector which represents a set of parameters. $\pi(\Phi)$ could be a posterior distribution where direct sampling is not possible and the Gibbs sampler is not operational as conditional distributions of different blocks of the parameters Φ are unknown. However, given a value for $\Phi = \Phi^*$ we are able to evaluate the density at Φ^* i.e. calculate $\pi(\Phi^*)$.

In this situation the MH algorithm can be used to take draws from $\pi(\Phi)$ using the following steps

Step 1 Specify a *candidate density* $q\left(\Phi^{G+1} \backslash \Phi^{G}\right)$ where G indexes the draw of the parameters Φ. One must be able to draw samples from this density. We discuss the exact specification of $q\left(\Phi^{G+1} \backslash \Phi^{G}\right)$ below.

Step 2 Draw a candidate value of the parameters $\Phi^{G+1}$from the candidate density $q\left(\Phi^{G+1} \backslash \Phi^{G}\right)$.

Step 3 Compute the probability of accepting Φ^{G+1} (denoted by α) using the expression

$$\alpha = \min\left(\frac{\pi\left(\Phi^{G+1}\right)/q\left(\Phi^{G+1} \backslash \Phi^{G}\right)}{\pi\left(\Phi^{G}\right)/q\left(\Phi^{G} \backslash \Phi^{G+1}\right)}, 1\right) \tag{2.2}$$

The numerator of this expression is the target density evaluated at the new draw of the parameters $\pi\left(\Phi^{G+1}\right)$ divided by the candidate density evaluated at the new draw of the parameters $q\left(\Phi^{G+1} \backslash \Phi^{G}\right)$. The denominator is the same expression evaluated at the previous draw of the parameters.

Step 4 If the acceptance probability α is large enough retain the new draw Φ^{G+1}, otherwise retain the old draw Φ^{G}. How do we decide if α is large enough in practice? We draw a number u from the standard uniform distribution. If $u < \alpha$. accept Φ^{G+1} otherwise keep Φ^{G}.[1].

Step 5 Repeat steps 2 to 4 M times and base inference on the last L draws. In other words, the empirical distribution using the last L draws is an approximation to target density. We discuss convergence of the MH algorithm below.

Note that one can think of the Gibbs sampler as a special case of the MH algorithm–i.e. a situation where the candidate density $q\left(\Phi^{G+1} \backslash \Phi^{G}\right)$ coincides with the target density and the acceptance probability assigned to every draw equals 1.

3. The Random Walk Metropolis Hastings Algorithm

The random walk MH algorithm offers a simple way of specifying the candidate density $q\left(\Phi^{G+1} \backslash \Phi^{G}\right)$ and is therefore widely used in applied work. As the name suggests, the random walk MH algorithm specifies the candidate generating density as a random walk

$$\Phi^{G+1} = \Phi^{G} + e \tag{3.1}$$

[1]This essentially means that we accept the draw with probability α if this experiment is repeated many times. For e.g if $\alpha = 0.1$, and if we 1000 replications we should expect 100 of the 1000 draws to have $u < \alpha$

where $e \sim N(0, \Sigma)$ is a $K \times 1$ vector. Note that $e_t = \Phi^{G+1} - \Phi^G$ is normally distributed. As the normal distribution is symmetric, the density $p\left(\Phi^{G+1} - \Phi^G\right)$ equals $p\left(\Phi^G - \Phi^{G+1}\right)$. In other words, $q\left(\Phi^{G+1} \backslash \Phi^G\right) = q\left(\Phi^G \backslash \Phi^{G+1}\right)$ under this random walk candidate density and the formula for the acceptance probability in equation 2.2 simplifies to

$$\alpha = \min\left(\frac{\pi\left(\Phi^{G+1}\right)}{\pi\left(\Phi^G\right)}, 1\right) \tag{3.2}$$

The random walk MH algorithm, therefore, works in the following steps:

Step 1 Specify a starting value for the parameters Φ denoted by Φ^0 and fix Σ the variance of shock to the random walk candidate generating density.

Step 2 Draw a new value for the parameters Φ^{New} using

$$\Phi^{New} = \Phi^{Old} + e \tag{3.3}$$

where $\Phi^{Old} = \Phi^0$ for the first draw

Step 3 Compute the acceptance probability

$$\alpha = \min\left(\frac{\pi\left(\Phi^{New}\right)}{\pi\left(\Phi^{Old}\right)}, 1\right) \tag{3.4}$$

If $\alpha > u \backsim U(0,1)$, then retain Φ^{New} and set $\Phi^{Old} = \Phi^{New}$, otherwise retain Φ^{Old}.

Step 4 Repeat steps 2 and 3 M times and use the last L draws for inference.

Note that Σ the variance of e_t is set by the researcher. A higher value for Σ could mean a lower rate of acceptances across the MH iterations (i.e. the acceptance rate is defined as the number of accepted MH draws divided by the total number of MH draws) but would mean that the algorithm explores a larger parameter space. In contrast, a lower value for Σ would mean a larger acceptance rate with the algorithm considering a smaller number of possible parameter values. The general recommendation is to choose Σ such that the acceptance rate is between 20% to 40%. We consider the choice of Σ in detail in the examples described below.[2]

3.1. Estimating a non-linear regression via the random walk MH algorithm. As an example, consider the estimation of the following non-linear regression model

$$Y_t = B_1\left(X_t^{B_2}\right) + v_t, v_t \sim N(0, \sigma^2) \tag{3.5}$$

and for the moment assume no prior information is used in estimating B_1, B_2 and σ^2 so the posterior distribution coincides with the likelihood function. Our aim is to draw samples from the marginal posterior distribution of the parameters. As the model is non linear, the results on the conditional distributions of the regression coefficients shown in Chapter 1 do not apply and the MH algorithm is needed. We proceed in the following steps:

Step 1 Set starting values for $\Phi = \{B_1, B_2, \sigma^2\}$. These starting values could be set, for e.g, by estimating a log linearised version of equation 3.5 via OLS. The variance of the candidate generating density Σ can be set as the OLS coefficient covariance matrix $\hat{\Omega}$ times a scaling factor λ i.e $\Sigma = \hat{\Omega} \times \lambda$.Note that $\hat{\Omega}$ provides a rough guess of how volatile each parameter is. The scaling factor lets the researcher control the acceptance rate (a higher value for λ would mean a lower acceptance rate). Note that in this simple model the choice of starting values may not be very important and the algorithm would probably converge rapidly to the posterior. However, in the more complicated (and realistic) models considered below this choice can be quite important.

Step 2 Draw a new set of parameters from the random walk candidate density

$$\Phi^{New} = \Phi^{Old} + e \tag{3.6}$$

Step 3 Compute the acceptance probability $\alpha = \min\left(\frac{\pi\left(\Phi^{New}\right)}{\pi\left(\Phi^{Old}\right)}, 1\right)$. Note that the target density $\pi(.)$ is the likelihood function in this example. The log likelihood function for this regression model is given by

$$\ln L\left(Y_t \backslash \Phi\right) = -\frac{T}{2}\ln 2\pi - \frac{T}{2}\sigma^2 - 0.5\left(\frac{\left(Y_t - B_1\left(X_t^{B_2}\right)\right)'\left(Y_t - B_1\left(X_t^{B_2}\right)\right)}{\sigma^2}\right) \tag{3.7}$$

Therefore the acceptance probability is simply the likelihood ratio

$$\alpha = \min\left(\exp\left(\ln L\left(Y_t \backslash \Phi^{New}\right) - \ln L\left(Y_t \backslash \Phi^{Old}\right)\right), 1\right) \tag{3.8}$$

where $\ln L\left(Y_t \backslash \Phi^{New}\right)$ is the log likelihood evaluated at the new draw of B_1, B_2, σ^2 and $\ln L\left(Y_t \backslash \Phi^{Old}\right)$ is the log likelihood at the old draw. If $\alpha < u \sim U(0,1)$ we retain the new draw and set $\Phi^{Old} = \Phi^{New}$.

Step 4 Repeat steps 2 and 3 M times. The last L draws of B_1, B_2, σ^2 provide an approximation to the marginal posterior distributions.

```
1 clear
2 %generate artificial data
3 T=100;
4 sigma=1;
5 b1=4;
6 b2=2;
7 e=randn(T,1)*sqrt(sigma);
8 X=rand(T,1);
9 Y=b1.*(X.^b2)+e;
10 %step 1 set starting values
11 Gammaold=[0;0;0.1];
12 %step 2 set SIGMA matrix via OLS estimation
13 yols=Y;
14 xols=[ones(T,1) X];
15 bols=inv(xols'*xols)*(xols'*yols);
16 eols=yols-xols*bols;
17 sols=((eols'*eols)/T);
18 vols=sols*inv(xols'*xols);
19 P=eye(3);                    %this is the variance of the metropolis
hastings random walk based partly on OLS estimates
20 P(1,1)=(vols(1,1));
21 P(2,2)=(vols(2,2));
22 P(3,3)=0.1;
23 K=0.1;
24 P=K*P;   Σ = Ω̂ × λ
25 REPS=5000;
26 true=repmat([b1 b2 sigma],REPS,1);
27 % step 3 metropolis Hastings algorithm
28 out=[];
29 naccept=0;
30 for j=1:REPS
31     %step 3a draw new Gamma
32     Gammanew=Gammaold+(randn(1,3)*chol(P))';   Φ^New = Φ^Old + e
33     %step 3b evaluate posterior at new draw
34     B1=Gammanew(1);B2=Gammanew(2);sigma2=Gammanew(3);
35     if sigma2<0
36         posteriorNEW=-1000000;
37     else
38     B=[B1;B2];
            lnL(Y_t\Φ) = −(T/2) ln 2π − (T/2)σ² − 0.5( ((Y_t − B_1(X_t^{B_2}))'(Y_t − B_1(X_t^{B_2}))) / σ² )
39     lik=-(T/2)*log(2*(22/7))-(T/2)*sigma2-0.5*(((Y-B1.*(X.^B2))'*(Y-
B1.*(X.^B2)))/sigma2); %likelihood function
40     posteriorNEW=lik; %posterior at the new draw
41     end
42     %step 3c evaluate posterior at old draw
43     B1=Gammaold(1);B2=Gammaold(2);sigma2=Gammaold(3);
44     B=[B1;B2];
45     lik=-(T/2)*log(2*(22/7))-(T/2)*sigma2-0.5*(((Y-B1.*(X.^B2))'*(Y-
B1.*(X.^B2)))/sigma2); %likelihood function
46     posteriorOLD=lik; %posterior at the old draw
47
48     %step 3d compute acceptance probability
49     accept=min([exp(posteriorNEW-posteriorOLD);1]);   %min(accept,1)
α = min(exp(lnL(Y_t\Φ^New) − lnL(Y_t\Φ^Old)), 1)
50
51     u=rand(1,1);  %random number from the uniform dist
```

FIGURE 1. Matlab code for example1

Figures 1 and 2 show the matlab code for this example (example1.m). Lines 3 to 9 generate artificial data for the non-linear regression model assuming that $B_1 = 4, B_2 = 2, \sigma^2 = 1$. Line 11 sets the starting values for these parameters arbitrarily. Lines 20 to 22 set the variance of the random walk candidate density as $\Sigma = \begin{pmatrix} \sigma^{B_1} & 0 & 0 \\ 0 & \sigma^{B_2} & 0 \\ 0 & 0 & 0.1 \end{pmatrix} \times$scaling factor where σ^{B_1} and σ^{B_2} are OLS estimates of the variance of B_1 and B_2. Line 29 sets the variable naccept which will count the number of accepted draws. Hence the acceptance rate is naccept/REPS. Line 30 starts the loop for the

[2]See Chib and Ramamurthy (2010) for a more efficient version of the basic algorithm described above.

```
52
53      if u<accept
54          Gammaold=Gammanew;  %accept draw
55          naccept=naccept+1;  %count number of acceptances
56      end
57      out=[out;Gammaold'];
58
59  end
60
61  plot([out true])
62  xlabel('Metropolis Hastings Draws');
63  legend('B 1','B 2','\sigma^{2}','True B 1','True B 2','True
\sigma^{2}');
```

If $\alpha > u$ we set $\Phi^{Old} = \Phi^{New}$

Published with MATLAB® 7.9

FIGURE 2. Matlab code for example1 (continued)

MH algorithm. Line 32 draws the new value of the parameters from the random walk candidate density. Note that there is nothing intrinsic in this step that stops the new value of σ^2 from being less than zero. Therefore lines 35 to 37 set the value of the log likelihood to a very small number if a negative σ^2 is drawn thus ensuring that this draw is not going to be accepted. Alternatively one can set the acceptance probability to 0 when a negative value for σ^2 is drawn. Lines 44 to 46 calculate the log likelihood at the old draw. Line 49 calculates the acceptance probability. Line 53 checks if the acceptance probability is bigger than u a number from the standard uniform distribution. If this is the case we retain the new draw of the parameters.Figure 3 shows all the draws of the model parameters. The algorithm is close to the true values of these parameters after a few hundred draws.

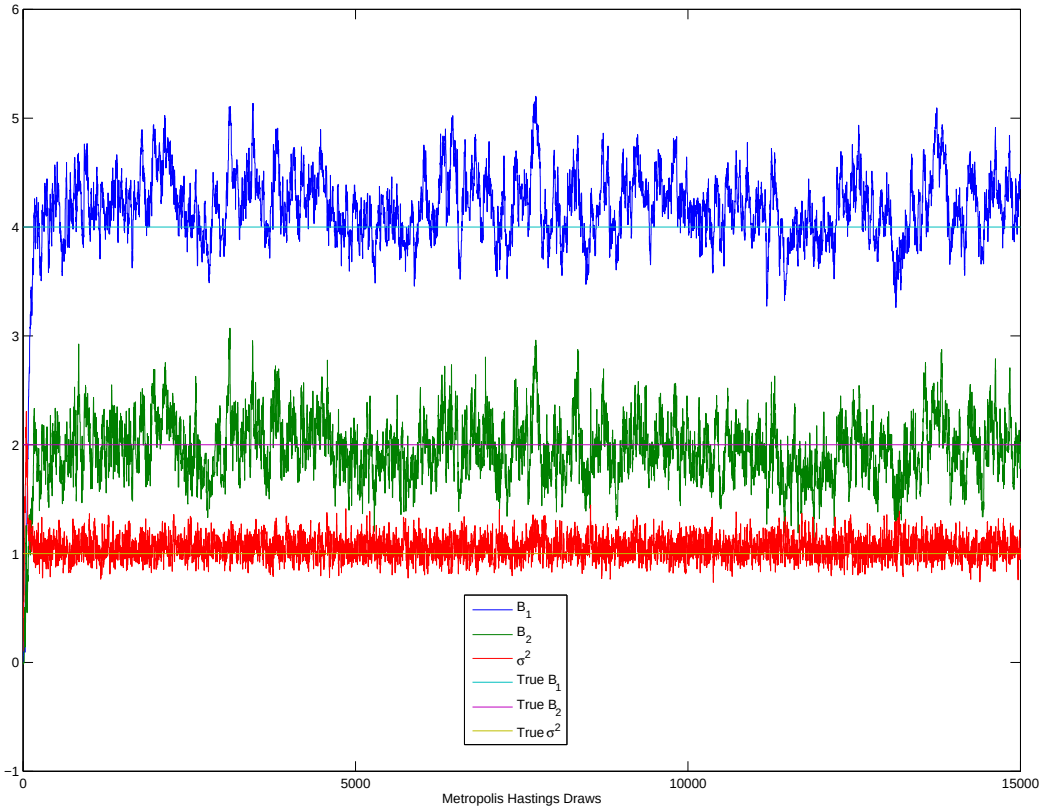

FIGURE 3. Draws of B_1, B_2, σ^2 using the MH algorithm in example 1

3.2. Estimating a non-linear regression via the random walk MH algorithm (incorporating prior distributions). We consider the same non-linear regression model examined in the previous section but now incorporate prior distribution for the regression parameters. We assume that the regression coefficients $B = \{B_1, B_2\}$ have a normal prior $p(B) \sim N(B_0, \Sigma_0)$. For convenience, we set a prior for the precision, the reciprocal of the variance. The Gamma prior $p\left(1/\sigma^2\right)$ with a prior scale parameter $\frac{\sigma^0}{2}$ and degrees of freedom $\frac{V^0}{2}$. The random walk MH algorithm now consists of the following steps:

algorithm is needed. We proceed in the following steps:

Step 1 Set the parameters of the prior distributions $p(B)$ and $p\left(1/\sigma^2\right)$. Set starting values for $\Phi = \{B_1, B_2, \sigma^2\}$. Finally set the variance of the candidate generating density Σ.

Step 2 Draw a new set of parameters from the random walk candidate density

$$\Phi^{New} = \Phi^{Old} + e \tag{3.9}$$

Step 3 Compute the acceptance probability $\alpha = \min\left(\frac{\pi\left(\Phi^{New}\right)}{\pi\left(\Phi^{Old}\right)}, 1\right)$. Note that the target density $\pi(.)$ is the *posterior density* in this example as we have prior distributions for our parameters. Recall from chapter 1 that the Bayes law states that he posterior distribution is proportional to the likelihood times the prior. Therefore we need to evaluate the likelihood and the prior distributions at the drawn value of the parameters and multiply them together. The log likelihood function for this regression model is given by

$$\ln L\left(Y_t \backslash \Phi\right) = -\frac{T}{2}\ln 2\pi - \frac{T}{2}\sigma^2 - 0.5\left(\frac{\left(Y_t - B_1\left(X_t^{B_2}\right)\right)'\left(Y_t - B_1\left(X_t^{B_2}\right)\right)}{\sigma^2}\right) \tag{3.10}$$

The prior density for the regression coefficients is just a normal density given by

$$P(B) = (2\pi)^{-K/2}\left|\Sigma_0\right|^{-\frac{1}{2}}\exp\left[-0.5\left(B - B_0\right)'\Sigma_0^{-1}\left(B - B_0\right)\right] \tag{3.11}$$

Note that this is evaluated at the new draw of the regression coefficients. If the new draw is very far from the prior mean B_0 and the prior is tight (the diagonal elements of Σ_0 are small) then $P(B)$ will evaluate to a small number.

Similarly, the log prior density for $1/\sigma^2$ is a Gamma distribution with a density function given by

$$P\left(1/\sigma^2\right) = C^* \frac{1}{\sigma^2}^{\frac{V_0}{2}-1} \exp\left(\frac{-\sigma_0}{2\sigma^2}\right) \tag{3.12}$$

where $C^* = \frac{1}{\Gamma\left(\frac{V_0}{2}\right)\left(\frac{2}{\sigma^0}\right)^{\frac{V_0}{2}}}$ and $\Gamma(.)$ denotes the Gamma function. The log posterior is given by

$$\ln H\left(\Phi \backslash Y_t\right) \propto \ln L\left(Y_t \backslash \Phi\right) + \ln P\left(B\right) + \ln P\left(1/\sigma^2\right)$$

Therefore the acceptance probability is simply the likelihood ratio

$$\alpha = \min\left(\exp\left(\ln H\left(\Phi^{New} \backslash Y_t\right) - \ln H\left(\Phi^{Old} \backslash Y_t\right)\right), 1\right) \tag{3.13}$$

where $\ln H\left(\Phi^{New} \backslash Y_t\right)$ is the log posterior evaluated at the new draw of B_1, B_2, σ^2 and $\ln H\left(\Phi^{Old} \backslash Y_t\right)$ is the log posterior evaluated at the old draw. If $\alpha < u \sim U(0,1)$ we retain the new draw and set $\Phi^{Old} = \Phi^{New}$.

Step 4 Repeat steps 2 and 3 M times. The last L draws of B_1, B_2, σ^2 provide an approximation to the marginal posterior distributions.

Figures 4 and 5 show the code for this example (example2.m). Relative to the previous example there are only two changes. First on lines 12 to 15 we set the parameters of the prior distributions for B and $1/\sigma^2$. Second, line 45 evaluates the log prior density for B at the new draw. Similarly, line 46 evaluates the log prior density for $1/\sigma^2$ at the new draw. The log posterior at the new draw is calculated on line 47. Lines 50 to 55 calculate the log posterior at the old draw of the parameters. The remaining code is identical to the previous example.

```
clear
%generate artificial data
T=100;
sigma=1;
b1=4;
b2=2;
e=randn(T,1)*sqrt(sigma);
X=rand(T,1);
Y=b1.*(X.^b2)+e;
%step 1 set starting values and priors
Gammaold=[0;0;0.1];
b0=zeros(2,1);
sigma0=eye(2)*100;  %P(b)~N(b0,sigma0)
s0=1;
v0=5;          %p(1/sigma)~Gamma(s0,v0)
%step 2 set SIGMA matrix via OLS estimation
yols=Y;
xols=[ones(T,1) X];
bols=inv(xols'*xols)*(xols'*yols);
eols=yols-xols*bols;
sols=((eols'*eols)/T);
vols=sols*inv(xols'*xols);
P=eye(3);                      %this is the variance of the metropolis
hastings random walk based partly on OLS estimates
P(1,1)=(vols(1,1));
P(2,2)=(vols(2,2));
P(3,3)=0.1;
K=0.2;
P=K*P;
REPS=15000;
true=repmat([b1 b2 sigma],REPS,1);
% step 3 metropolis Hastings algorithm
out=[];
naccept=0;
for j=1:REPS
    %step 3a draw new Gamma
    Gammanew=Gammaold+(randn(1,3)*chol(P))';
    %step 3b evaluate posterior at new draw
    B1=Gammanew(1);B2=Gammanew(2);sigma2=Gammanew(3);
    if sigma2<0
        posteriorNEW=-1000000;
    else
    B=[B1;B2];
    resid=Y-B1.*(X.^B2);
    lik=-(T/2)*log(2*pi*sigma2)-0.5*(((resid)'*(resid))/sigma2);
%likelihood function
```

$$P(B) = (2\pi)^{-K/2}|\Sigma_0|^{-\frac{1}{2}}\exp[-0.5(B-B_0)'\Sigma_0^{-1}(B-B_0)]$$

```
        normalprior=log(mvnpdf(B,b0,sigma0)); %evaluate prior for B1
and B2
```

$$P(1/\sigma^2) = C^*\frac{1}{\sigma^2}^{\frac{V_0}{2}-1}\exp\left(\frac{-\sigma_0}{2\sigma^2}\right)$$

```
    gammaprior=gampdf1(v0,s0,1/sigma2); %evaluate prior for 1/sigma
    posteriorNEW=lik+normalprior+gammaprior; %posterior at the new
draw
    end
    %step 3c evaluate posterior at old draw
    B1=Gammaold(1);B2=Gammaold(2);sigma2=Gammaold(3);
    B=[B1;B2];
    resid=Y-B1.*(X.^B2);
```

FIGURE 4. Matlab code for example 2

```
53     lik=-(T/2)*log(2*pi*sigma2)-0.5*(((resid)'*(resid))/sigma2);
%likelihood function
54         normalprior=log(mvnpdf(B,b0,sigma0)); %evaluate prior for B1
and B2
55     gammaprior=gampdf1(v0,s0,1/sigma2); %evaluate prior for 1/sigma
56     posteriorOLD=lik+normalprior+gammaprior; %posterior at the old
draw
57
58     %step 3d compute acceptance probability
59     accept=min([exp(posteriorNEW-posteriorOLD);1]);   %min(accept,1)
60
61     u=rand(1,1);  %random number from the uniform dist
62
63     if u<accept
64         Gammaold=Gammanew;  %accept draw
65         naccept=naccept+1;  %count number of acceptances
66     end
67     out=[out;Gammaold'];
68
69 end
70
71 plot([out true])
72 xlabel('Metropolis Hastings Draws');
73 legend('B_1','B_2','\sigma^{2}','True B_1','True B_2','True
\sigma^{2}');
```

Published with MATLAB® 7.9

FIGURE 5. Matlab code for example 2 (continued)

3.3. The random walk MH algorithm for a state space model. In this section we consider the estimation of a regression with time-varying parameters using the MH algorithm. Note that Gibbs sampling is feasible for this model. Our reason for using the MH algorithm is related to the fact the steps involved in dealing with this model are very similar to those required when estimating a DSGE model. In particular, the choice of starting values is no longer trivial.

The model we consider has the following state space representation

$$\begin{aligned} y_t &= c_t + b_t x_t + v_t, v_t \tilde{} N(0,R) \\ \begin{pmatrix} c_t \\ b_t \end{pmatrix} &= \begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix} + \begin{pmatrix} F_1 & 0 \\ 0 & F_2 \end{pmatrix} \begin{pmatrix} c_{t-1} \\ b_{t-1} \end{pmatrix} + \begin{pmatrix} e_{1t} \\ e_{2t} \end{pmatrix}, \\ & \begin{pmatrix} e_{1t} \\ e_{2t} \end{pmatrix} \tilde{} N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} Q_1 & 0 \\ 0 & Q_2 \end{pmatrix} \right) \end{aligned} \tag{3.14}$$

The random walk MH algorithm for this model works exactly as before. At each iteration we calculate the log posterior for the model at the old and the new draw of the parameters $\Phi = \{\mu_1, \mu_2, F_1, F_2, R, Q_1, Q_2\}$. Calculation of the posterior involves evaluating the prior distributions and the log likelihood of the model. Note that the likelihood function of this state space model is evaluated using the Kalman filter. As discussed in Hamilton (1994) (page 385) if the shocks to the state space model (v_t, e_{1t}, e_{2t}) are distributed normally, then the density of the data $f(y_t \backslash x_t)$ is given as

$$f(y_t \backslash x_t) = (2\pi)^{-1/2} \left| f_{t\backslash t-1} \right|^{-1/2} \times \exp\left(-0.5 \eta'_{t\backslash t-1} f^{-1}_{t\backslash t-1} \eta_{t\backslash t-1} \right) \tag{3.15}$$

for $t = 1....T$ with the log likelihood of the model given by

$$\ln f(y_t \backslash \Phi) = \sum_{t=1}^{T} \ln f(y_t \backslash x_t) \tag{3.16}$$

Here $\eta_{t\backslash t-1}$ is the prediction error from the prediction step of the Kalman filter and $f_{t\backslash t-1}$ is the variance of the prediction error (see Chapter 3).

Figures 6 and 7 show a matlab function (likelihoodTVP.m) which uses the Kalman filter to calculate the likelihood for this model and will be used in the main code discussed below. Line 4 checks if the variances (stored as the last three elements of theta) are positive and $F_1$and F_2 do not sum to a number greater than 1 (this is a rough way to check for stability). Lines 5 to 7 form the matrix $\begin{pmatrix} F_1 & 0 \\ 0 & F_2 \end{pmatrix}$ while lines 8 to 10 form the matrix $\begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}$. Line 13 specifies the matrix R while lines 14 to 16 specify the matrix $\begin{pmatrix} Q_1 & 0 \\ 0 & Q_2 \end{pmatrix}$. The Kalman filter recursions on lines 20 to 39 are as described in Chapter 3. Line 40 uses the prediction error and the variance of the prediction error to calculate $f(y_t \backslash x_t) = (2\pi)^{-1/2} \left| f_{t\backslash t-1} \right|^{-1/2} \times \exp\left(-0.5 \eta'_{t\backslash t-1} f^{-1}_{t\backslash t-1} \eta_{t\backslash t-1} \right)$. Line 42 adds this for each observation (if there are no numerical problems). Line 47 returns the negative of the likelihood function (we are going to minimise this below).

The MH algorithm for this model is given by the following steps:

Step 1 Set priors for the coefficients and variances of the state space model. We assume that μ_1, μ_2, F_1, F_2 have a normal prior while the reciprocal of R, Q_1, Q_2 have a Gamma prior.

Step 2 Set a starting value for the parameters $\Phi = \{\mu_1, \mu_2, F_1, F_2, R, Q_1, Q_2\}$ and the variance of the shock to the random walk candidate generating density. We set the starting value for Φ as the estimate Φ^{ML} by numerically maximising the log posterior. The mode of the posterior provides a reasonable point to the start the MH algorithm and implies that fewer iterations may be required for the algorithm to converge.[3]The estimate of the covariance of Φ^{ML} can be used to set the variance of the random walk candidate density. Note that the covariance of Φ^{ML} is given by the inverse of the hessian of the log posterior with respect to the model parameters. Denoting this estimated variance by $\hat{\Omega}$, the variance of the shock to the candidate generating density is set as $\Sigma = \hat{\Omega} \times \lambda$ where λ is a scaling factor chosen by the researchers such that the acceptance rate is between 20% and 40%.

Step 3 Draw a new set of parameters from the random walk candidate density

$$\Phi^{New} = \Phi^{Old} + e \tag{3.17}$$

Step 4 Compute the acceptance probability $\alpha = \min\left(\frac{\pi(\Phi^{New})}{\pi(\Phi^{Old})}, 1 \right)$. As in the previous example the target density is the posterior distribution. The log of the posterior distribution is calculated as the sum of the log likelihood and the sum of the log priors. As described above, the log likelihood is calculated by using the Kalman filter. If $\alpha > u \sim U(0,1)$ then we keep Φ^{New} otherwise we retain the old draw.

[3]Note also that if the posterior is multi-modal (which may be the case for complicated models) the numerical maximum will be a rough approximation to the posterior mode.

```
1 function [out,beta tt]=likelihoodTVP(theta,y,x)
2 %extract parameters of the state space
3 out=1000000000;
4 if sum(theta(5:end)<0)==0 && sum(abs(theta(1:2))>1)==0 only calculate
likelihood if variances are positive and F_1 + F_2 ≤ 1
5 F=zeros(2,2);
```

$$\begin{pmatrix} F_1 & 0 \\ 0 & F_2 \end{pmatrix}$$

```
6 F(1,1)=theta(1);
7 F(2,2)=theta(2);
8 mu=zeros(1,2);
```

$$\begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}$$

```
9 mu(1,1)=theta(3);
10 mu(1,2)=theta(4);
11
12
```

R

```
13 r=(theta(5));
```

$$\begin{pmatrix} Q_1 & 0 \\ 0 & Q_2 \end{pmatrix}$$

```
14 Q=zeros(2,2);
15 Q(1,1)=(theta(6));
16 Q(2,2)=(theta(7));
17 t=rows(y);
```

$$\ln f(y_t \backslash \Phi) = \sum_{t=1}^{T} \ln f(y_t \backslash x_t)$$

```
18 lik=0; will hold
19 %filter
20 beta0=zeros(1,2);
21 p00=eye(2);
22 beta tt=[];
23 ptt=zeros(t,2,2);
24 beta11=beta0;
25 p11=p00;
26 for i=1:t
27     H=x(i,:);
28     %Prediction
29 beta10=mu+beta11*F';
30 p10=F*p11*F'+Q;
31 yhat=(H*(beta10)')';
32 eta=y(i,:)-yhat;
33 feta=(H*p10*H')+r;
34 %updating
35 K=(p10*H')*inv(feta);
36 beta11=(beta10'+K*eta')';
37 p11=p10-K*(H*p10);
38 beta tt=[beta tt;beta11];
39 ptt(i,:,:)=p11;
```

FIGURE 6. The log likelihood for the time-varying parameter model in Matlab

Step 5 Repeat steps 3 and 4 M times. The last L draws of Φ provide an approximation to the marginal posterior distributions.

Figures 8, 9 and 10 show the code for the MH algorithm for this model. Line 2 of the code adds the optimization software *csminwel* written by Chris Sims and freely available at http://sims.princeton.edu/yftp/optimize/mfiles/. This matlab function minimises a user supplied function. Lines 5 to 23 create artificial data for the state space model assuming that $\mu_1 = 0.1, \mu_2 = -0.1, F_1 = 0.95, F_2 = 0.95, R = 2, Q_1 = 0.1, Q_2 = 0.1$. Lines 25 to 36 set the parameters for the prior distributions. Lines 37 to 39 maximise the log posterior of the model using *csminwel.* Line 39 called csminwel using the code:

$$f(y_t \backslash x_t) = (2\pi)^{-1/2} |f_{t\backslash t-1}|^{-1/2} \times \exp(-0.5 \eta'_{t\backslash t-1} f^{-1}_{t\backslash t-1} \eta_{t\backslash t-1})$$ in logs

```
40 liki=-0.5*log(2*pi)-0.5*log(det(feta))+(-0.5*(eta)*inv(feta)*(eta'));
41 if isreal(liki) & (1-isinf(liki))
42     lik=lik+liki;
43 else
44     lik=lik-10; if log f(y\x) cannot be computed set it equal to -10
45 end
46 end
47 out=-lik; return the negative of the likelihood function
48 if isnan(out)|| 1-isreal(out) || isinf(out)
49     out=1000000000; if log lik cannot be computed return a large
number
50 end
51 end
```

Published with MATLAB® 7.9

FIGURE 7. The log likelihood for the time-varying parameter model in Matlab (continued)

[FF,AA,gh,hess,itct,fcount,retcodeh] =
csminwel('posterior',theta0,eye(length(theta0))*.1,[],1e-15,1000,y,x,F0,VF0,MU0,VMU0,R0,VR0,Q0,VQ0);

The inputs to the function are (1) the name of the function that calculates the log posterior. This is called posterior.m in our example. Note that this example evaluates the log likelihood using likelihoodTVP.m. The function then evaluates the log prior for each parameter. The sum of these is the log joint prior. The sum of the log joint prior and the log likelihood is the log posterior. Note that posterior.m returns the negative of the log posterior. Therefore csminwel minimises the minimum of the negative log posterior which is equivalent to maximising the log posterior. (2) the initial values of the model parameters theta0. (3) the initial hessian matrix which can be left as default. (4) a

```
clear
addpath('sims Optimization'); includes function to minimize -logL written by Chris Sims
addpath('functions');
%create artificial data for time-varying paramteer model
T=300;
%generate artificial data on a time-varying parameter model
N=2;
Q=eye(N,N)*0.1;
R=2;
F(1,1)=0.95;
F(2,2)=0.95;
e=randn(T,2);
v=randn(T,1);
x=[randn(T,1) ones(T,1)];
y=zeros(T,1);
b=zeros(T,2);
MU=[0.1 -0.1];
for j=2:T
    b(j,:)=MU+b(j-1,:)*F'+e(j,:)*chol(Q);
    y(j,:)=x(j,:)*b(j,:)'+v(j,:)*sqrt(R);
end

TRUE=[diag(F);MU';R;diag(Q)];
%**********step 1 set priors for each parameter
% F~N(F0,VF0)
F0=ones(2,1);
VF0=eye(2)*2;
%MU~N(MU0,VMU0);
MU0=zeros(2,1);
VMU0=eye(2);
%R~IG(R0,VR0)
R0=1;
VR0=1;
%Q(i,i)~IG(Q0,VQ0)
Q0=0.1;
VQ0=1;
%**************step 2 estimate model via maximum likelihood
theta0=ones(7,1).*0.1;
[FF,AA,gh,hess,itct,fcount,retcodeh] = csminwel('posterior',theta0,eye(length(theta0))*.1,[],1e-
15,1000,y,x,F0,VF0,MU0,VMU0,R0,VR0,Q0,VQ0);
%**************step 2 set scale factor for the metropolis hastings
K=0.4;  %scaling factor
P=(chol(hess*K)); %compute variance of the random walk
Gammaold=AA; posterior mode estimates
REPS=50000;
BURN=30000;
naccept=0;
out1=zeros(REPS-BURN,7);
out2=zeros(REPS-BURN,1);
%compute posterior at old draw
          %compute -1*likelihood at old draw
```


FIGURE 8. Matlab code for the TVP model

```
        lik=likelihoodTVP(Gammaold,y,x);
        %evaluate prior for each set of parameters
        F=Gammaold(1:2);
        MU=Gammaold(3:4);
        R=Gammaold(5);
        Q=Gammaold(6:7);

        %prior for F
        Fprior=-log(det(VF0))-log(2*pi)/2-((F-F0)'*inv(VF0)*(F-F0));
        %prior for MU
        MUprior=-log(det(VMU0))-log(2*pi)/2-((MU-MU0)'*inv(VMU0)*(MU-MU0));
        %prior for R
         Rprior=log(2) - gammaln(VR0/2) + (VR0/2)*log(VR0*R0^2/2) - ( (VR0+1)/2 )*log(R^2) - VR0*R0^2/(2*(R)^2);
        %prior for Q
         Qprior=0;
        for i=1:2
         Qprior=Qprior+(log(2) - gammaln(VQ0/2) + (VQ0/2)*log(VQ0*Q0^2/2) - ( (VQ0+1)/2 )*log(Q(i)^2) - VQ0*Q0^2/(2*(Q(i))^2));
        end

        %joint prior is the sum of these
        priorold=Fprior+MUprior+Rprior+Qprior;
        posteriorOLD=-lik+priorold;posterior at old draw
        jj=1;
for j=1:REPS
    %step 1 draw new Gamma
    Gammanew=Gammaold+(randn(1,7)*P)';

    %step 2 check variances positive and elements of F sum to less than 1
    check=sum(Gammanew(5:end)<0) && sum(Gammanew(1:2)>1);
    if check
         posteriorNEW=-1000000;
    else
        %compute -1*likelihood at new draw
        lik=likelihoodTVP(Gammanew,y,x);
       F=Gammanew(1:2);
        MU=Gammanew(3:4);
        R=Gammanew(5);
        Q=Gammanew(6:7);

        %prior for F
        Fprior=-log(det(VF0))-log(2*pi)/2-((F-F0)'*inv(VF0)*(F-F0));
        %prior for MU
        MUprior=-log(det(VMU0))-log(2*pi)/2-((MU-MU0)'*inv(VMU0)*(MU-MU0));
        %prior for R
         Rprior=log(2) - gammaln(VR0/2) + (VR0/2)*log(VR0*R0^2/2) - ( (VR0+1)/2 )*log(R^2) - VR0*R0^2/(2*(R)^2);
        %prior for Q
         Qprior=0;
        for i=1:2
         Qprior=Qprior+(log(2) - gammaln(VQ0/2) + (VQ0/2)*log(VQ0*Q0^2/2) - ( (VQ0+1)/2 )*log(Q(i)^2) - VQ0*Q0^2/(2*(Q(i))^2));
         end

         %joint prior is the sum of these
```

FIGURE 9. Matlab code for the TVP model (continued)

```
103
104         priornew=Fprior+MUprior+Rprior+Qprior;
105         posteriorNEW=-lik+priornew;
106     end
107
108
109         accept=min([exp(posteriorNEW-posteriorOLD);1]);   %min(accept,1)
110
111     u=rand(1,1);  %random number from the uniform dist
112
113     if u<accept
114         Gammaold=Gammanew;  %accept draw
115         posteriorOLD=posteriorNEW; more efficient than calculating posterior at old draw all the time
116         naccept=naccept+1;  %count number of acceptances
117     end
118
119
120       ARATE=naccept/j; acceptance rate
121
122
123
124
125
126
127
128       if j>BURN
129       out1(jj,:)=Gammaold';
130       out2(jj,:)=posteriorOLD;
131       jj=jj+1;
132       end
133 end
134 subplot(3,3,1);
135 plot([out1(:,1) repmat(TRUE(1),size(out1,1),1)]);
136 title('F_{1}');
137 subplot(3,3,2);
138 plot([out1(:,2) repmat(TRUE(2),size(out1,1),1)]);
139 title('F_{2}');
140 subplot(3,3,3);
141 plot([out1(:,3) repmat(TRUE(3),size(out1,1),1)]);
142 title('\mu_{1}');
143 subplot(3,3,4);
144 plot([out1(:,4) repmat(TRUE(4),size(out1,1),1)]);
145 title('\mu_{2}');
146 subplot(3,3,5);
147 plot([out1(:,5) repmat(TRUE(5),size(out1,1),1)]);
148 title('R');
149 subplot(3,3,6);
150 plot([out1(:,6) repmat(TRUE(6),size(out1,1),1)]);
151 title('Q_{1}');
152 subplot(3,3,7);
153 plot([out1(:,7) repmat(TRUE(7),size(out1,1),1)]);
154 title('Q_{2}');
```

FIGURE 10. Matlab code for TVP model continued

function for calculating analytical derivitives. If this is unavailable then we enter an empty matrix [] as done above. (5) The tolerance level to stop the iterative procedure. This should be left as default. (6) The maximum number of iterations set to a 1000 in the example above. All the remaining arguments (y,x,F0,VF0,MU0,VMU0,R0,VR0,Q0,VQ0) are passed directly to the function posterior.m and are inputs for that function. The function returns (1) FF the value of the function at the minimum. (2) AA the value of the parameters at the minimum and (3) hess the inverse hessian of the function being minimised

Line 42 of the code sets the variance of the random walk candidate generating density as a scalar times the parameter variance obtained from the optimisation using csminwel. Line 43 sets the initial value of the parameters as the posterior mode estimates.

Line 51 calculates the log likelihood at the initial value of the parameters. Lines 59 to 68 evaluate the log prior distributions for the parameters of the state space model. Line 69 calculates the log joint prior as the sum of these prior distributions. Line 70 calculates the log posterior as the sum of the log likelihood and the log joint prior.

Line 74 draws the new value of the parameters from the random walk candidate generating density. Line 82 calculates the log likelihood at the new draw (assuming that the drawn variances are positive and the elements of F sum to less than 1). Lines 83 to 100 evaluate the log joint prior at the new draw and line 101 calculates the posterior. Line 109 calculates the acceptance probability. If this is bigger than a number from the standard uniform distribution then the new draw of the parameters is accepted. In this case Line 115 also sets posteriorOLD to posteriorNEW– it automatically updates the value of the posterior at the old draw eliminating the need to compute the posterior at the old draw at every iteration (as we have done in the examples above).

Line 120 computes the acceptance rate (the ratio of the number of accepted draws and the total draws). Once past the burn-in stage we save the draws of the model parameters. Figure 11 shows the retained draws of the parameters along with the true values.

3.4. The random walk MH algorithm used in a Threshold VAR model. In this section, we consider how the MH algorithm is used in the estimation of a Threshold VAR model (TVAR). The TVAR is defined as

$$\begin{aligned} Y_t &= c_1 + \sum_{j=1}^{P} \beta_1 Y_{t-j} + v_t, VAR(v_t) = \Omega_1 \text{ if } S_t \leq Y^* \\ Y_t &= c_2 + \sum_{j=1}^{P} \beta_2 Y_{t-j} + v_t, VAR(v_t) = \Omega_2 \text{ if } S_t > Y^* \end{aligned}$$

where Y_t is a matrix of endogenous variables, $S_t = Y_{j,t-d}$ (i.e. a lag of one of the endogenous variables) is the threshold variable and Y^* is the threshold level. Note that if Y^* and d are known, then the TVAR is simply two

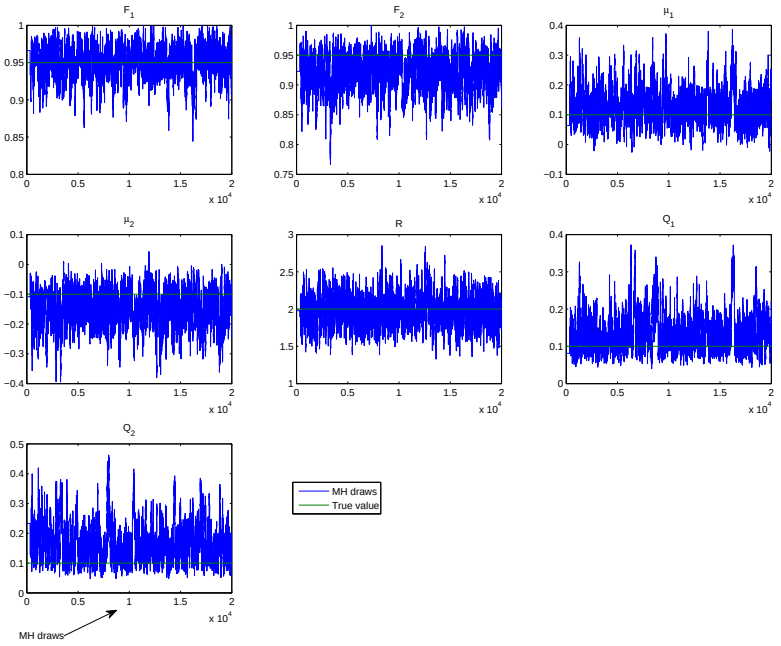

FIGURE 11. MH draws for the TVP model

VAR models defined over the appropriate data samples using $S_t \leq Y^*$ and $S_t > Y^*$. This observation allows us to devise a Gibbs algorithm (with a MH step). In what follows below we assume the delay parameter d to be known. See Chen and Lee (1995) for the extension of the algorithm to the case where d is estimated.

Step 1 Set Priors. In the application below, we assume $p(Y^*) \tilde{} N(\bar{Y}^*, \sigma_{Y^*})$. We set a natural conjugate prior for the VAR parameters in both regimes using dummy observations. See the prior used in section 6. Set an initial value for Y^*. One way to do this is to use the mean or median of $Y_{i,t-d}$.

Step 2 Seperate the data into two regimes. The first regime includes all observations such that $S_t \leq Y^*$. Call this sample $Y_{1,t}$. The second regime includes all observations such that $S_t > Y^*$. Call this sample $Y_{2,t}$.

Step 3 Sample the VAR parameters $b_i = \{c_i, \beta_i\}$ and Ω_i in each regime $i = 1, 2$. Let X denote the right hand side variables of the VAR. The conditional distribution is exactly as defined in chapter 2 above and is given by

$$\begin{aligned} H(b_i \backslash \Omega_i, Y_t, Y^*) &\tilde{} N(vec(B_i^*), \Omega_i \otimes (X_i^{*\prime} X_i^*)^{-1}) \\ H(\Omega_i \backslash b_i, Y_t, , Y^*) &\tilde{} IW(S_i^*, T_i^*) \end{aligned} \tag{3.18}$$

where

$$\begin{aligned} B_i^* &= (X_i^{*\prime} X_i^*)^{-1} (X_i^{*\prime} y_i^*) \\ S_i^* &= (y_i^* - X_i^* b)' (y_i^* - X_i^* b_i) \end{aligned}$$

where $y_i^* = [Y_{i,t}; Y_D]$ and $X_i^* = [X_{i,t}; X_D]$ with Y_D, X_D the dummy observations that define the prior for the left and the right hand side of the VAR respectively.

Step 4 Use a MH step to sample Y^*. Draw a new value of the threshold from the random walk

$$Y_{new}^* = Y_{old}^* + e, e \tilde{} N(0, \Sigma)$$

Then compute the acceptance probability

$$\alpha = \frac{F(Y \backslash b_i, \Omega_i, Y_{new}^*) p(Y_{new}^*)}{F(Y \backslash b_i, \Omega_i, Y_{old}^*) p(Y_{old}^*)}$$

where $F(Y \backslash b_i, \Omega_i, Y_{new}^*)$ is the likelihood of the VAR computed as the product of the likelihoods in the two regimes. The log likelihood in each regime (ignoring constants) is

$$\left(\frac{T}{2}\right) \log |\Omega_i^{-1}| - 0.5 \sum_{t=1}^{T} \left[\left(Y_{i,t} - X_{i,t} \tilde{b}_i\right)' \Omega_i^{-1} \left(Y_{i,t} - X_{i,t} \tilde{b}_i\right) \right]$$

with $\tilde{b}_i$ equivalent to b_i reshaped to be conformable with $X_{i,t}$. Then draw $u \tilde{} U(0,1)$. If $u < \alpha$ accept Y_{new}^* else retain Y_{old}^*. The scale Σ can be tuned to ensure an acceptance rate between 20% and 40%.

As an example we consider a TVAR where Y contains US data on GDP growth, CPI inflation, a short term interest rate and a financial conditions index (FCI) calculated by the Chicago Fed. The threshold variable is assumed to be the second lag of FCI and examine the impulse response of the variables to a unit increase in FCI (a deterioration

mytemp

```
1 clear
2 addpath('functions');
3 data=xlsread('\data\tvardata2.xlsx');
4 L=2;  %lag length
5 tard=2;  %delay
6 tarvar=4;  % threshold variable is the column number tarvar in data
7 MaxTrys=1000;
8 tarscale=0.1;  %scaling parameter for RW Metropolis algorithm
9 REPS=10000;
10 BURN=8000;
11 HORZ=40;  %impulse response horizon
12 %prepare data
13 Y=data;
14 N=cols(Y);
15 ncrit=(N*L+1);
16 %take lags
17 X=[];
18 for j=1:L
19 X=[X lag0(data,j) ];
20 end
21 X=[X ones(rows(X),1)];
22 %compute threshold variable
23 Ystar=lag0(Y(:,tarvar),tard);  Ystar is threshold variable
24 Y=Y(max([L,tard(1)])+1:end,:);
25 X=X(max([L,tard(1)])+1:end,:);
26 Ystar=Ystar(max([L,tard(1)])+1:end,:);
27 tarmean=mean(Ystar);  %mean of the prior on the threshold is the mean value of the threshold
variable
28 tarvariance=10;   %variance of the prior
29 % Additional priors for VAR coefficients
30 lamdaP  = 1;
31 tauP    = 10*lamdaP;
32 epsilonP= 1/10000;
33 muP=mean(Y)';
34 sigmaP=[];
35 deltaP=[];
36 e0=[];
37 for i=1:N
38     ytemp=Y(:,i);
39     xtemp=[lag0(ytemp,1) ones(rows(ytemp),1)];
40     ytemp=ytemp(2:end,:);
41     xtemp=xtemp(2:end,:);
42     btemp=xtemp\ytemp;
43     etemp=ytemp-xtemp*btemp;
44     stemp=etemp'*etemp/rows(ytemp);
45     if abs(btemp(1))>1
46         btemp(1)=1;
47     end
48     deltaP=[deltaP;btemp(1)];
49     sigmaP=[sigmaP;stemp];
50     e0=[e0 etemp];
51 end
52 %dummy data to implement priors see http://ideas.repec.org/p/ecb/ecbwps/20080966.html
53 [yd,xd] = create_dummies(lamdaP,tauP,deltaP,epsilonP,L,muP,sigmaP,N);
54 T=rows(Y);
55 %append
56   Y0=[Y;yd];
57   X0=[X;xd];
58
```

mytemp.html[15/07/2014 13:49:15]

FIGURE 12. Matlab code for TVAR model

of financial conditions) in the two regimes. The matlab code is in the file named thresholdvarNFCI.m and displayed in figures 12, 13 and 14. In this example the prior $p(Y^*) \tilde{} N(\bar{Y}^*, \sigma_{Y^*})$ is set by using the mean of NFCI as $\bar{Y}^*$ and $\sigma_{Y^*} = 10$ (line 28). Lines 30 to 53 set the natural conjugate prior for the VAR parameters. Lines 80 to 87 seperate the sample into two regimes. Lines 89 to 128 draw the VAR coefficients and covariance in each regime. The MH step to draw the threshold variable starts on line 134 with a draw from the random walk candidate density. Then the log posterior $\ln\left(F\left(Y \backslash b_i, \Omega_i, Y^*_{new}\right) p\left(Y^*_{new}\right)\right)$ is computed on line 136 while $\ln\left(F\left(Y \backslash b_i, \Omega_i, Y^*_{old}\right) p\left(Y^*_{old}\right)\right)$ is computed on line 137. The acceptance probability is computed on line 138.

mytemp

```
sigma1=eye(N); %starting value for sigma
sigma2=eye(N);
beta0=vec(X0\Y0);
beta01=beta0;
beta02=beta0;
tar=tarmean; %initial value of the threshold
tarold=tar;
naccept=0;




%gibbs algorithm
jgibbs=1;
for igibbs=1:REPS




%step 1: Seperate into two regimes
  e1=Ystar<=tar;
  e2=Ystar>tar;

  Y1=Y(e1,:);
  X1=X(e1,:);

  Y2=Y(e2,:);
  X2=X(e2,:);

  %step 2 Sample Coefficients and variance regime 1

  Y0=[Y1;yd];
  X0=[X1;xd];
%conditional mean of the VAR coefficients
mstar1=vec(X0\Y0);  %ols on the appended data
xx=X0'*X0;
ixx1=xx\eye(cols(xx));
 [ beta1,PROBLEM1] = getcoef( mstar1,sigma1,ixx1,MaxTrys,N,L );
  if PROBLEM1
      beta1=beta01;
   else
      beta01=beta1;
   end

   %draw covariance
   e=Y0-X0*reshape(beta1,N*L+1,N);
  scale=e'*e;
  sigma1=iwpQ(rows(Y0),inv(scale));


  %step 3 Sample Coefficients and variance in regime 2

   Y0=[Y2;yd];
  X0=[X2;xd];
 %conditional mean of the VAR coefficients
 mstar2=vec(X0\Y0);  %ols on the appended data
 xx=X0'*X0;
 ixx2=xx\eye(cols(xx));
  [ beta2,PROBLEM2] = getcoef( mstar2,sigma2,ixx2,MaxTrys,N,L );
```

mytemp.html[15/07/2014 13:49:15]

FIGURE 13. Matlab code for TVAR model

mytemp

```
    if PROBLEM2
        beta2=beta02;
    else
        beta02=beta2;
    end

    %draw covariance
    e=Y0-X0*reshape(beta2,N*L+1,N);
   scale=e'*e;
   sigma2=iwpQ(rows(Y0),inv(scale));



   %step 4 Sample Threshold via a Random Walk Metropolis Step

   tarnew=tarold+randn(1,1)*sqrt(tarscale);  Candidate draw
        getvarpost.m computes log posterior used for acceptance probability
     postnew=getvarpost(Y,X,beta1,beta2,sigma1,sigma2,L,tarnew,tarmean,tarvariance,Ystar,ncrit);
     postold=getvarpost(Y,X,beta1,beta2,sigma1,sigma2,L,tarold,tarmean,tarvariance,Ystar,ncrit);
    accept=exp(postnew-postold);
    u=rand(1,1);
    if u<accept
        tarold=tarnew;
        naccept=naccept+1;
    end
    tar=tarold;
    arate=naccept/igibbs;
if igibbs>100 && igibbs<1100
       if arate<0.2
           tarscale=tarscale*0.99;
       elseif arate>0.4
           tarscale=tarscale*1.01;
       end
end

disp(sprintf(' Replication %s of %s acceptance %s. ', ...
           num2str(igibbs), num2str(REPS),num2str(arate) ));

       if igibbs>BURN
           %impulse response analysis
           A01=chol(sigma1);
           A02=chol(sigma2);
           irf1=irfsim(reshape(beta1,N*L+1,N),N,L,A01,[0 0 0 1],HORZ);
           irf2=irfsim(reshape(beta2,N*L+1,N),N,L,A02,[0 0 0 1],HORZ);
           irf1=irf1./irf1(1,4);
           irf2=irf2./irf2(1,4);

           %save results
           irfmat1(jgibbs,:,:)=irf1;
           irfmat2(jgibbs,:,:)=irf2;
           smat(:,jgibbs)=e2;
           jgibbs=jgibbs+1;
       end



end


figure(1)
```

mytemp.html[15/07/2014 13:49:15]

FIGURE 14. Matlab code for TVAR model

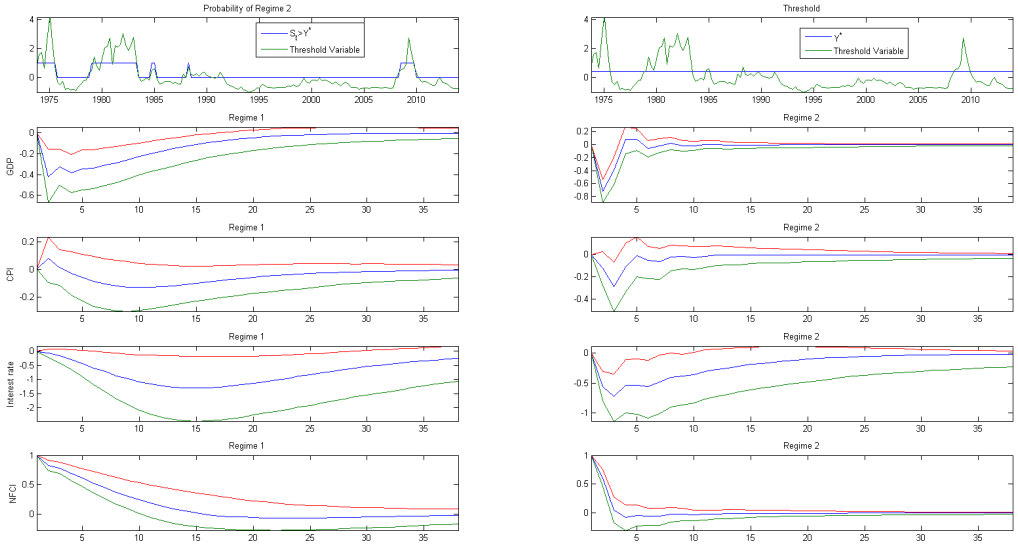

FIGURE 15. Results from the TVAR model for the US

The top right panel of figure 15 plots the estimated threshold and the threshold variable and shows that regime 2 persisted in the 1980s, the early 1990s and then during the recent recession. There is some evidence that the negative impact of an FCI shock on GDP growth is larger in regime 2.

4. The independence Metropolis Hastings algorithm

The independence MH algorithm differs from the random walk MH algorithm in that the candidate generating density is not specified as a random walk. Therefore, the new draw of the parameters does not depend directly on the previous draw. The candidate density is specified as

$$q\left(\Phi^{G+1} \backslash \Phi^{G}\right)=q\left(\Phi^{G+1}\right) \tag{4.1}$$

Note that now, in general, the formula for the acceptance probability does not simplify and is given by

$$\alpha=\min\left(\frac{\pi\left(\Phi^{G+1}\right)/q\left(\Phi^{G+1} \backslash \Phi^{G}\right)}{\pi\left(\Phi^{G}\right)/q\left(\Phi^{G} \backslash \Phi^{G+1}\right)},1\right) \tag{4.2}$$

The independence MH algorithm is therefore more general than the random walk MH algorithm. Unlike the random walk MH algorithm, the candidate generating density in the independence MH algorithm has to be tailored to the particular problem at hand. We examine an application to stochastic volatility models below.

Apart from the change in the form of the candidate generating density the steps of the algorithm remain the same:

Step 1 Set starting values for the model parameters.
Step 2 Draw a candidate value of the parameters $\Phi^{G+1}$from the candidate generating density $q\left(\Phi^{G+1}\right)$
Step 3 Compute the acceptance probability

$$\alpha=\min\left(\frac{\pi\left(\Phi^{G+1}\right)/q\left(\Phi^{G+1}/\Phi^{G}\right)}{\pi\left(\Phi^{G}\right)/q\left(\Phi^{G}/\Phi^{G+1}\right)},1\right) \tag{4.3}$$

Step 4 If $u \sim U(0,1)$ is less than α retain Φ^{G+1}. Otherwise retain the old draw.
Step 5 Repeat steps 2 to 4 M times and base inference on the last L draws. In other words, the empirical distribution using the last L draws is an approximation to target density.

4.1. Estimation of stochastic volatility models via the independence MH algorithm. A simple stochastic volatility model for a $T \times 1$ data series y_t is given by

$$\begin{aligned} y_t &= \varepsilon_t \sqrt{\exp\left(\ln h_t\right)} \\ \ln h_t &= \ln h_{t-1}+v_t \\ & v_t \tilde{} N(0,g) \end{aligned} \tag{4.4}$$

where h_t is time-varying variance. Note that this is a state space model where the observation equation is non-linear in the state variable h_t and therefore the Carter and Kohn algorithm does not apply. Jacquier *et al.* (2004) instead suggest applying an independence MH algorithm at each point in time to sample from the conditional distribution of h_t which is given by $f(h_t \backslash h_{-t}, y_t)$ where the subscript $-t$ denotes all other dates than t.Jacquier *et al.* (2004) argue that because the transition equation of the model is a random walk, the knowledge of h_{t+1} and h_{t-1} contains all relavant information about h_t. Therefore, the conditional distribution of h_t can be simplified as

$$f(h_t \backslash h_{-t}, y_t) = f(h_t \backslash h_{t-1}, h_{t+1}, y_t) \tag{4.5}$$

Jacquier *et al.* (2004) show that this density has the following form

$$f(h_t \backslash h_{t-1}, h_{t+1}, y_t) = h_t^{-0.5} \exp\left(\frac{-y_t^2}{2h_t}\right) \times h_t^{-1} \exp\left(\frac{-(\ln h_t - \mu)^2}{2\sigma_h}\right) \tag{4.6}$$

with

$$\mu = \frac{(\ln h_{t+1} + \ln h_{t-1})}{2} \tag{4.7}$$

$$\sigma_h = \frac{g}{2} \tag{4.8}$$

That is $f(h_t \backslash h_{t-1}, h_{t+1}, y_t)$ is a product of a normal density $h_t^{-0.5} \exp\left(\frac{-y_t^2}{2h_t}\right)$ and a log normal density $h_t^{-1} \exp\left(\frac{-(\ln h_t - \mu)^2}{2\sigma_h}\right)$.

To sample from $f(h_t \backslash h_{t-1}, h_{t+1}, y_t)$, Jacquier *et al.* (2004) suggest a date by date application of the independence MH algorithm with the candidate density defined as the second term in equation 4.6

$$q\left(\Phi^{G+1}\right) = h_t^{-1} \exp\left(\frac{-(\ln h_t - \mu)^2}{2\sigma_h}\right) \tag{4.9}$$

The acceptance probability in this case is given by

$$\alpha = \min\left(\frac{\pi\left(\Phi^{G+1}\right)/q\left(\Phi^{G+1}/\Phi^{G}\right)}{\pi\left(\Phi^{G}\right)/q\left(\Phi^{G}/\Phi^{G+1}\right)}, 1\right) \tag{4.10}$$

$\rightarrow$

$$\alpha = \frac{\left[h_{t,new}^{-0.5} \exp\left(\frac{-y_t^2}{2h_{t,new}}\right) \times h_{t,new}^{-1} \exp\left(\frac{-(\ln h_{t,new} - \mu)^2}{2\sigma_h}\right)\right] / h_{t,new}^{-1} \exp\left(\frac{-(\ln h_{t,new} - \mu)^2}{2\sigma_h}\right)}{\left[h_{t,old}^{-0.5} \exp\left(\frac{-y_t^2}{2h_{t,old}}\right) \times h_{t,old}^{-1} \exp\left(\frac{-(\ln h_{t,old} - \mu)^2}{2\sigma_h}\right)\right] / h_{t,old}^{-1} \exp\left(\frac{-(\ln h_{t,old} - \mu)^2}{2\sigma_h}\right)} \tag{4.11}$$

where the subscript *new* denotes the new draw and the subscript *old* denotes the old draw. Equation 4.11 simplifies to give

$$\alpha = \frac{h_{t,new}^{-0.5} \exp\left(\frac{-y_t^2}{2h_{t,new}}\right)}{h_{t,old}^{-0.5} \exp\left(\frac{-y_t^2}{2h_{t,old}}\right)} \tag{4.12}$$

Therefore, for each t one generates a value of h_t using the candidate density in equation 4.9 and then calculates the acceptance probability using equation 4.12. Note however that this algorithm is not operational for the first and the last date in the sample as the calculation of $\mu = \frac{(\ln h_{t+1} + \ln h_{t-1})}{2}$ requires knowledge of h_{t+1} and h_{t-1}.

Jacquier *et al.* (2004) suggest sampling the initial value of h_t denoted by h_0 using the following procedure. Starting with the following prior for $\ln h_0 \tilde{} N(\bar{\mu}, \bar{\sigma})$ Jacquier *et al.* (2004) show that the posterior for $\ln h_0$ is given by

$$f(h_0 \backslash h_1) = h_0^{-1} \exp\left(\frac{-(\ln h_0 - \mu_0)^2}{2\sigma_0}\right) \tag{4.13}$$

where

$$\begin{aligned} \sigma_0 &= \frac{\bar{\sigma} g}{\bar{\sigma} + g} \\ \mu_0 &= \sigma_0 \left(\frac{\bar{\mu}}{\bar{\sigma}} + \frac{\ln h_1}{g}\right) \end{aligned}$$

Therefore the algorithm starts by sampling h_0 from equation 4.13 and accepting the draw (as the data for this observation y_0 is not defined).

Jacquier *et al.* (2004) suggest sampling the final value of h_t (with $t = T$)using the following modified candidate generating density

$$q\left(\Phi^{G+1}\right) = h_t^{-1} \exp\left(\frac{-(\ln h_t - \mu)^2}{2\sigma_h}\right) \tag{4.14}$$

where

$$\begin{aligned} \mu &= \ln h_{t-1} \\ \sigma_h &= g \end{aligned} \tag{4.15}$$

The algorithm for the stochastic volatility model consists of the following steps[4]:

Step 1 Obtain a starting value for $h_t, t = 0....T$ as $\hat{\varepsilon}_t^2$ and set the prior $\bar{\mu}, \bar{\sigma}$ (e.g $\bar{\mu}$ could be the log of OLS estimate of the variance of ε_t and $\bar{\sigma}$ could be set to a big number to reflect the uncertainty in this initial guess). Set an inverse Gamma prior for g i.e. $p(g) \sim IG\left(\frac{g_0}{2}, \frac{v_0}{2}\right)$ Set a starting value for g.

Step 2 Time 0 Sample the initial value of h_t denoted by h_0 from the log normal density

$$f\left(h_0 \backslash h_1\right) = h_0^{-1} \exp\left(\frac{-\left(\ln h_0 - \mu_0\right)^2}{2\sigma_0}\right)$$

where the mean $\mu_0 = \sigma_0\left(\frac{\bar{\mu}}{\bar{\sigma}} + \frac{\ln h_1}{g}\right)$ and $\sigma_0 = \frac{\bar{\sigma} g}{\bar{\sigma}+g}$.

ALGORITHM 4. *To sample from the log normal density* $z \sim \log normal\left(\mu, \sigma\right)$ *sample* z_0 *from the normal density* $N\left(\mu, \sigma\right)$. *Then* $z = \exp\left(z_0\right)$.

ep 2 Time 1 to T-1 For *each date* t=1 to T-1 draw a new value for h_t from the candidate density (call the draw $h_{t,new}$)

$$q\left(\Phi^{G+1}\right) = h_t^{-1} \exp\left(\frac{-\left(\ln h_t - \mu\right)^2}{2\sigma_h}\right)$$

where $\mu = \frac{\left(\ln h_{t+1} + \ln h_{t-1}\right)}{2}$ and $\sigma_h = \frac{g}{2}$. Compute the acceptance probability

$$\alpha = \min\left(\frac{h_{t,new}^{-0.5} \exp\left(\frac{-y_t^2}{2h_{t,new}}\right)}{h_{t,old}^{-0.5} \exp\left(\frac{-y_t^2}{2h_{t,old}}\right)}, 1\right)$$

Draw $u \tilde{} U(0,1)$. If $u < \alpha$ set $\boldsymbol{h}_t = h_{t,new}$. Otherwise retain the old draw.

Step 2 Time T For the last time period $t = T$ compute $\mu = \ln h_{t-1}$ and $\sigma_h = g$ and draw $h_{t,new}$ from the candidate density

$$q\left(\Phi^{G+1}\right) = h_t^{-1} \exp\left(\frac{-\left(\ln h_t - \mu\right)^2}{2\sigma_h}\right)$$

Compute the acceptance probability

$$\alpha = \min\left(\frac{h_{t,new}^{-0.5} \exp\left(\frac{-y_t^2}{2h_{t,new}}\right)}{h_{t,old}^{-0.5} \exp\left(\frac{-y_t^2}{2h_{t,old}}\right)}, 1\right)$$

Draw $u \tilde{} U(0,1)$. If $u < \alpha$ set $\boldsymbol{h}_t = h_{t,new}$. Otherwise retain the old draw.

Step 3 Given a draw for h_t compute the residuals of the transition equation $v_t = \ln h_t - \ln h_{t-1}$. Draw g from the inverse Gamma distribution with scale parameter $\frac{v_t' v_t + g_0}{2}$ and degrees of freedom $\frac{T+v_0}{2}$. Note that this is an example of a combination of Metropolis and Gibbs sampling algorithms.

Step 4 Repeat steps 2 and 3 M times. The last L draws of h_t and g provide an approximation to the marginal posterior distributions.

Figures 16, 17 and 18 present the Matlab code for the stochastic volatility model applied to annual UK inflation over the period 1914q1 to 2011q4 (example4.m). Lines 14 and 15 of the code set the prior for g. Lines 16 and 17 set the prior $\ln h_0 \tilde{} N\left(\bar{\mu}, \bar{\sigma}\right)$ where $\bar{\mu}$ is set equal to the log of the variance of the first 10 observations in the sample. Line 23 calculates a rough starting value for h_t as the square of the first difference of y_t. Lines 35 and 36 calculate σ_0 and μ_0 and line 38 draws h_0 from the log normal density. Line 41 starts a loop from period 1 to T-1. Note that line 42 selects h_{t+1} as the lead value of h_t using the last draw of h_t. Line 47 and 48 calculate the mean and variance of the candidate density and line 49 draws the candidiate value of h_t. Line 54 calculates the acceptance probability in logs. Lines 68 to 84 repeat this for the final observation in the sample period. Line 84 calculate the residuals of the transition equation as $v_t = \ln h_t - \ln h_{t-1}$. Line 85 draws g from the inverse Gamma distribution.

[4]Note that the Jacquier *et al.* (2004) algorithm is a single-move algorithm– the stochastic volatility is drawn one period at a time. This may mean that this algorithm requires a large number of draws before convergence occurs. Kim *et al.* (1998) develop an algorithm to sample the entire time-series of the stochastic volatility jointly and show that this multi-move algorithm is more efficient.

```
%a stochastic volatility model for UK inflation
clear
addpath('functions');
%oad inflation data
Y=xlsread('\data\inflation.xlsx');
Y=((log(Y)-log(lag0(Y,4))))*100;
Y=Y(5:end,:);
T=rows(Y);
TT0=10;  %training sample

%Independence metropolis hastings algorithm for svol model
%step 1 priors for g~IG(V0,T0) and initial conditions for the
stochastic
%volatility
```

$$p(g) \sim IG(g_0, v_0)$$

```
V0=0.01;  %prior scale
T0=1;       %prior degrees of freedom
```

$\bar{\mu}$

```
mubar=log(std(Y(1:TT0))^2);
```

$\bar{\sigma}$

```
sigmabar=10;
%remove training sample
Y=Y(TT0+1:end,:);
T=rows(Y);
%step 2 starting values for stochastic volatility and
```

Obtain a starting value for $h_t, t = 0 \ldots T$ as $\hat{\varepsilon}_t^2$

```
hlast=diff(Y).^2;
hlast=[hlast(1:2);hlast]+0.0001;  %small number added to ensure no
value is zero
g=1;
REPS=30000;
BURN=25000;
out=[];
for j=1:REPS
%step 3 data by date metropolis hastings algorithm to draw the
stochastic
%volatility
hnew=zeros(T+1,1);
i=1;
%time period 0
```

h_1

```
hlead=hlast(i+1);
```

$$\sigma_0 = \frac{\bar{\sigma} g}{\bar{\sigma} + g}$$

```
ss = sigmabar*g/(g + sigmabar);   %variance
mu = ss*(mubar/sigmabar + log(hlead)/g);  %mean
```

$$\mu_0 = \sigma_0 \left(\frac{\bar{\mu}}{\bar{\sigma}} + \frac{lnh_1}{g} \right)$$

```
%draw from lognormal  using mu and ss
h = exp(mu + (ss^.5)*randn(1,1));
```

Sample the initial value of h_t denoted by h_0 from the log normal density

```
hnew(i)=h;
%time period 1 to t-1
for i=2:T
```

h_{t+1}

```
    hlead=hlast(i+1);
```

FIGURE 16. Matlab code for the stochastic volatility model

h_{t-1}

```
43     hlag=hnew(i-1);
44     yt=Y(i-1);
45
46 %mean and variance of the proposal log normal density
```

$$\mu = \frac{(\ln h_{t+1} + \ln h_{t-1})}{2}$$

```
47 mu = (log(hlead)+log(hlag))/2;
```

$$\sigma_h = \frac{g}{2}$$

```
48 ss = g/2;
49 %candidate draw from lognormal
50 htrial = exp(mu + (ss^.5)*randn(1,1));
51 %acceptance probability in logs
52 lp1 = -0.5*log(htrial) - (yt^2)/(2*htrial);  %numerator
53 lp0 = -0.5*log(hlast(i)) - (yt^2)/(2*hlast(i));   %denominator
54 accept = min([1;exp(lp1 - lp0)]);  %ensure accept<=1
```

$$\alpha = \min\left(\frac{h_{t,new}^{-0.5}\exp\left(\frac{-y_t^2}{2h_{t,new}}\right)}{h_{t,old}^{-0.5}\exp\left(\frac{-y_t^2}{2h_{t,old}}\right)}, 1\right)$$

```
55 u = rand(1,1);
56 if u <= accept;
57    h = htrial;
58 else
59    h = hlast(i);
60 end
61 hnew(i)=h;
62 end
63 %time period T
64 i=T+1;
65 yt=Y(i-1);
66 hlag=hnew(i-1);
67 %mean and variance of the proposal density
```

$$\mu = (\ln h_{t-1})$$

```
68 mu = log(hlag);   % only have ht-1
```

$$\sigma_h = g$$

```
69 ss = g;
70 %candidate draw from lognormal
71 htrial = exp(mu + (ss^.5)*randn(1,1));
72 %acceptance probability
73 lp1 = -0.5*log(htrial) - (yt^2)/(2*htrial);
74 lp0 = -0.5*log(hlast(i)) - (yt^2)/(2*hlast(i));
75 accept = min([1;exp(lp1 - lp0)]);  %ensure accept<=1
76 u = rand(1,1);
77 if u <= accept;
78    h = htrial;
79 else
80    h = hlast(i);
81 end
82 hnew(i)=h;
83 %step 4 draw g from the inverse Gamma distribution
84 errors=diff(log(hnew));
```

Draw g from the inverse Gamma distribution with scale parameter $v_t' v_t + g_0$ and degrees of freedom $T + v_0$

```
85 g=IG(T0,V0,errors);  %draw from the inverse Gamma distribution
86 %step 5 update vale of H
87 hlast=hnew;
```

FIGURE 17. Matlab code for the stochastic volatility model continued

```
88 %save
89 if j>BURN
90 out=[out hlast];
91 end
92 end
93 TT=1917.5:0.25:2011;
94 subplot(1,2,1);
95 plot(TT(1:end),Y);
96 title('Annual CPI inflation for the UK');
97 axis tight
98 subplot(1,2,2);
99 plot(TT,[prctile(out(2:end,:)',[50 18 84])']);
100 title('Estimated stochastic volatility');
101 axis tight
102 legend('Estimated posterior median','lower bound','upper
bound','true');
```

Published with MATLAB® 7.9

FIGURE 18. Matlab code for the stochastic volatility model continued

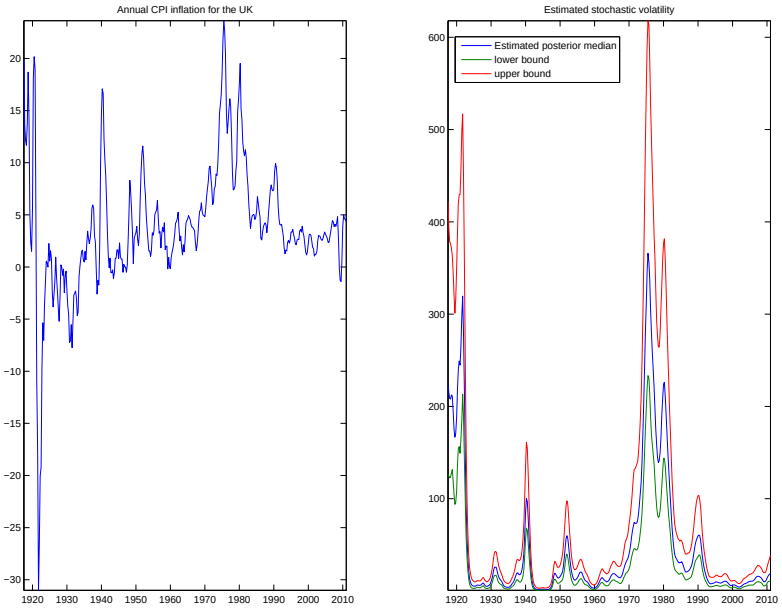

FIGURE 19. Estimated stochastic volatility of UK inflation

The right panel of figure 19 plots the estimated stochastic volatility of UK inflation.

We now a consider an extended version of this stochastic volatility model for inflation. The model now assumes a time-varying AR(1) specification for inflation with stochastic volatility in the error term. This model is given as

$$y_t = c_t + b_t y_{t-1} + \varepsilon_t \sqrt{\exp(\ln h_t)} \tag{4.16}$$

Letting $B = \{c, b\}$ the coefficients in the regression evolve as

$$B_t = B_{t-1} + e_t \tag{4.17}$$

where $e_t \sim N(0, Q)$. As before, the variance of the error term h_t evolves as

$$\begin{aligned} \ln h_t &= \ln h_{t-1} + v_t \\ & v_t \tilde{} N(0, g) \end{aligned} \tag{4.18}$$

This model can be easily estimated by combining the Carter and Kohn algorithm with the Metropolis algorithm described above. The steps are as follows:

Step 1 Set a inverse Wishart prior for Q. The prior scale matrix can be set as $Q_0 = k \times Q_{ols} \times T_0$ where T_0 is the length of training sample, Q_{ols} is the variance covariance matrix of B obtained via OLS using the traning sample and k is a scaling factor set to a small number. Obtain a starting value for $h_t, t = 0....T$ as $\hat{\varepsilon}_t^2$ and set the prior $\bar{\mu}, \bar{\sigma}$ (e.g $\bar{\mu}$ could be the log of OLS estimate of the variance of ε_t using the training sample and $\bar{\sigma}$ could be set to a big number to reflect the uncertainty in this initial guess). Set an inverse Gamma prior for g i.e. $p(g) \sim IG(g_0, v_0)$ Set a starting value for g and Q.

Step 2 Time 0 Conditional on g and B_t sample the initial value of h_t denoted by h_0 from the log normal density

$$f(h_0 \backslash h_1) = h_0^{-1} \exp\left(\frac{-(\ln h_0 - \mu_0)^2}{2\sigma_0} \right)$$

where the mean $\mu_0 = \sigma_0 \left(\frac{\bar{\mu}}{\bar{\sigma}} + \frac{\ln h_1}{g} \right)$ and $\sigma_0 = \frac{\bar{\sigma} g}{\bar{\sigma} + g}$.

ep 2 Time 1 to T-1 For *each date* t=1 to T-1 draw a new value for h_t (conditional on g and B_t) from the candidate density (call the draw $h_{t,new}$)

$$q\left(\Phi^{G+1}\right) = h_t^{-1} \exp\left(\frac{-(\ln h_t - \mu)^2}{2\sigma_h} \right)$$

where $\mu = \frac{(\ln h_{t+1} + \ln h_{t-1})}{2}$ and $\sigma_h = \frac{g}{2}$. Compute the acceptance probability (note that the residuals ε_t are used in the expression below rather than y_t as in the previous example)

$$\alpha = \min\left(\frac{h_{t,new}^{-0.5} \exp\left(\frac{-\varepsilon_t^2}{2h_{t,new}} \right)}{h_{t,old}^{-0.5} \exp\left(\frac{-\varepsilon_t^2}{2h_{t,old}} \right)}, 1 \right)$$

Draw $u \tilde{} U(0,1)$. If $u < \alpha$ set $\boldsymbol{h}_t = h_{t,new}$. Otherwise retain the old draw.

Step 2 Time T For the last time period $t = T$ compute $\mu = \ln h_{t-1}$ and $\sigma_h = g$ and draw $h_{t,new}$ from the candidate density

$$q\left(\Phi^{G+1}\right) = h_t^{-1} \exp\left(\frac{-\left(\ln h_t - \mu\right)^2}{2\sigma_h}\right)$$

Compute the acceptance probability

$$\alpha = \min\left(\frac{h_{t,new}^{-0.5} \exp\left(\frac{-\varepsilon_t^2}{2h_{t,new}}\right)}{h_{t,old}^{-0.5} \exp\left(\frac{-\varepsilon_t^2}{2h_{t,old}}\right)}, 1\right)$$

Draw $u \tilde{} U(0,1)$. If $u < \alpha$ set $\boldsymbol{h}_t = h_{t,new}$. Otherwise retain the old draw.

Step 3 Given a draw for h_t compute the residuals of the transition equation $v_t = \ln h_t - \ln h_{t-1}$. Draw g from the inverse Gamma distribution with scale parameter $\frac{v_t' v_t + g_0}{2}$ and degrees of freedom $\frac{T+v_0}{2}$. Note that this is an example of a combination of Metropolis and Gibbs sampling algorithms.

Step 4 Conditional on h_t and Q sample B_t using the Carter and Kohn algorithm as described in Chapter 3. This algorithm remains apart from the minor difference that the variance of the error to observation equation is different at each point in time. This is easily incorporated into the Kalman filter by selecting the appropriate variance at each point in time.

Step 5 Sample Q from the inverse Wishart distribution (conditional on B_t) with scale matrix $\left(B_t - B_{t-1}\right)'\left(B_t - B_{t-1}\right) + Q_0$ and degrees of freedom $T_0 + T$.

Step 6 Repeat steps 2 and 5 M times. The last L draws of h_t, g, B_t and Q provide an approximation to the marginal posterior distributions.

```
1 %a time-varying parameter model with stochastic volatility model
2 clear
3 addpath('functions');
4 %Load inflation data
5 Y=xlsread('\data\inflation.xlsx');
6 Y=((log(Y)-log(lag0(Y,4))))*100;
7 Y=Y(5:end,:);
8 T=rows(Y);
9 TT0=10;  %training sample
10 X=[lag0(Y,1) ones(T,1)];
11 Y=Y(2:end);
12 X=X(2:end,:);
13 %Independence metropolis hastings algorithm for svol model
14 %step 1 priors for g~IG(V0,T0) and initial conditions for the
stochastic
15 %volatility
```

$$p(g) \sim IG(g_0, v_0)$$

```
16 V0=0.01;  %prior scale
17 T0=1;       %prior degrees of freedom
18 Y0=Y(1:TT0);
19 X0=X(1:TT0,:);
20 B0=X0\Y0;
21 E0=Y0-X0*B0;
22 S0=(E0'*E0)/T0;
23 VV0=S0*inv(X0'*X0);
```

$\bar{\mu}$

```
24 mubar=log(std(E0)^2);
```

$\bar{\sigma}$

```
25 sigmabar=10;
26 %step 2 set starting values for time varying coefficient beta
27 beta0=B0;   %state variable  b[t-1/t-1]
28 p00=VV0;          %variance of state variable p[t-1/t-1]
29 %step 3 set prior for Q
```

$k \times Q_{ols} \times T_0$

```
30 Q0=(VV0*T0)*1e-4;
31 Q=Q0;  %intial values
32 %remove training sample
33 Y=Y(TT0+1:end,:);
34 X=X(TT0+1:end,:);
35 T=rows(Y);
36 %step 4 starting values for stochastic volatility and
37 hlast=diff(Y).^2;
38 hlast=[hlast(1:2);hlast]+0.0001;  %small number added to ensure no
value is zero
39 errors=diff(Y);
40 errors=[errors(1);errors];  %rough estimate for the errors of
observation equation
41 g=1;
42 REPS=50000;
43 BURN=45000;
44 out=[];
45 out1=[];
46 out2=[];
47 for j=1:REPS
48 %step 5 data by date metropolis hastings algorithm to draw the
stochastic
49 %volatility
50 hnew=zeros(T+1,1);
```

FIGURE 20. Matlab code for the time-varying parameter AR model with stochastic volatility

The matlab code for this example (example5.m) is shown in figures 20, 21 and 22. Lines 18 to 23 of the code estimate an AR(1) model via OLS on a training sample of 10 observations. Line 24 sets $\bar{\mu}$ as the log of the error variance using this OLS residuals. Lines 27 and 28 set the initial value of the time varying coefficients and the associated variance as the OLS estimates. Line 30 sets the prior scale matrix Q_0 using the OLS estimate of the coefficient covariance. Lines 37 to 40 set an initial value for h_t and ε_t and line 47 starts the algorithm. Lines 48 to 101 sample h_t using the independence MH step described in the previous example. The only change is that the residuals from the observation equation ε_t are used to evaluate the densities $h_{t,new}^{-0.5} \exp\left(\frac{-\varepsilon_t^2}{2h_{t,new}}\right)$ and $h_{t,old}^{-0.5} \exp\left(\frac{-\varepsilon_t^2}{2h_{t,new}}\right)$ when calculating the acceptance probability. Line 104 samples g from the inverse Gamma distribution. Line 108 samples

```
i=1;
%time period 0
hlead=hlast(i+1);
ss = sigmabar*g/(g + sigmabar);   %variance
mu = ss*(mubar/sigmabar + log(hlead)/g);  %mean
%draw from lognormal  using mu and ss
h = exp(mu + (ss^.5)*randn(1,1));
hnew(i)=h;
%time period 1 to t-1
for i=2:T
    hlead=hlast(i+1);
    hlag=hnew(i-1);
    yt=errors(i-1);  %note change
```

note that the residuals ε_t are used in the expression below rather than y_t as in the previous example

```

%mean and variance of the proposal log normal density
mu = (log(hlead)+log(hlag))/2;
ss = g/2;
%candidate draw from lognormal
htrial = exp(mu + (ss^.5)*randn(1,1));
%acceptance probability in logs
lp1 = -0.5*log(htrial) - (yt^2)/(2*htrial);  %numerator
lp0 = -0.5*log(hlast(i)) - (yt^2)/(2*hlast(i));   %denominator
accept = min([1;exp(lp1 - lp0)]);  %ensure accept<=1
u = rand(1,1);
if u <= accept;
   h = htrial;
else
   h = hlast(i);
end
hnew(i)=h;
end
%time period T
i=T+1;
yt=errors(i-1);
```

note that the residuals ε_t are used in the expression below rather than y_t as in the previous example

```
hlag=hnew(i-1);
%mean and variance of the proposal density
mu = log(hlag);   % only have ht-1
ss = g;
%candidate draw from lognormal
htrial = exp(mu + (ss^.5)*randn(1,1));
%acceptance probability
lp1 = -0.5*log(htrial) - (yt^2)/(2*htrial);
lp0 = -0.5*log(hlast(i)) - (yt^2)/(2*hlast(i));
accept = min([1;exp(lp1 - lp0)]);  %ensure accept<=1
u = rand(1,1);
if u <= accept;
   h = htrial;
else
   h = hlast(i);
end
hnew(i)=h;
%step 6 draw g from the inverse Gamma distribution
gerrors=diff(log(hnew));
g=IG(T0,V0,gerrors);  %draw from the inverse Gamma distribution
%step 7 update vale of H
hlast=hnew;
%step 8 draw the time varying coefficients using CARTER and KOHN
algorithm
```

FIGURE 21. Matlab code for the time-varying parameter AR model with stochastic volatility continued

the time-varying coefficients using the Carter Kohn algorithm. For simplicity, the code for this algorithm is moved into a seperate function carterkohn1.m saved in the functions folder. This code is identical to the examples discussed in the previous chapter apart from the minor difference that the value of the variance of the errors of the observation equation at time t is set to h_t. See line 18 in carterkohn1.m. Note that this function also returns the updated value of the error term ε_t. The inputs to this function are as follows: (1) the initial state $B_{0\backslash 0}$ (2) Variance of the initial state (3) the time-varying variance of shock to the observation equation h_t (4) Q (5) Y_t the dependent variable and (6) X_t the independent variables. Conditional on a value for B_t line 112 samples Q from the inverse Wishart distribution.

Figure 23 shows the estimated stochastic volatility and the time-varying coefficients.

Conditional on h_t and Q sample B_t using the Carter and Kohn algorithm as described in Chapter 3

```
[beta,errors]=carterkohn1(beta0',p00,hlast,Q,Y,X);
%step 9 draw Q
errorsq=diff(beta);
scaleQ=(errorsq'*errorsq)+Q0;
Q=iwpQ(T+TT0,inv(scaleQ));
```

Sample Q from the inverse Wishart distribution (conditional on B_t)

```
%save
if j>BURN
out=[out hlast];
out1=[out1 beta(:,1)];
out2=[out2 beta(:,2)];
end
end
TT=1917.75:0.25:2011;
subplot(2,2,1);
plot(TT,[prctile(out(2:end,:)',[50 18 84])'])
legend('Estimated posterior median','lower bound','upper bound');
title('Stochastic Volatility');
axis tight
subplot(2,2,2);
plot(TT,[prctile(out1(1:end,:)',[50 18 84])'])
legend('Estimated posterior median','lower bound','upper bound');
title('Time-Varying AR(1) Coefficient');
axis tight
subplot(2,2,3);
plot(TT,[prctile(out2(1:end,:)',[50 18 84])' ])
legend('Estimated posterior median','lower bound','upper bound');
title('Time-Varying constant');
axis tight
subplot(2,2,4);
plot(TT,[prctile((out2(1:end,:)./(1-out1(1:end,:)))',[50 18 84])' Y
])
legend('Estimated posterior median','lower bound','upper bound');
title('Long Run Mean of Inflation c_{t}/(1-b_{t})');
axis tight
```

Published with MATLAB® 7.9

FIGURE 22. Matlab code for the time-varying parameter AR model with stochastic volatility continued

5. A VAR with time-varying coefficients and stochastic volatility

We re-examine an extended version of the time-varying parameter VAR model shown in the previous chapter. The extension involves allowing the variance covariance matrix of the error terms to be time-varying. This model has been used in several recent studies (see for e.g. Primiceri (2005)) and is especially suited to examining the time-varying transmission of structural shocks to the economy.

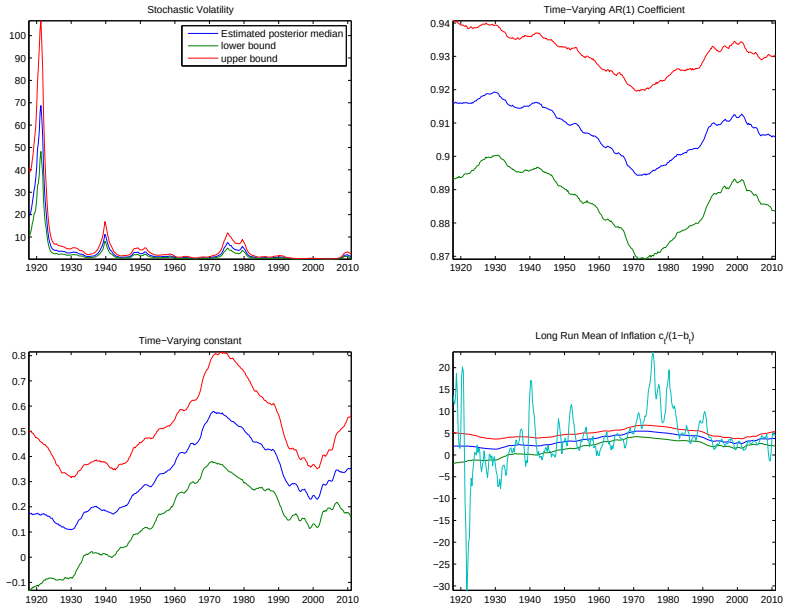

FIGURE 23. Estimates from the time-varying AR model with stochastic volatility

We consider the following VAR model with time-varying parameters

$$\begin{aligned} Y_t &= c_t + \sum_{j=1}^{P} B_{j,t} Y_{t-j} + v_t, VAR\left(v_t\right) = R_t \\ \beta_t &= \{c_t, B_{1,t} .... B_{P,t}\} \\ \beta_t &= \beta_{t-1} + e_t, VAR(e_t) = Q \end{aligned} \tag{5.1}$$

The covariance matrix of the error term v_t i.e. R_t has time-varying elements. For simplicity most studies consider the following structure for R_t

$$R_t = A_t^{-1} H_t A_t^{-1\prime} \tag{5.2}$$

where A_t is a lower triangular matrix with elements $a_{ij,t}$ and H_t is a diagonal matrix with diagonal elements $h_{i,t}$. For example for a three variable VAR

$$A_t = \begin{pmatrix} 1 & 0 & 0 \\ a_{12,t} & 1 & 0 \\ a_{13,t} & a_{23,t} & 1 \end{pmatrix}, H_t = \begin{bmatrix} h_{1,t} & 0 & 0 \\ 0 & h_{2,t} & 0 \\ 0 & 0 & h_{3,t} \end{bmatrix}$$

where

$$a_{ij,t} = a_{ij,t-1} + V_t, VAR(V_t) = D$$

and

$$\ln h_{i,t} = \ln h_{i,t-1} + z_{i,t}, VAR(z_{i,t}) = g_i$$

for $i = 1..3$. Therefore, this model has two sets of time varying 'coefficients' β_t and $a_{ij,t}$ and a stochastic volatility model for the diagonal elements $h_{i,t}$. As in the previous example, this VAR model can be estimated by combining the Carter and Kohn algorithm to draw β_t and $a_{ij,t}$ with the independence MH algorithm for the stochastic volatility. Before we describe the algorithm, it is worth noting the following relationship

$$A_t v_t = \varepsilon_t \tag{5.3}$$

where $VAR\left(\varepsilon_t\right) = H_t$. For a three variable VAR this relationship implies the following set of equations

$$\begin{pmatrix} 1 & 0 & 0 \\ a_{12,t} & 1 & 0 \\ a_{13,t} & a_{23,t} & 1 \end{pmatrix} \begin{pmatrix} v_{1,t} \\ v_{2,t} \\ v_{3,t} \end{pmatrix} = \begin{pmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \\ \varepsilon_{3,t} \end{pmatrix} \tag{5.4}$$

or expanding

$$\begin{aligned} v_{1,t} &= \varepsilon_{1,t} \\ v_{2,t} &= -a_{12,t} v_{1,t} + \varepsilon_{2,t} \\ v_{3,t} &= -a_{13,t} v_{1,t} - a_{23,t} v_{2,t} + \varepsilon_{3,t} \end{aligned} \tag{5.5}$$

where $VAR(\varepsilon_{2,t}) = h_{2t}$ and $VAR(\varepsilon_{3,t}) = h_{3t}$ and

$$a_{12,t} = a_{12,t-1} + V_{1t}, VAR(V_{1t}) = D_1 \tag{5.6}$$

$$\begin{pmatrix} a_{13,t} \\ a_{23,t} \end{pmatrix} = \begin{pmatrix} a_{13,t-1} \\ a_{23,t-1} \end{pmatrix} + \begin{pmatrix} V_{2t} \\ V_{3t} \end{pmatrix}, VAR(\begin{pmatrix} V_{2t} \\ V_{3t} \end{pmatrix}) = D_2 \tag{5.7}$$

Therefore, $a_{ij,t}$ are time varying coefficients on regressions involving the VAR residuals and can be sampled using the method described in the previous example. The Gibbs and MH algorithm for estimating this three variable time-varying VAR model consists of the following steps

Step 1a Set a prior for Q and starting values for the Kalman filter. The prior for Q is inverse Wishart $p(Q) \sim IW(Q_0, T_0)$. Note that this prior is quite crucial as it influences the amount of time-variation allowed for in the VAR model. In other words, a large value for the scale matrix Q_0 would imply more fluctuation in β_t. This prior is typically set using a training sample. The first T_0 observations of the sample are used to estimate a standard fixed coefficient VAR via OLS such that $\beta_0 = (X_{0t}'X_{0t})^{-1}(X_{0t}'Y_{0t})$ with a coefficient covariance matrix given by $p_{0\backslash 0} = \Sigma_0 \otimes (X_{0t}'X_{0t})^{-1}$ where $X_{0t} = \{Y_{0t-1}, ...Y_{0t-p}, 1\}$, $\Sigma_0 = \frac{(Y_{0t}-X_{0t}\beta_0)'(Y_{0t}-X_{0t}\beta_0)}{T_0-K}$ and the subscript 0 denotes the fact that this is the training sample. The scale matrix Q_0 is set equal to $p_{0\backslash 0} \times T_0 \times \tau$ where τ is a scaling factor chosen by the researcher. Some studies set $\tau = 3.510^{-4}$ i.e. a small number to reflect the fact that the training sample in typically short and the resulting estimates of $p_{0\backslash 0}$ maybe imprecise. Note that one can control the apriori amount of time-variation in the model by varying τ. Set a starting value for Q. The initial state is set equal to $\beta_{0\backslash 0} = vec(\beta_0)'$ and the intial state covariance is given by $p_{0\backslash 0}$.

Step 1b Set the prior for D_1 and D_2. The prior for D_1 is inverse Gamma $p(D_1) \sim IG(D_{10}, T_0)$ and the prior for D_2 is inverse Wishart $p(D_2) \sim IW(D_{20}, T_0)$. Benati and Mumtaz (2006) set $D_{10} = 0.001$ and $D_{20} = \begin{pmatrix} 0.001 & 0 \\ 0 & 0.001 \end{pmatrix}$. Let $C = \Sigma_0^{1/2}$ and let $C0$ denote the inverse of the matrix C with the diagonal normalised to 1. The initial values for $a_{ij,t}$ (i.e. the initial state $a_{ij,0\backslash 0}$) are the non-zero elements of $C0$ with the variance of the initial state set equal to $abs(a_{ij}) \times 10$ (as in Benati and Mumtaz (2006)). Set a starting value for $a_{ij,t}$.

Step 1c Obtain a starting value for $h_{i,t}, t = 0....T$ and $i = 1..3$ as $\hat{v}_{it}^2$ and set the prior $\bar{\mu}_i, \bar{\sigma}$. $\bar{\mu}_i$ can be set equal to the log of the ith diagonal element of Σ_0 and $\bar{\sigma}$ to a large number. Set an inverse Gamma prior for g_i i.e. $p(g_i) \sim IG(g_0, v_0)$. Set a starting value for g_i.

Step 2 Conditional on A_t, H_t and Q draw β_t using the Carter and Kohn algorithm. The algorithm exactly as described for the time-varying VAR without stochastic volatility in Chapter 3 with the difference that the variance of v_t changes at each point in time and this needs to be taken into account when running the Kalman filter.

Step 3 Using the draw for β_t calculate the residuals of the transition equation $\beta_t - \beta_{t-1} = e_t$ and sample Q from the inverse Wishart distribution using the scale matrix $e_t'e_t + Q_0$ and degrees of freedom $T + T_0$.

Step 4 Draw $a_{ij,t}$ the elements of A_t using the Carter and Kohn algorithm (conditional on β_t, H_t, D_1 and D_2). The state space formulation for $a_{12,t}$ is

$$\begin{aligned} v_{2,t} &= -a_{12,t}v_{1,t} + \varepsilon_{2,t}, VAR(\varepsilon_{2,t}) = h_{2,t} \\ a_{12,t} &= a_{12,t-1} + V_{1t}, VAR(V_{1t}) = D_1 \end{aligned}$$

The state space formulation for $a_{13,t}$ and $a_{23,t}$ is

$$\begin{aligned} v_{3,t} &= -a_{13,t}v_{1,t} - a_{23,t}v_{2,t} + \varepsilon_{3,t}, VAR(\varepsilon_{3,t}) = h_{3,t} \\ \begin{pmatrix} a_{13,t} \\ a_{23,t} \end{pmatrix} &= \begin{pmatrix} a_{13,t-1} \\ a_{23,t-1} \end{pmatrix} + \begin{pmatrix} V_{2t} \\ V_{3t} \end{pmatrix}, VAR(\begin{pmatrix} V_{2t} \\ V_{3t} \end{pmatrix}) = D_2 \end{aligned}$$

Note that these two formulations are just time-varying regressions in the residuals and the Carter and Kohn algorithm is applied to each seperately to draw $a_{12,t}, a_{13,t}$ and $a_{23,t}$.

Step 5. Conditional on a draw for $a_{12,t}, a_{13,t}$ and $a_{23,t}$ calculate the residuals V_{1t}, V_{2t} and V_{3t}. Draw D_1 from the inverse Gamma distribution with scale parameter $\frac{V_{1t}'V_{1t}+D_{1,0}}{2}$ and degrees of freedom $\frac{T+T_0}{2}$. Draw D_2 from the inverse Wishart distribution with scale matrix $V_{2t}'V_{2t} + D_{2,0}$ and degrees of freedom $T + T_0$.

Step 6 Using the draw of A_t from step 4 calculate $\varepsilon_t = A_t v_t$ where $\varepsilon_t = \begin{pmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \\ \varepsilon_{3,t} \end{pmatrix}$. Note that ε_t are contemporaneously uncorrelated. We can therefore draw $h_{i,t}$ for $i = 1..3$ separately by simply applying the independence MH algorithm described above for each ε_t (conditional on a draw for g_i).

Step 7 Conditional on a draw for $h_{i,t}$ for $i = 1..3$ draw g_i from the inverse Gamma distribution with scale parameter $\frac{(\ln h_{i,t} - \ln h_{i,t-1})'(\ln h_{i,t} - \ln h_{i,t-1}) + g_0}{2}$ and degrees of freedom $\frac{T+v_0}{2}$.

Step 8 Repeat steps 2 and 7 M times. The last L draws provide an approximation to the marginal posterior distributions of the model parameters.

```
1 clear
2 addpath('functions');
3 % a TVP-VAR with stochastic volatility using dlog(GDP) dlog(CPI) and R
for the US 1962 2004
4 %load data
5 data=xlsread('\data\usdata.xls')/100;
6 N=size(data,2);
7 L=2;   %number of lags in the VAR
8 Y=data;
9 X=[ lag0(Y,1) lag0(Y,2) ones(size(Y,1),1) ];
10 Y=Y(3:end,:);
11 X=X(3:end,:);
12 %step 1 set starting values and priors using a pre-sample of 10 years
13 T0=40;
14 y0=Y(1:T0,:);
15 x0=X(1:T0,:);
16 b0=x0\y0;
17 e0=y0-x0*b0;
18 sigma0=(e0'*e0)/T0;
19 V0=kron(sigma0,inv(x0'*x0));
20 %priors for the variance of the transition equation
```

The scale matrix Q_0 is set equal to $p_{0\backslash 0} \times T_0 \times \tau$

```
21 Q0=V0*T0*3.5e-04;  %prior for the variance of the transition equation
error
```

$$p_{0\backslash 0} = \Sigma_0 \otimes (X_{0t}' X_{0t})^{-1}$$

```
22 P00=V0;                                          % variance of
the intial state vector  variance of state variable p[t-1/t-1]
```

$$\beta_{0\backslash 0} = vec(\beta_0)'$$

```
23 beta0=vec(b0)';      % intial state vector
%state variable  b[t-1/t-1]
24 %priors and starting values for aij
25 C0=chol(sigma0);
26 C0=C0./repmat(diag(C0),1,N);
```

let $C0$ denote the inverse of the matrix C with the diagonal normalised to 1

```
27 C0=inv(C0)';
```

$a_{ij,0\backslash 0}$

```
28 a10=C0(2,1);   %intial state vector
29 a20=C0(3,1:2);  %intial state vector second equation
```

the variance of the initial state set equal to $abs(a_{ij}) \times 10$

```
30 pa10=abs(a10)*10;  %variance of the state vector
31 pa20=diag(a20)*10;  %variance of the state vector
```

$$D_1 = 0.001 \text{ and } D_2 = \begin{pmatrix} 0.001 & 0 \\ 0 & 0.001 \end{pmatrix}$$

```
32 D10=10^(-3);  %prior scale matrix for D1
33 D20=10^(-3)*eye(2); %prior scale matrix for D2
34 %remove intial Sample
35 Y=Y(T0+1:end,:);
36 X=X(T0+1:end,:);
37 T=rows(X);
38 %priors and starting values for the stochastic vol
```

Obtain a starting value for $h_{it}, t = 0....T$ and $i = 1..3$ as $\hat{v}_{it}^2$

FIGURE 24. Matlab code for the time-varying VAR with stochastic volatility

The matlab code for estimating this model (example6.m) is shown in figures 24, 25, 26 and 27. We consider a time-varying VAR model with two lags using US data on GDP growth, CPI inflation and the Federal Funds rate over the period 1954Q3 to 2010Q2 in this code. Lines 25 to 27 set the initial values for the elements of A_t by calculating the matrix $C0$. Lines 30 amd 31 set the variance around these initial values. Lines 32 and 33 set the prior scale matrices D_{10} and D_{20}. Lines 38 to 45 set the priors and starting values for the stochastic volatility models for the transformed VAR residuals ε_t. Lines 59 to 112 contain the Carter and Kohn algorithm to sample the VAR coefficients β_t. The only change relative to the example in chapter 3 is on lines 69 to 72. Line 70 using the function chofac.m to reshape the value of $a_{ij,t}$ at time t into a lower triangular matrix. Line 72 calculates the VAR error covariance

```
39 hlast=(diff(Y).^2)+0.0001;
40 hlast=[hlast(1:2,:);hlast];  %rough intial guess for svol
41 g=ones(3,1);  %rough guess for the variance of the transition
equation
42 g0=0.01^2;  %scale parameter for inverse gamma
43 Tg0=1;
44 mubar=log(diag(sigma0));
45 sigmabar=10;
46 %initialise parameters
47 Q=Q0;
48 D1=D10;
49 D2=D20;
50 a1=repmat(a10,T,1);
51 a2=repmat(a20,T,1);
52 %Gibbs sampling algorithm Step 2
53 reps=100000;
54 burn=99000;
55 mm=1;
56 for m=1:reps
57 m
58 %%Step 2a Set up matrices for the Kalman Filter
59 ns=cols(beta0);
60 F=eye(ns);
61 mu=0;
62 beta tt=[];          %will hold the filtered state variable
63 ptt=zeros(T,ns,ns);    % will hold its variance
64 beta11=beta0;
65 p11=P00;
66 % %%%%%%%%%%%Step 2b run Kalman Filter
67 for i=1:T
68    x=kron(eye(N),X(i,:));
69 a=[a1(i) a2(i,:)];
70 A=chofac(N,a');
71 H=diag(hlast(i+1,:));
                                    R_t = A_t^{-1} H_t A_t^{-1\prime}
72 R=inv(A)*H*inv(A)';
73     %Prediction
74 beta10=mu+beta11*F';
75 p10=F*p11*F'+Q;
76 yhat=(x*(beta10)')';
77 eta=Y(i,:)-yhat;
78 feta=(x*p10*x')+R;
79 %updating
80 K=(p10*x')*inv(feta);
81 beta11=(beta10'+K*eta')';
82 p11=p10-K*(x*p10);
83 ptt(i,:,:)=p11;
84 beta tt=[beta tt;beta11];
85 end
86 %%%%%%%%%%%end of Kalman
Filter%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
87 %step 2c Backward recursion to calculate the mean and variance of the
distribution of the state
88 %vector
89 chck=-1;
90 while chck<0
91 beta2 = zeros(T,ns);   %this will hold the draw of the state variable
92 wa=randn(T,ns);
93 error=zeros(T,N);
94 roots=zeros(T,1);
95 i=T;  %period t
```

FIGURE 25. Matlab code for the time-varying VAR with stochastic volatility

matrix for that time period and this is used in Kalman filter equations. Line 116 samples Q from the inverse Wishart distribution. Line 121 uses the Carter and Kohn algorithm to sample $a_{12,t}$ (where for simplicity the code for the algorithm is in the function carterkohn1.m). Line 122 samples $a_{13,t}$ and $a_{23,t}$ using the same function. Lines 124 and 125 sample D_1 from the inverse Gamma distribution. Lines 127 and 128 sample D_2 from the inverse Wishart distribution. Lines 131 to 136 calculate $\varepsilon_t = A_t v_t$. Lines 138 to 142 use the independence MH algorithm to draw $h_{i,t}, i = 1..3$ using these ε_t. The code for the algorithm is identical to the two previous examples but is included in the function getsvol.m for simplicity. This function takes in the following inputs (1) the previous draw of $h_{i,t}$ (2) g_i (3) $\bar{\mu}$ (4) $\bar{\sigma}$ (5) ε_t and returns a draw for $h_{i,t}$. Lines 145 to 148 draw g_i from the inverse Gamma distribution.

```
96 p00=squeeze(ptt(i,:,:));
97 beta2(i,:)=beta_tt(i:i,:)+(wa(i:i,:)*chol(p00));   %draw for beta in
period t from N(beta tt,ptt)
98 error(i,:)=Y(i,:)-X(i,:)*reshape(beta2(i:i,:),N*L+1,N);  %var
residuals
99 roots(i)=stability(beta2(i,:)',N,L);
100 %periods t-1..to .1
101 for i=T-1:-1:1
102 pt=squeeze(ptt(i,:,:));
103 bm=beta tt(i:i,:)+(pt*F'*inv(F*pt*F'+Q)*(beta2(i+1:i+1,:)-
beta_tt(i,:)*F')')';  %update the filtered beta for information
contained in beta[t+1]
%i.e. beta2(i+1:i+1,:) eq 8.16 pp193 in Kim Nelson
104 pm=pt-pt*F'*inv(F*pt*F'+Q)*F*pt;  %update covariance of beta
105 beta2(i:i,:)=bm+(wa(i:i,:)*chol(pm));  %draw for beta in period t
from N(bm,pm)eq 8.17 pp193 in Kim Nelson
106 error(i,:)=Y(i,:)-X(i,:)*reshape(beta2(i:i,:),N*L+1,N);  %var
residuals
107 roots(i)=stability(beta2(i,:)',N,L);
108 end
109 if sum(roots)==0
110     chck=1;
111 end
112 end
113 % step 3 sample Q from the IW distribution
114 errorq=diff(beta2);
115 scaleQ=(errorq'*errorq)+Q0;
116 Q=iwpQ(T+T0,inv(scaleQ));
117 %step4 sample aij using the carter kohn algorithm
```

The state space formulation for $a_{12,t}$ is

$$v_{2,t} = -a_{12,t}v_{1,t} + \varepsilon_{2,t}, VAR(\varepsilon_{2,t}) = h_{2,t}$$
$$a_{12,t} = a_{12,t-1} + V_{1t}, VAR(V_{1t}) = D_1$$

The state space formulation for $a_{13,t}$ and $a_{23,t}$ is

$$v_{3,t} = -a_{13,t}v_{1,t} - a_{23,t}v_{2,t} + \varepsilon_{3,t}, VAR(\varepsilon_{3,t}) = h_{3,t}$$
$$\begin{pmatrix} a_{13,t} \\ a_{23,t} \end{pmatrix} = \begin{pmatrix} a_{13,t-1} \\ a_{23,t-1} \end{pmatrix} + \begin{pmatrix} V_{2t} \\ V_{3t} \end{pmatrix}, VAR(\begin{pmatrix} V_{2t} \\ V_{3t} \end{pmatrix}) = D_2$$

```
118 v3=error(:,3);
119 v2=error(:,2);
120 v1=error(:,1);
121 [a1,trash]=carterkohn1(a10,pa10,hlast(:,2),D1,v2,-v1);
122 [a2,trash]=carterkohn1(a20,pa20,hlast(:,3),D2,v3,[-v1 -v2]);
123 %step 5 sample D1 and D2
124 a1errors=diff(a1);
125 D1=IG(T0,D10,a1errors);  %draw from the inverse Gamma distribution
126 a2errors=diff(a2);
127 scaleD2=(a2errors'*a2errors)+D20;
128 D2=iwpQ(T+T0,inv(scaleD2)); %draw from inverse Wishart
129 %step 6 sample h_i seperately for i=1,3
130  %step 6a calculate epsilon=A*v
```

Using the draw of A_t from step 4 calculate $\varepsilon_t = A_t v_t$

```
131  epsilon=[];
132  for i=1:T
133      a=[a1(i) a2(i,:)];
134      A=chofac(N,a');
135      epsilon=[epsilon;error(i,:)*A'];
136  end
137  %sample stochastic vol for each epsilon using the MH algorithm
```

FIGURE 26. Matlab code for the time-varying VAR with stochastic volatility

Figure 28 plots the estimated impulse response to a monetary policy shock (identified via sign restrictions) and the estimated stochastic volatility.

6. Convergence of the MH algorithm

```
  hnew=[];
  for i=1:N
      htemp=getsvol(hlast(:,i),g(i),mubar(i),sigmabar,epsilon(:,i));
      hnew=[hnew htemp];
  end
hlast=hnew;
%step 7 Sample G for IG distribution
for i=1:N
    gerrors=diff(log(hnew(:,i)));
g(i)=IG(Tg0,g0,gerrors);  %draw from the inverse Gamma distribution
end
if m>burn
    %save output from Gibbs sampler
    out1(mm,1:T,:)=beta2;
    out2(mm,1:T,1:N)=hlast(2:end,:);
    out3(mm,1:N*(N*L+1),1:N*(N*L+1))=Q;
    out4(mm,1:T,1:(N*(N-1))/2)=[a1 a2];
    out5(mm,1)=D1;
    out6(mm,1:2,1:2)=D2;
    out7(mm,1:N)=g';
    mm=mm+1;
end
end
%save results
save tvp.mat out1 out2 out3 out4 out5 out6 out7
%compute irf to a policy shock using sign restrictions
horz=40;% impulse response horizon
irfmat=zeros(size(out1,1),T,horz,N); %empty matrix to save impulse
response to a policy shock
for i=1:size(out1,1);

    for j=1:size(out1,2)

       H=diag(squeeze(out2(i,j,:)));
       a=squeeze(out4(i,j,:));
        A=chofac(N,a);
        sigma=inv(A)*H*inv(A)';  %covariance matrix
   %sign restrictions
        chck=-1;
        while chck<0
        K=randn(N,N);
        QQ=getQR(K);
        A0hat=chol(sigma);
        A0hat1=(QQ*A0hat);  %candidate draw
        for m=1:N
        %check signs in each row
        e1=A0hat1(m,1)<0;  %Response of Y
        e2=A0hat1(m,2)<0;  %Response of P
        e3=A0hat1(m,3)>0;  %Response of R

        if e1+e2+e3==3
            MP=A0hat1(m,:);
            chck=10;
        else
            %check signs but reverse them
         e1=-A0hat1(m,1)<0;  %Response of Y
        e2=-A0hat1(m,2)<0;  %Response of P
        e3=-A0hat1(m,3)>0;  %Response of R

        if e1+e2+e3==3
            MP=-A0hat1(m,:);
```

FIGURE 27. Matlab code for the time-varying VAR with stochastic volatility

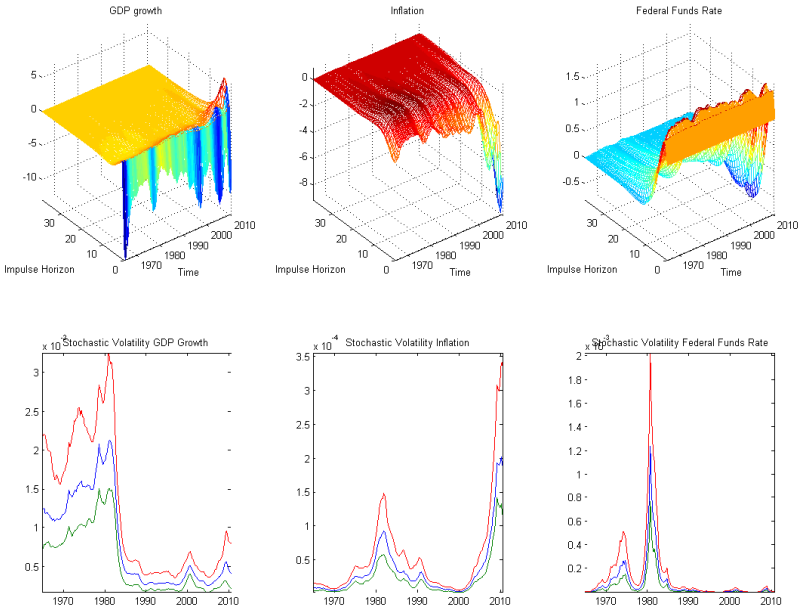

FIGURE 28. Response to a monetary policy shock from the time-varying VAR with stochastic volatility (Top panel) and the estimated stochastic volatility (bottom panel)

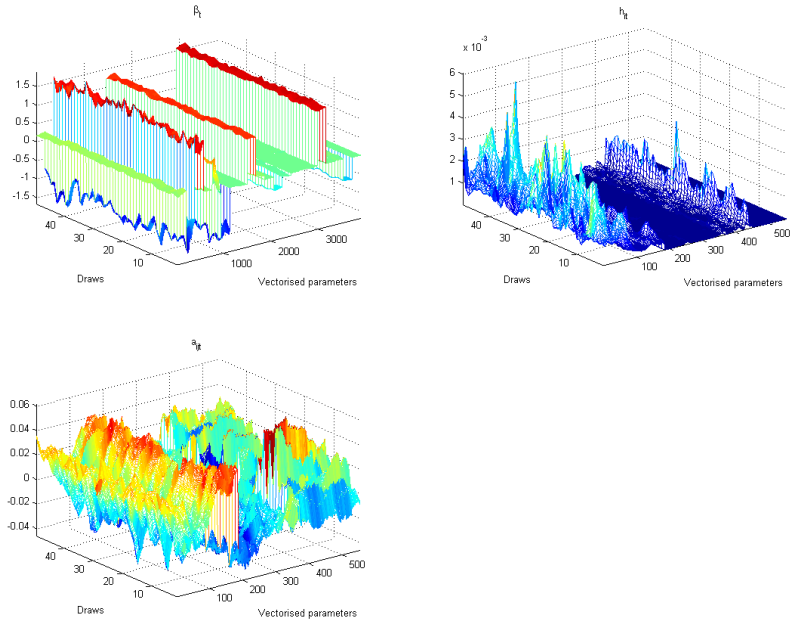

FIGURE 29. Recursive mean for key parameters of the time-varying VAR model

Most of the methods for checking convergence of the Gibbs sampler (see Chapter 1) can be applied immediately to output from the MH algorithm. Several studies present simple statistics such as recursive means of the MH draws and the autocorrelation functions to test if the algorithm has converged. As an example we present the recursive means of the retained draws for the time-varying parameter VAR considered in the previous section. As described above this model is estimated using a mixture of Gibbs and MH steps. Figure 29 presents the recursive means calculated every 20 draws for β_t, h_{it} and $a_{ij,t}$. The X-axis of each panel represents these parameterised vectorised. The Y-axis represents the draws. The recursive means usggest convergence for β_t, $a_{ij,t}$ but indicate some variation in the means for h_{it} possibly suggesting that more draws are required for this model.

Gelman and Rubin (1992) suggest a diagnostic for monitoring the convergence of multiple MH chains (for estimating the same model) started from different starting values. For every parameter of interest ξ Gelman and Rubin (1992) calculate the *within chain variance* as

$$\begin{aligned} W_T &= \frac{1}{M}\sum_{m=1}^{M}\frac{1}{T}\sum_{t=1}^{T}\left(\bar{\xi}_{t,m}-\bar{\xi}_m\right)^2, \\ \bar{\xi}_m &= \frac{1}{T}\sum_{t=1}^{T}\xi_{tm}, \bar{\xi}=\frac{1}{M}\sum_{m=1}^{M}\bar{\xi}_m \end{aligned}$$

where T denotes the total number of iterations in each of the M MH algorithms.

Gelman and Rubin (1992) calculate the between chain variance

$$B_T = \frac{1}{M}\sum_{m=1}^{M}\left(\bar{\xi}_m-\bar{\xi}\right)^2$$

They argue that W_T underestimates the variance of ξ (before convergence) as the MH algorithm has not explored the parameter space. In contrast, $\sigma_T^2=\frac{T-1}{T}W_T+B_T$ overestimates this variance due to dispersed starting values. If the MH algorithm has converged then W_T and σ_T^2 should be similar. Gelman and Rubin (1992) suggest calculating the statistic

$$R_T = \frac{\sigma_T^2+\frac{B_T}{M}}{W_T}\frac{v_T}{v_T-2}$$

where $v_T=\frac{2\left(\sigma_T^2+\frac{B_T}{M}\right)^2}{W_T}$ and checking if this is close to 1 which would indicate convergence of the MH algorithm.

7. Further Reading

- Koop (2003) chapter 5 provides an excellent description of the Metropolis Hastings algorithm.

8. Appendix: Computing the marginal likelihood using the Gelfand and Dey method

Gelfand and Dey (1994) introduce a method for computing the marginal likelihood that is particularly convenient to use when employing the Metropolis Hastings algorithm. This method is based on the following result.

$$E\left[\frac{f(\Phi)}{F(Y\backslash\Phi)\times P(\Phi)}\backslash Y\right]=\frac{1}{F(Y)} \tag{8.1}$$

where $F(Y\backslash\Phi)$ denotes the likelihood function, $P(\Phi)$ is the prior distribution, $F(Y)$ is the marginal likelihood and $f(\Phi)$ is any pdf with support Θ defined within the region of the posterior. The proof of equation 8.1 can be obtained by noting that $E\left[\frac{f(\Phi)}{F(Y\backslash\Phi)\times P(\Phi)}\backslash Y\right]=\int\frac{f(\Phi)}{F(Y\backslash\Phi)\times P(\Phi)}\times H(\Phi\backslash Y)\,d\Phi$ where $H(\Phi\backslash Y)$ is the posterior distribution. Note that $H(\Phi\backslash Y)=\frac{F(Y\backslash\Phi)\times P(\Phi)}{F(Y)}$ and the density $f(\Phi)$ integrates to 1 leaving us with the right hand side in equation 8.1.

We can approximate the marginal likelihood as $\frac{1}{M}\sum_{j=1}^{M}\frac{f(\Phi_j)}{F(Y\backslash\Phi_j)\times P(\Phi_j)}$ where Φ_j denotes draws of the parameters from Metropolis Hastings algorithm and $F(Y\backslash\Phi_j)\times P(\Phi_j)$ is the posterior evaluated at each draw. Geweke (1998) recommends using a truncated normal distribution for $f(\Phi)$. This distribution is truncated at the tails to ensure that $f(\Phi)$ is bounded from above, a requirement in Gelfand and Dey (1994). In particular, Geweke (1998) suggest using

$$f(\Phi)=\frac{1}{p(2\pi)^{k/2}}\left|\hat{\Sigma}\right|^{-1/2}\exp\left[-0.5\left(\Phi_j-\hat{\Phi}\right)\hat{\Sigma}^{-1}\left(\Phi_j-\hat{\Phi}\right)'\right]\times I\left(\Phi_j\in\hat{\Theta}\right) \tag{8.2}$$

where $\hat{\Phi}$ is the posterior mean, $\hat{\Sigma}$ is the posterior covariance and k is the number of parameters. The indicator function $I\left(\Phi_j\in\hat{\Theta}\right)$ takes a value of 1 if

$$\left[\left(\Phi-\hat{\Phi}\right)\hat{\Sigma}^{-1}\left(\Phi-\hat{\Phi}\right)'\right]\leq\chi^2_{1-p}(k)$$

where $\chi^2_{1-p}(k)$ is the inverse χ^2 cumulative distribution function with degrees of freedom k and probability p. Thus $\chi^2_{1-p}(k)$ denotes the value that exceeds $1-p\%$ of the samples from a χ^2 distribution with k degrees of freedom. The indicator function $I\left(\Phi_j\in\hat{\Theta}\right)$ therefore removes 'extreme' values of Φ_j. For more details, see Koop (2003) page 104.

In figures 30 and 31 we estimate the marginal likelihood for a linear regression model via the Gelfand and Dey method. The model is exactly used in the appendix to Chapter 1 and is based on artifical data. A simple random walk Metropolis Hastings algorithm is used to approximate the posterior on lines 35 to 62 and we save the log posterior evaluated at each draw and each draw of the parameters. Lines 65 and 66 calculate the posterior mean and variance. We set $1-p=0.1$ on line 68. In practice, different value of $1-p$ can be tried to check robustness of the estimate. On line 70 we evaluate the inverse χ^2 CDF. Line 71 to 78, loop through the saved draws of the parameters. On 73 we calculate $\left(\Phi-\hat{\Phi}\right)\hat{\Sigma}^{-1}\left(\Phi-\hat{\Phi}\right)'$. If this is less than or equal to $\chi^2_{1-p}(k)$ we evaluate $\frac{f(\Phi_j)}{F(Y\backslash\Phi_j)\times P(\Phi_j)}$ in logs, adding the constant lpost_mode to prevent overflow.

```
clear;
clc
addpath('functions')
%generate artificial data
T=100;
X=[ones(T,1) randn(T,1)];
btrue=[1;0.5];
sigmatrue=0.2;
Y=X*btrue+randn(T,1)*sqrt(sigmatrue);
%set priors
T0=3;
D0=2.5;
B0=zeros(2,1);
Sigma0=eye(2)*(4);
%step 2 set SIGMA matrix via OLS estimation
yols=Y;
xols=X;
bols=inv(xols'*xols)*(xols'*yols);
eols=yols-xols*bols;
sols=((eols'*eols)/T);
vols=sols*inv(xols'*xols);
K=0.1;
P=eye(3);                    %this is the variance of the metropolis
hastings random walk based partly on OLS estimates
P(1,1)=(vols(1,1));
P(2,2)=(vols(2,2));
P(3,3)=0.1;
%analytical computation of the marginal likelihood
mlika=mlikols(B0,Sigma0,T0,D0,Y,X);
disp('Analytical log Marginal Likelihood');
disp(log(mlika));
sigma2=1;
Gammaold=[0;0;1]; %starting values
%compute posterior
posteriorOLD=postols(Y,X,Gammaold,B0,Sigma0,T0,D0);
reps=15000;   %total numbers of MH iterations
burn=4000;   %percent of burn-in iterations
outpost=[];  %will hold posterior
outparam=[]; %will hold parameters
naccept=0;
for i=1:reps
Gammanew=Gammaold+(randn(1,3)*chol(P*K))';
  sigma2=Gammanew(3);
    if sigma2<0
        posteriorNEW=-1000000;
    else
        posteriorNEW=postols(Y,X,Gammanew,B0,Sigma0,T0,D0);
    end
        accept=min([exp(posteriorNEW-posteriorOLD);1]);
%min(accept,1)

    u=rand(1,1);  %random number from the uniform dist

    if u<accept
        Gammaold=Gammanew;  %accept draw
        naccept=naccept+1;  %count number of acceptances
        posteriorOLD=posteriorNEW;

    end
if i>burn
```

FIGURE 30. Matlab code to calculate the marginal likelihood via the Gelfand and Dey Method

```
    outpost=[outpost;posteriorOLD];
     outparam=[outparam;Gammaold'];
end
end
%calculate the marginal likelihood using Gelfand and Dey method
%posterior mean and variance
pmean=mean(outparam);
pvar=cov(outparam);
lpost_mode=max(outpost);
p=0.1; %critical value of the Chi-squared distribution
npara=size(outparam,2); %number of parameters
critval = chi2inv(p,npara);
    tmp = 0;
    for i = 1:size(outparam,1);
        deviation  = (outparam(i,:)-pmean)*inv(pvar)*((outparam(i,:)-
pmean))';
        if deviation <= critval;
            lftheta = -log(p)-
(npara*log(2*pi)+log(det(pvar))+deviation)/2;
            tmp = tmp + exp(lftheta - outpost(i)+lpost mode);
        end;
    end;
mlik=lpost mode-log(tmp/size(outparam,1));
disp('Gelfand and Dey log Marginal Likelihood');
disp(mlik);
```

Published with MATLAB® 7.9

FIGURE 31. Matlab code to calculate the marginal likelihood via the Gelfand and Dey Method (continued)

Bibliography

Banbura, Marta, Domenico Giannone and Lucrezia Reichlin, 2007, Bayesian VARs with Large Panels, *CEPR Discussion Papers 6326*, C.E.P.R. Discussion Papers.
URL: *http://ideas.repec.org/p/cpr/ceprdp/6326.html*

Bauwens, Luc, Michel Lubrano and Jean François Richard, 1999, *Bayesian inference in dynamic econometric models*, Oxford University Press.

Benati, Luca and Haroon Mumtaz, 2006, US evolving Macroeconomic Dynamics: A structural investigation., Mimeo European Central Bank.

Bernanke, Ben, Jean Boivin and Piotr S. Eliasz, 2005, Measuring the Effects of Monetary Policy: A Factor-augmented Vector Autoregressive (FAVAR) Approach, *The Quarterly Journal of Economics* **120**(1), 387–422.
URL: *http://ideas.repec.org/a/tpr/qjecon/v120y2005i1p387-422.html*

Canova, Fabio, 2007, *Methods for Applied Macroeconomic Research*, Princeton University Press, Princeton.

Carriero, Andrea, George Kapetanios and Massimiliano Marcellino, 2010, Forecasting Government Bond Yields with Large Bayesian VARs, *Working Papers 662*, Queen Mary, University of London, School of Economics and Finance.
URL: *http://ideas.repec.org/p/qmw/qmwecw/wp662.html*

Carter, C. K. and R. Kohn, 1994, On Gibbs sampling for state space models, *Biometrika* **81**(3), 541–553.
URL: *http://biomet.oxfordjournals.org/content/81/3/541.abstract*

Casella, George and Edward I. George, 1992, Explaining the Gibbs Sampler, *The American Statistician* **46**(3), pp. 167–174.
URL: *http://www.jstor.org/stable/2685208*

Chen, Cathy W. S. and Jack C. Lee, 1995, BAYESIAN INFERENCE OF THRESHOLD AUTOREGRESSIVE MODELS, *Journal of Time Series Analysis* **16**(5), 483–492.
URL: *http://dx.doi.org/10.1111/j.1467-9892.1995.tb00248.x*

Chib, Siddhartha, 1993, Bayes regression with autoregressive errors : A Gibbs sampling approach, *Journal of Econometrics* **58**(3), 275–294.
URL: *http://ideas.repec.org/a/eee/econom/v58y1993i3p275-294.html*

Chib, Siddhartha, 1995, Marginal Likelihood from the Gibbs Output, *Journal of the American Statistical Association* **90**(432), 1313–1321.
URL: *http://www.jstor.org/stable/2291521*

Chib, Siddhartha and Srikanth Ramamurthy, 2010, Tailored randomized block MCMC methods with application to DSGE models, *Journal of Econometrics* **155**(1), 19–38.
URL: *http://ideas.repec.org/a/eee/econom/v155y2010i1p19-38.html*

Cogley, Timothy and Thomas J. Sargent, 2002, Evolving Post-World War II U.S. Inflation Dynamics, *NBER Macroeconomics Annual 2001, Volume 16*, NBER Chapters, National Bureau of Economic Research, Inc, pp. 331–388.
URL: *http://ideas.repec.org/h/nbr/nberch/11068.html*

Doan, Thomas, Robert B. Litterman and Christopher A. Sims, 1983, Forecasting and Conditional Projection Using Realistic Prior Distributions, *NBER Working Papers 1202*, National Bureau of Economic Research, Inc.
URL: *http://ideas.repec.org/p/nbr/nberwo/1202.html*

Fry, Renee and Adrian Pagan, 2007, Some Issues in Using Sign Restrictions for Identifying Structural VARs, *NCER Working Paper Series 14*, National Centre for Econometric Research.
URL: *http://ideas.repec.org/p/qut/auncer/2007-8.html*

Gelfand, A. E. and D. K. Dey, 1994, Bayesian Model Choice: Asymptotics and Exact Calculations, *Journal of the Royal Statistical Society. Series B (Methodological)* **56**(3), pp. 501–514.
URL: *http://www.jstor.org/stable/2346123*

Gelman, Andrew and Donald B. Rubin, 1992, Inference from Iterative Simulation Using Multiple Sequences, *Statistical Science* **7**(4), pp. 457–472.
URL: *http://www.jstor.org/stable/2246093*

Geweke, John, 1991, Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments, *Technical report*.

Geweke, John, 1998, Using simulation methods for Bayesian econometric models: inference, development, and communication, *Technical report*.

Hamilton, J. D., 1994, *Time series analysis*, Princeton University Press, Princeton.

Jacquier, E, N Polson and P Rossi, 2004, Bayesian analysis of stochastic volatility models, *Journal of Business and Economic Statistics* **12**, 371–418.

Kadiyala, K Rao and Sune Karlsson, 1997, Numerical Methods for Estimation and Inference in Bayesian VAR-Models, *Journal of Applied Econometrics* **12**(2), 99–132.
URL: *http://ideas.repec.org/a/jae/japmet/v12y1997i2p99-132.html*

Kim, C-J. and C. R. Nelson, 1999, *State-space models with regime switching*, MIT Press, Cambridge, Massachusetts.

Kim, Sangjoon, Neil Shephard and Siddhartha Chib, 1998, Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models, *Review of Economic Studies* **65**(3), 361–93.
URL: *http://ideas.repec.org/a/bla/restud/v65y1998i3p361-93.html*

Koop, Gary, 2003, *Bayesian Econometrics*, Wiley.

Negro, Marco Del and Frank Schorfheide, 2004, Priors from General Equilibrium Models for VARS, *International Economic Review* **45**(2), 643–673.
URL: *http://ideas.repec.org/a/ier/iecrev/v45y2004i2p643-673.html*

Primiceri, G, 2005, Time varying structural vector autoregressions and monetary policy, *The Review of Economic Studies* **72**(3), 821–852.

Ramirez, Juan Rubio, Daniel Waggoner and Tao Zha, 2010, Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference, *Review of Economic Studies* **77**(2), 665–696.
URL: *http://ideas.repec.org/a/bla/restud/v77y2010i2p665-696.html*

Robertson, John C. and Ellis W. Tallman, 1999, Vector autoregressions: forecasting and reality, *Economic Review* (Q1), 4–18.
URL: *http://ideas.repec.org/a/fip/fedaer/y1999iq1p4-18nv.84no.1.html*

Robertson, John C, Ellis W Tallman and Charles H Whiteman, 2005, Forecasting Using Relative Entropy, *Journal of Money, Credit and Banking* **37**(3), 383–401.
URL: *http://ideas.repec.org/a/mcb/jmoncb/v37y2005i3p383-401.html*

Sims, Christopher A and Tao Zha, 1998, Bayesian Methods for Dynamic Multivariate Models, *International Economic Review* **39**(4), 949–68.
URL: *http://ideas.repec.org/a/ier/iecrev/v39y1998i4p949-68.html*

Villani, Mattias, 2009, Steady-state priors for vector autoregressions, *Journal of Applied Econometrics* **24**(4), 630–650.
URL: *http://dx.doi.org/10.1002/jae.1065*

Waggoner, Daniel F. and Tao Zha, 1999, Conditional Forecasts In Dynamic Multivariate Models, *The Review of Economics and Statistics* **81**(4), 639–651.
URL: *http://ideas.repec.org/a/tpr/restat/v81y1999i4p639-651.html*

Zellner, Arnold, 1971, *An introduction to Bayesian Inference in Econometrics*, Wiley.