

THE FEDERAL RESERVE BANK of KANSAS CITY
ECONOMIC RESEARCH DEPARTMENT

Estimating VAR's Sampled at Mixed or Irregular Spaced Frequencies: A Bayesian Approach

Ching Wai (Jeremy) Chiu, Bjørn Eraker, Andrew T. Foerster, Tae
Bong Kim and Hernán D. Seoane

December 2011

RWP 11-11



RESEARCH WORKING PAPERS

Estimating VAR's Sampled at Mixed or Irregular Spaced Frequencies : A Bayesian Approach*

Ching Wai (Jeremy) Chiu[†] Bjørn Eraker[‡] Andrew T. Foerster[§]

Tae Bong Kim[¶] Hernán D. Seoane^{||}

December 16, 2011

Abstract

Economic data are collected at various frequencies but econometric estimation typically uses the coarsest frequency. This paper develops a Gibbs sampler for estimating VAR models with mixed and irregularly sampled data. The approach allows efficient likelihood inference even with irregular and mixed frequency data. The Gibbs sampler uses simple conjugate posteriors even in high dimensional parameter spaces, avoiding a non-Gaussian likelihood surface even when the Kalman filter applies. Two applications illustrate the methodology and demonstrate efficiency gains from the mixed frequency estimator: one constructs quarterly GDP estimates from monthly data, the second uses weekly financial data to inform monthly output.

Keywords: Gibbs Sampling, Mixed Frequency Data, VAR, Bayesian Estimation

JEL classification: C11, C22, G10, E27

*Previously titled "Bayesian Mixed Frequency VAR's." The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Kansas City or the Federal Reserve System.

[†]Duke University

[‡]University of Wisconsin, School of Business

[§]Corresponding author: Andrew Foerster, Federal Reserve Bank of Kansas City, 1 Memorial Drive, Kansas City, MO 64198, 816-881-2751, andrew.foerster@kc.frb.org.

[¶]Korea Development Institute

^{||}Universidad Carlos III de Madrid

1 Introduction

Economic data are rarely collected at the same instances in time. Data from liquid markets are available almost continuously, while aggregate macro data in many cases are available only at monthly, quarterly, or annual frequencies. Mixed and irregular sampling frequencies represent a significant challenge to time-series econometricians.

This paper develops Bayesian estimation of mixed frequency Vector Autoregressions (VAR's). The method is a simple, yet very powerful algorithm for Markov-Chain-Monte-Carlo sampling from the posterior distributions of the VAR parameters. The algorithm works in the presence of mixed frequency or irregularly spaced observations. The posterior is conditioned on data observed at mixed frequencies rather than simply data observed at the coarsest frequency. The method follows from the assumption that the econometrician simply does not observe the high frequency realizations of the low frequency data, and can accordingly treat these data as missing values. Consequently, and consistent with the standard utilization of missing values in Bayesian econometrics, the Bayesian Mixed Frequency (BMF) algorithm developed is a Gibbs sampler that produces alternate draws from the missing data and the unknown parameters in the model. Under typical assumptions about normally distributed exogenous shocks, the VAR's linear structure allows for draws from Gaussian conditional distributions for estimating the missing data, along with draws from Gaussian and inverse Wishart conditional posterior distributions for the parameters in the model. Since this Gibbs sampler requires only simulation from known densities, it is extremely simple to implement.

There has been much work addressing the issue of mixed frequency data from a variety of different approaches. An early contribution is the Kalman filtering approach introduced by Harvey & Pierse (1984), which notes that for linear VAR models, missing observations can be incorporated by simply skipping a term from the updating equation whenever an observation is missing. The VAR's linear and Gaussian form makes it straightforward to formulate a state-space form. However, the Kalman filter approach is potentially cumbersome when the missing data occur at irregular frequencies, especially if there are multiple series with missing data at differing frequencies. In addition, the Kalman filter

yields a likelihood function that is non-linear and non-Gaussian over a potentially very large parameter space; analyzing such likelihood functions often proves difficult both from frequentist and Bayesian viewpoints. The BMF approach, by contrast, handles irregular and multiple missing series with ease, and the Gibbs sampling from standard densities makes the analysis of the resulting posterior densities very tractable.

Another approach, suggested by Miller & Chin (1996), uses monthly data to improve quarterly variable forecasts. The method is an iterative procedure that first uses quarterly observed variables to construct quarterly forecasts, then uses monthly observed variables to construct quarterly forecasts, and finally combines the two forecasts using estimated weights. Corrado & Greene (1988) show that adding monthly information via a monthly pooling procedure can improve quarterly forecasts. The BMF method, in comparison, uses all the relevant information to make multi-frequency forecasting for each variable in the VAR, so in the context of monthly and quarterly data, considers forecasts of monthly variables as well. This difference implies that BMF exploits all the available information to forecast any variable in the model, which offers the advantage of producing additional forecasts but also allows for inference based upon the effects of quarterly variables on monthly ones. Other papers that use bridging type models include Baffigi *et al.* (2004), and those that use bridging with factors, such as Giannone *et al.* (2008) and Angelini *et al.* (2008).

A growing body of work considering the estimation of mixed frequency models is the work on MIDAS (MIXed DATA Sampling) described in Ghysels *et al.* (2004), Andreou *et al.* (2010), Ghysels *et al.* (2007), among others. The MIDAS method allows regressions of a low frequency variable onto high frequency variables. For example, Ghysels *et al.* (2004) study the predictability of stock returns over relative low frequencies (monthly or quarterly) from high frequency volatility estimates, Andreou *et al.* (2009) consider the importance of daily data for forecasting monthly or quarterly real data, and Bai *et al.* (2010) expand MIDAS to deal with state-space models. While the MIDAS approach differs substantially from the Kalman filter approach of Harvey & Pierse (1984), it potentially suffers from the same pitfalls: handling observations that are irregularly spaced requires altering the estimated

equations as in Francis *et al.* (2011), and larger systems may lead to significant numerical burdens.

In contrast to these methods, the approach taken in this paper is from a Bayesian perspective, and will consequently treat lower frequency data as missing. The missing data approach to higher frequency data has a history from both a Bayesian and frequentist perspective. Chow & Lin (1971) discuss how to interpolate time series using related series. Sims & Zha (2006b) and Leeper & Zha (2003), for example, use quarterly GDP interpolated to monthly intervals in monthly VARs. Other mixed-frequency VAR approaches use stock-flow relationships for interpolation, such as Zdrozny (1988), Mittnik & Zdrozny (2005), or Mariano & Murasawa (2010). The BMF approach, on the other hand, follows the Bayesian approach to missing data, similar to, for example, Kim *et al.* (1998).

The traditional approach for dealing with mixed frequency data is to discard high frequency data and simply perform estimation on the coarsest frequency data. This estimation strategy potentially discards information contained in the higher frequency data, yet is used often in macro time series econometrics, especially within the context of VAR estimation, making it a useful benchmark. Indeed, a number of Bayesian and frequentist applications, including studying the effects of monetary policy, oil, or uncertainty shocks, include VAR's estimated at a monthly frequency despite the availability of higher-frequency data. The coarse estimation can be used to identify parameters in the VAR even if the econometrician assumes that the true VAR evolves at some higher frequency than that used for estimation because Gaussian VAR's are closed under temporal aggregation.

In addition to developing the methodology, this paper demonstrates the advantages of the BMF estimation method using numerical simulations and actual data. For numerical simulations over a range of parameter constellations, BMF uniformly dominates estimation using coarse sampling from the frequentist perspective of mean square deviations from the true values. After considering simulated data, two applications highlight the advantages of BMF using actual data. The first involves a monthly and quarterly set of data on the real economy, and the second involves combining monthly real economic variables with high-frequency financial variables. In both contexts, BMF outperforms the coarsely sampled estimator in that the posterior standard deviations are smaller when using BMF. Which

posterior standard deviations decrease the most depend on the application, it can either help accuracy for the low or the high frequency variables. The BMF approach also improves the estimation of impulse response functions, as the decrease in parameter uncertainty associated with BMF is typically reflected in tighter confidence bands for the impulse response functions. Among other things, this result implies that BMF can allow for sharper conclusions about the impact of economic policies or the effects of shocks.

The remainder of the paper is organized as follows: Section 2 discusses the construction of a Gibbs sampler for the model. Section 3 presents simulation based evidence for the advantages of using the BMF approach. Sections 4 and 5 present two examples of applications of mixed frequency estimation; the first example uses a monthly and quarterly model of the macroeconomy, and the second example uses weekly financial data along with monthly data on output. Finally, Section 6 concludes.

2 Econometric Methodology

This section discusses the main algorithm of data augmentation and estimation in the presence of missing data. While the general method applies easily to VAR's with multiple lags, this paper focuses on a first order VAR for simplicity. Consequently, the model is

$$y_t = A + By_{t-1} + \varepsilon_t, \quad \varepsilon_t \sim N(0, \Sigma) \quad (1)$$

where $\dim(y_t) = N$. Denote $y_t = (x_t, z_t)$ where $\dim(x_t) = N_x$ and $\dim(z_t) = N_z$ such that $N_z + N_x = N$ and suppose x_t is a fully observed variable and z_t is a variable with missing data. Again, for simplicity, assume in the following that z and x are recorded at two frequencies, but note that the method applies to a multi-frequency dataset. In the example application in Section 4, x_t is observed monthly and z_t is observed quarterly; Section 5's example has x_t observed weekly and z_t observed monthly. In the case of monthly and quarterly observations, the missing data are $\{\hat{z}_1, \hat{z}_2, \hat{z}_4, \hat{z}_5, \hat{z}_7, \dots\}$, where \hat{z}_t denotes a sampled observation at time t . Let \hat{z} denote the vector of observed and sampled

data, $\hat{z}_{\setminus t}$ denote all elements of \hat{z} except the t -th ones, and let $\hat{Y}^{(i)}$ denote the full collection of observed and sampled data at iteration i .

The BMF estimator is an application of Bayesian Gibbs sampling, which requires iterating over objects of interest, sampling those objects from known distributions conditional upon the remaining objects. In the current setup, the objects of interest are the missing observations, the matrices A and B , and the covariance matrix Σ . Given prior distributions and initial values of the parameters, the i -th iteration of the MCMC algorithm reads

- Step 1 : for $t = 1, \dots, T$, draw missing data $\hat{z}_t^{(i)} \mid x, \hat{z}_{\setminus t}^{(i-1)}, A^{(i-1)}, B^{(i-1)}, \Sigma^{(i-1)}$, where $\hat{z}_{\setminus t}^{(i-1)}$ is the vector of most recently updated missing values and $A^{(i-1)}, B^{(i-1)}, \Sigma^{(i-1)}$ are the latest draws of A, B, Σ , respectively. For example, in the case of consecutive updating, $\hat{z}_{\setminus t}^{(i-1)} = (\hat{z}_1^{(i)}, \hat{z}_2^{(i)}, \dots, \hat{z}_{t-1}^{(i)}, \hat{z}_{t+1}^{(i-1)}, \dots, \hat{z}_T^{(i-1)})$.¹
- Step 2 : draw $A^{(i)}, B^{(i)} \mid \hat{Y}^{(i)}, \Sigma^{(i-1)}$
- Step 3 : draw $\Sigma^{(i)} \mid \hat{Y}^{(i)}, A^{(i)}, B^{(i)}$

The new step in the procedure is Step 1, which is drawing missing data given the parameters in the model and the fully observed data. Except for this first step, the procedure is a standard Normal linear model which yields Normal and Wishart posterior distributions. The next two subsections discuss the algorithm in detail. Since Steps 2 and 3 are familiar in that they are straightforward draws from known distributions, they are addressed first, and then the Step 1 in detail shows how to sample missing data given a set of parameters.

¹For notational simplicity, this notation assumes that all z_t are missing; implicitly the updating equation for non-missing data is just $\hat{z}_t^{(i)} = \hat{z}_t^{(i-1)}$.

In addition, the exact timing of updating is flexible; a possible alternative is to use the entire vector of missing values from the previous iteration, so $\hat{z}_{\setminus t}^{(i-1)} = (\hat{z}_1^{(i-1)}, \hat{z}_2^{(i-1)}, \dots, \hat{z}_{t-1}^{(i-1)}, \hat{z}_{t+1}^{(i-1)}, \dots, \hat{z}_T^{(i-1)})$. However, this timing tends to be less efficient.

2.1 Drawing Parameters

As noted above, conditional on a set of observed and augmented data, Steps 2 and 3 in the Gibbs sampling algorithm are simply draws from Normal and Wishart distributions.

2.1.1 Step 2: Updating the Parameter Matrices A and B

For each iteration, drawing parameters comes from a posterior density composed of the likelihood function and the prior. For notational simplicity, the entire set of data can be treated as fully observed. Consider the following variables:

$$\begin{aligned}
 Y_{N \times (T-1)} &\equiv \begin{bmatrix} x_2 & x_3 & \cdots & x_T \\ z_2 & z_3 & \cdots & z_T \end{bmatrix} \\
 \beta_{N \times (N+1)} &\equiv \begin{bmatrix} A & B \end{bmatrix} \equiv \begin{bmatrix} A_x & B_{xx} & B_{xz} \\ A_z & B_{zx} & B_{zz} \end{bmatrix} \\
 X_{(N+1) \times (T-1)} &\equiv \begin{bmatrix} 1 & 1 & \cdots & 1 \\ x_1 & x_2 & \cdots & x_{T-1} \\ z_1 & z_2 & \cdots & z_{T-1} \end{bmatrix}
 \end{aligned}$$

The likelihood is then

$$\begin{aligned}
 L(\beta|y_T, \dots, y_1, \Sigma) &\propto \prod_{t=2}^T \exp \left\{ -\frac{1}{2} (y_t - A - B y_{t-1})' \Sigma^{-1} (y_t - A - B y_{t-1}) \right\} \\
 &= \exp \left[-\frac{1}{2} \text{tr} \left((Y - \beta X)' \Sigma^{-1} (Y - \beta X) \right) \right] \\
 &\propto \exp \left[-\frac{1}{2} \text{tr} \left((XX') \left(\beta - YX' (XX')^{-1} \right)' \Sigma^{-1} \left(\beta - YX' (XX')^{-1} \right) \right) \right]
 \end{aligned} \tag{2}$$

Proportionality in the last step is obtained by removing all the constant terms and using the trace

properties.²

The conditional posterior for β is therefore

$$p(\beta|y_T, \dots, y_1, \Sigma) \sim MN\left(YX'(XX')^{-1}, (XX')^{-1}, \Sigma\right), \quad (3)$$

a matrix normal density. The vectorized form of β thus has a multivariate Normal density given by

$$p(\text{vec}(\beta)|y_T, \dots, y_1, \Sigma) \sim N(\mu_\lambda, \Sigma_\lambda) \quad (4)$$

where the mean is

$$\mu_\lambda = \text{vec}\left(YX'(XX')^{-1}\right) \quad (5)$$

and the variance is given by

$$\Sigma_\lambda = (XX')^{-1} \otimes \Sigma. \quad (6)$$

Assuming a conjugate prior³

$$\pi(\text{vec}(\beta)) \sim N(\mu_\Omega, \Sigma_\Omega), \quad (7)$$

the posterior density for $\text{vec}(\beta)$ is therefore

$$p(\text{vec}(\beta)|y_T, \dots, y_1, \Sigma) \sim N\left((\Sigma_\lambda^{-1} + \Sigma_\Omega^{-1})^{-1}(\Sigma_\Omega^{-1}\mu_\Omega + \Sigma_\lambda^{-1}\mu_\lambda), (\Sigma_\lambda^{-1} + \Sigma_\Omega^{-1})^{-1}\right). \quad (8)$$

Consequently, conditional on the set of data and a covariance matrix Σ , sampling the parameters β simply requires sampling from a normal distribution.

² $\text{tr}\left(\begin{smallmatrix} ABC & - & DEF \\ n \times n & & n \times n \end{smallmatrix}\right) = \text{tr}\left(\begin{smallmatrix} A & B & C \\ n \times m & m \times k & k \times n \end{smallmatrix}\right) - \text{tr}\left(\begin{smallmatrix} D & E & F \\ n \times m & m \times k & k \times n \end{smallmatrix}\right) = \text{tr}\left(\begin{smallmatrix} CAB \\ k \times k \end{smallmatrix}\right) - \text{tr}\left(\begin{smallmatrix} FDE \\ k \times k \end{smallmatrix}\right) = \text{tr}\left(\begin{smallmatrix} CAB & - & FDE \\ k \times k & & k \times k \end{smallmatrix}\right)$

³The examples later use a zero vector for the prior mean and ten times the identity matrix for the prior variance, which constitutes an uninformative prior.

2.1.2 Step 3: Updating the Covariance Matrix Σ

Given values of the VAR matrices $A^{(i)}$ and $B^{(i)}$ and the data $\hat{Y}^{(i)}$, the conditional posterior for Σ is of the conjugate inverse Wishart form and follows straightforwardly. Let

$$\epsilon_t^{(i)} = Y_t^{(i)} - A^{(i)} - B^{(i)}Y_{t-1}^{(i)} \quad (9)$$

denote the estimated error term at iteration i of the Gibbs sampler. The conditional posterior distribution of Σ is

$$p(\Sigma | Y^{(i)}, A^{(i)}, B^{(i)}) \propto \Sigma^{-T/2} \exp\left(-\sum_t \epsilon_t^{(i)} \Sigma^{-1} \epsilon_t^{(i)'}\right) p(\Sigma) \quad (10)$$

where $p(\Sigma)$ denotes the prior. Using an inverse Wishart prior with mean Ψ and degrees of freedom m the conditional posterior is

$$\Sigma | Y^{(i)}, A^{(i)}, B^{(i)} \sim IW\left(\sum_t \epsilon_t^{(i)} \epsilon_t^{(i)'} + \Psi, T + m\right), \quad (11)$$

another inverse Wishart. Algorithms for sampling from IW distributions are readily available.

2.2 Step 1: Sampling the Latent Data

Step 1 of the Gibbs sampler requires drawing the latent data from its conditional posterior distribution. It is convenient to draw a single t -th element in one operation, so the goal is to draw $\hat{z}_t^{(i)} | x, \hat{z}_{\setminus t}^{(i-1)}, A^{(i-1)}, B^{(i-1)}, \Sigma^{(i-1)}$. Separating the VAR into its components:

$$\begin{bmatrix} x_t \\ z_t \end{bmatrix} = \begin{bmatrix} A_x \\ A_z \end{bmatrix} + \begin{bmatrix} B_{xx} & B_{xz} \\ B_{zx} & B_{zz} \end{bmatrix} \begin{bmatrix} x_{t-1} \\ z_{t-1} \end{bmatrix} + \begin{bmatrix} u_t \\ v_t \end{bmatrix} \quad (12)$$

where

$$\begin{bmatrix} u_t \\ v_t \end{bmatrix} \sim N\left(0, \begin{bmatrix} \Sigma_{xx} & \Sigma_{xz} \\ \Sigma_{zx} & \Sigma_{zz} \end{bmatrix}\right). \quad (13)$$

The Appendix shows that the conditional density for z_t is the multivariate normal

$$\hat{z}_t \mid \hat{z}_{\setminus t}, x, \Theta \sim N(M, W_1) \quad (14)$$

where the mean is given by

$$M \equiv W_1^{-1} W_2 \quad (15)$$

where the covariance matrix is

$$W_1 = \Sigma^{zz} + B'_{zz} \Sigma^{zz} B_{zz} + B'_{xz} \Sigma^{xx} B_{xz} + B'_{zz} \Sigma^{zx} B_{xz} + B'_{xz} \Sigma^{xz} B_{zz} \quad (16)$$

and

$$\begin{aligned} W_2 = & -\Sigma^{zx} (x_t - A_x - B_{xx}x_{t-1} - B_{xz}z_{t-1}) \\ & + \Sigma^{zz} (A_z + B_{zx}x_{t-1} + B_{zz}z_{t-1}) \\ & + B'_{xz} \Sigma^{xx} (x_{t+1} - A_x - B_{xx}x_t) \\ & + B'_{zz} \Sigma^{zx} (x_{t+1} - A_x - B_{xx}x_t) \\ & + B'_{xz} \Sigma^{xz} (z_{t+1} - A_z - B_{zx}x_t) \\ & + B'_{zz} \Sigma^{zz} (z_{t+1} - A_z - B_{zx}x_t). \end{aligned} \quad (17)$$

and the inverse of Σ is defined as

$$\Sigma^{-1} = \begin{bmatrix} \Sigma^{xx} & \Sigma^{xz} \\ \Sigma^{zx} & \Sigma^{zz} \end{bmatrix}. \quad (18)$$

It is now straightforward to construct Gibbs sampling to draw \hat{z}_t , since it is also conditionally normal. One possibility is to draw the elements in a consecutive order. Another approach is to draw odd and even elements of z alternately, which can easily be implemented in a vectorized programming environment.

2.3 Coarse Sampling Estimation

The standard approach to mixed frequency estimation is to delete the high frequency data such that the VAR is estimated at whichever frequency is jointly available. Thus, in estimating a model with, in the case of the example in Section 4, monthly and quarterly data, one would sample both variables at the quarterly frequency. In choosing a quarterly sampling frequency for the monthly data, one throws away information contained in the higher frequency data.

It should be noted that, in the context of many macroeconomic applications that use mixed frequency data, many simply perform estimation at the lowest frequency. Within the literature on monetary policy shocks, papers such as Christiano *et al.* (1996), Christiano *et al.* (1999), Sims & Zha (2006a), Sims & Zha (2006b), and Banbura *et al.* (2010), ignore the high frequency movements in interest rates and financial variables and estimate monthly or quarterly VARs. In studying the effects of oil price shocks, Kilian & Park (2009), Kilian *et al.* (2009), and Kilian & Vigfusson (2011) discard information in high-frequency price movements, again estimating using a monthly frequency. Monthly estimation of VARs is also used by Bloom (2009) to study the effects of uncertainty shocks, even though some of the relevant asset pricing data are available at much higher frequencies. In each of these applications, discarding data at high frequencies and estimating using the lowest sampling interval is standard procedure.

The estimator based solely on coarse data is not an unreasonable estimator. In particular, it is true that the estimator can be used to estimate the true values of the parameters in a VAR even if the true VAR evolves at a higher frequency than that used for estimation. This fact follows because the model $Y_t = A + BY_{t-1} + \epsilon_t$ is closed under temporal aggregation, so that $Y_{t+n} = A_n + B_n Y_t + \epsilon_{t_n}$ where the new coefficients A_n and B_n are given by

$$A_n = \sum_{s=1}^n B^{s-1} A \tag{19}$$

$$B_n = B^n \tag{20}$$

and the covariance matrix of the error ϵ_{t_n} is

$$\Sigma_n = \sum_{s=1}^n B^{s-1} \Sigma B^{s-1'}. \quad (21)$$

Estimation of the lower frequency VAR produces an estimate of B_n , which can then produce an estimate \hat{B} based on some n^{th} root of the estimate \hat{B}_n . Then, using (19) and (21) produces estimates of A and Σ . If, for example, B_q is an estimate of B obtained using quarterly sampled data, inverting (20) gives the implied monthly estimate $B_q^{\frac{1}{3}}$. Equations (19)-(21) thus allow the comparison of estimates obtained through BMF with those produced by coarse estimation. Such comparisons are best done by transforming the posterior simulations and then computing the quantity of interest. For example, to compare posterior standard deviations from coarse estimation at a quarterly frequency to BMF at a monthly one, the sample standard deviation should be computed based upon the converted draws of the Gibbs sampler, so computing $\frac{1}{G} \sum_{i=1}^G (B^{(i)})^{\frac{1}{3}}$ for G iterations. This conversion gives an estimate of the monthly implied standard deviation from quarterly coarse estimates for easy comparisons across the two methodologies.

3 Simulation Results

Having presented the methodology, this section examines BMF using simulated data. The purpose is to analyze how BMF fares relative to estimation at the coarsest frequency when the objective is to recover parameter estimates, say the posterior mean, that are as close as possible in some sense to the truth. This is very much a frequentist way of thinking, and so the exercise should accordingly be interpreted as a small sample study of the posterior mean as a frequentist parameter estimate.

Table 1 reports the root mean square error of the parameters estimated using BMF versus estimation after discarding the high frequency data for four different parameter constellations, each with two sample sizes. The parameterizations use two different coefficient matrices and two different covariances among the errors. The data are generated by a monthly bivariate VAR, where the variable

x is observed every period, and the variable z is observed every third period $t = 1, 4, 7, \dots, 3T$, so T is the number of quarters in the sample, meaning there are $3T$ months. In the shorter sample, there are $T = 20$ quarters, and in the longer sample there are $T = 80$.

As can be seen, BMF attains smaller root mean square error for all the parameter estimates regardless of sample size, correlation and diagonal elements of matrix B . Several additional features should be mentioned. First, *ceteris paribus*, BMF improvements are higher for small sample than for larger sample estimation. In this case the quarterly data provide limited information about the VAR coefficients, and adding monthly data to the estimation produces sizable gains.

Second, for small sample estimation, the smaller the correlation of x and z and the smaller the diagonal elements of B , the better performance of BMF is achieved. As can be seen in the first block of Table 1, for a sample size of $T = 20$ and correlation of zero, the improvement of BMF goes from 17.5% to 72.8%. However, for the same sample size and B , but with a correlation of 95%, the improvement of BMF goes from 9.43% to 26.4%. We can see similar results when the diagonal elements of B are increased from 0.90 to 0.99. BMF is still an improvement but the relative gains in root mean square error are smaller, running 4.18% to 33% for the case of no correlation and from 2.16% to 10.1% for a correlation of 95%. Finally, with high diagonal elements of B and bigger sample size, BMF has a better performance but relatively smaller than in the previous cases. However, it is still important in many cases. Summarizing, Table 1 shows that BMF is always an improvement. The improvement is higher for the smaller sample size and the smaller correlation between the variables of the system.

Hence, this section shows that BMF has significant gains compared to the usual estimation strategy of using the coarsest frequency data in terms of root mean squared errors. Having provided a comparison on simulated data, the following two sections illustrate two experiments widely studied in the economic literature: first, estimating GDP at monthly frequency from monthly data, and second, exploring the information contained in financial variables for real variables.

4 Application I: Monthly and Quarterly Data

4.1 Data Description

This section provides an example implementation of the algorithm to data collected at monthly and quarterly frequencies. The primary objective is to formulate a model that allows analysis of GDP at a frequency higher than the quarterly data readily available. Since GDP is a widely cited indicator of economic performance, using higher frequency data may help to provide more timely snapshots of the economy than quarterly data allows.

This application is of considerable interest in the mixed frequency estimation literature. Kuzin *et al.* (2010) use quarterly GDP and a set of monthly indicators to compute monthly estimates of quarterly GDP in the Euro area following the MIDAS estimation strategy, Mittnik & Zadzorny (2005) pursue a similar objective using the Kalman Filter. Aruoba *et al.* (2009) use a dynamic factor model to estimate economic activity. Barhoumi *et al.* (2008) compare several estimation strategies that generate monthly GDP estimates using quarterly GDP and monthly indicators, Diron (2008) assesses the ability of real-time monthly data to help forecast quarterly GDP, and Marcellino & Musso (2010) study the estimation of real-time output gaps.

To keep the exercise simple and transparent, this application combines monthly data on industrial production, inflation, and the unemployment rate with quarterly real GDP data for the US. The data are the twelve-month change in industrial production and inflation, the four-quarter change in real GDP, and the unemployment rate, all expressed as percentage points. For the mixed frequency data, the timing assumption is that every month, the monthly data are observed, but the quarterly data are observed only during the last month of each quarter. Since the timing assumes that a period equals a month, the analysis below converts the quarterly results to a monthly frequency using the method described in Section 2.3. The data run from Jan-1948 to Jun-2011, for a total of 762 months or 254 quarters. Summary statistics for the variables are presented in Table 2.

4.2 Estimation Results

Tables 3, 4, and 5 display the estimates of the constant A , the coefficient matrix B , and the variance-covariance matrix Σ , respectively, for the BMF estimator using monthly observations versus the coarse estimation strategy of discarding the first two months of each quarter and therefore using only quarterly observations. In all cases, the ordering of the variables follows that in Table 2: industrial production growth, inflation, unemployment rate, and real GDP growth.

Table 3 shows the estimates of the constant terms in the VAR, converted to a monthly frequency. The two estimates produce somewhat different posterior means for the constant terms, but they are all within one standard deviation across the methods. Noticeably, however, the posterior standard deviations are uniformly smaller for the monthly BMF than for the quarterly estimates, so BMF produces more precise estimates.

Table 4 shows the estimates of the coefficient matrices B , again with the quarterly estimates converted to their monthly counterpart. Similar to the estimates for the constant in Table 3, the posterior means across methods are close with the two methods. Again, the standard deviations generated by BMF are uniformly smaller than those produced by quarterly estimation, which reflects the fact that including monthly data provides more information about the persistence of the process.

Figure 1 displays the kernel density estimates of the marginal posterior densities for the implied monthly VAR coefficients using the monthly BMF and versus quarterly estimation. As noted in Table 4, the point estimates do not vary greatly across methods, the bigger difference is in the posterior standard deviations. The posteriors using the monthly BMF have sharper peaks, reflecting the lower standard deviations. The most notable differences come in the fourth row and fourth column of plots, which are the estimates associated with the infrequently observed quarterly real GDP. So BMF improves the precision of the estimates for all variables, but has the biggest difference for the coarsely observed quarterly GDP.

Table 5 reports estimates of the residual covariance matrix Σ from monthly BMF and from quarterly estimation. In this case, there is a very substantial reduction in the posterior standard deviations for

each of the coefficients. In general, the decrease in posterior standard deviations is the most noticeable for the monthly variables, showing that including the higher frequency data tends to produce more precise estimates of the covariances to the shocks.

4.3 Impulse Responses

The previous subsection discussed the improvement in estimates from incorporating the higher frequency data, but the question remains whether the lower standard deviation translates to changes in dynamics implied by the VAR. Figure 2 displays a comparison of impulse response functions between the monthly BMF estimate and the quarterly estimate converted to its monthly counterpart. The plots are very similar for the most part, suggesting that both estimates yield similar responses of the economy to the exogenous shocks. In most cases, the responses are almost identical, with both methodologies yielding humped-shaped responses for certain variables and monotonic responses from others. The cases that differ the most, which are the effects of GDP shocks, have generally the same shape, and the responses are within each others' confidence bands.

Despite the similarities in the shapes of the impulse responses between the two estimates, the accompanied confidence bands suggest that the responses produced using BMF are often more precise than those produced with just quarterly data. The most marked increase in precision occurs in identifying the effects of the various shocks on the monthly variables. Since these are the most frequently observed data, the ability to include these more frequent data in the estimation using BMF produces uniformly tighter confidence bands. For the shocks to the monthly variables, the decrease in the size of the confidence bands is relatively small, which reflects the fact that since those data are observed at the same frequency, adding more frequent observations doesn't have as much of an impact. The more significant reduction in the width of the confidence bands occurs for the effects of GDP on the monthly variables. With the coarsely sampled estimation method, since the data are observed quarterly, comovements of the variables, and hence the effects of shocks, are relatively difficult to distinguish. On the other hand, with BMF the shocks are more easily traced out since the comovements of variables are

observed at a monthly frequency. Hence it is not surprising that the improvement of BMF is strongest among the variables for which BMF allows more observations.

The use of the BMF estimates, therefore, serves as a significant improvement when analyzing impulse response functions. The actual response functions differ only marginally, but the confidence bands for the BMF estimates suggest stronger estimates of the effects of exogenous shocks. The inclusion of the additional, more frequent observations that the BMF estimate makes it much easier to identify the responses of those more frequent observations.

5 Application II: Weekly and Monthly Data

5.1 Data Description

The application in the previous section showed how to use BMF to combine monthly and quarterly data, this application turns to using higher frequency financial data to inform inference about aspects of the real economy. Financial data are often available on an almost continual basis: asset prices and interest rates change even by the minute. This example examines how high-frequency data may help investigate the impact of financial variables or asset prices on output. As noted above, industrial production is available at a monthly frequency, and measures output in a set of subsectors in the economy. Since these production sectors may be especially influenced by changes in interest rates or oil prices, the high frequency data are measures of the level and slope of the yield curve, as well as spot oil prices.

Interest rate conditions may affect production decisions, and oil is often an essential input to production, so it is natural to consider these variables along with IP. A number of papers such as Kilian & Park (2009), Kilian *et al.* (2009) or Kilian & Vigfusson (2011), to name a few, estimate monthly VARs in order to study the affects of oil shocks, even though the spot oil price changes much more frequently. So the analysis uses the spot real price of West Texas Intermediate crude oil, and a measure of the intercept and slope of the yield curve of interest rates. The slope of the yield curve

is defined as the difference in yields between the seven- and one-year zero coupons. The slope gives expectations about future interest rates and because it has been frequently noted that an inverted yield curve (negative slope) tend to precede recessions. All constant maturity zero-coupon yield data are from the dataset by Gurkaynak *et al.* (2006). While these variables are available at extremely high frequencies, the analysis below focuses on weekly data. The data run from the first week of Jan-1986 to the last week of Jul-2011, for a total of 1336 weekly observations and 307 monthly observations. Summary statistics for the variables are presented in Table 6.

In addition to being able to address the impact of interest rates and oil on industrial production, the choice of weekly intervals presents an interesting challenge for mixed-frequency data. The assumption of timing is the following: the last business day of each week (usually Friday but occasionally Thursday), the yield curve and oil spot prices are observed, and the last Friday (or Thursday) of each month, the twelve-month growth rate of industrial production is observed. The challenge is that most months will have four weekly observations per month, but there will be some months that have five weeks associated with them. While BMF can handle this irregularly observed data with ease, using a method such as the Kalman filter or MIDAS would require either ignoring the fifth week in these months or changing the structure of the estimated equations in these months. Since the base period considered is a week, the analysis below converts the monthly estimates to their weekly counterparts following the method described in Section 2.3.⁴

5.2 Estimation

Tables 7 and 8 display the estimates of the constant A and the coefficient matrix B , respectively, for the BMF estimator using weekly observations and versus the discarding all but the last week of the month and therefore using only monthly observations. In all cases, the ordering of the variables

⁴In practice, the conversion of irregularly spaced data can be cumbersome, since, in this case, there are not a fixed number of weeks per month, and therefore equations (19), (20), and (21) cannot be directly applied. So to convert the monthly estimates to weekly estimates, the first step is to convert the monthly to their "whole-sample" counterparts (307 months per sample) and then to their weekly counterparts (1 sample per 1336 weeks). Because of the difficulty of inverting (21) in this case, the reported results here focus on the estimates of A and B .

follows that in Table 6: yield curve intercept, yield curve slope, oil price, and industrial production.

Table 7 contains the estimates of the constant terms in the VAR, converted to a weekly frequency. Three of the posterior means are similar across methods, the third variable, oil, shows a marked change in posterior estimate and reduction in posterior standard deviation, but even this change is small relative to the standard deviations.

Table 8 shows the estimates of the B coefficients, again with the monthly estimates converted to their weekly counterpart. As with the previous application, most of the posterior means are similar across methodologies, and BMF tends to have smaller posterior standard deviations. The notable exception to the similar posterior means are the estimates associated with oil, the third variable in the VAR. In the previous example, BMF provided the biggest reduction in standard errors for the coarsely observed GDP, but here the biggest reduction is associated with the finely observed oil price. The reason for this result is apparent from comparing the descriptive statistics in Table 6, which show that the oil variable has by far the biggest standard deviation and the lowest autocorrelation by a significant margin. Consequently, when oil is included at a high frequency, the inclusion adds a lot more information about the dynamics of the VAR at a weekly interval.

Figure 3 displays the kernel density estimates of the marginal posterior densities for the implied weekly VAR coefficients using the weekly BMF versus when only monthly data are used. The results seen in Table 8 are striking in these plots: the posterior estimates tend to be fairly similar, but there is significant gain in the precision of estimates associated with the third variable, which is the spot price of oil. Moving from monthly data to weekly has a large gain in precision of the posterior estimates associated with oil, as reflected by the high peaks of the marginal posterior in the third column and third row of subplots.

5.3 Impulse Responses

After noting the gains in accuracy from the parameter estimates, especially for the oil variable, Figure 4 shows the effects of incorporating weekly data on the impulse response functions. As with the previous

application, the confidence bands are smaller for the BMF estimator. And reflecting the smaller standard errors associated with the parameters corresponding to the effects of oil, the narrowing of confidence bands is most pronounced for the oil variables. The shapes of the impulse responses are nearly identical across the two methodologies, but slight shifts in impulse responses from the weekly data may lead to different conclusions. For example, the "Oil on Oil" response is not very persistent for either case, but the relative persistence is greater for BMF. In addition, the BMF confidence intervals for oil cover zero in the case of "Oil on Level" and "Oil on Slope" meaning that the significance of the shocks may be reversed with the inclusion of higher frequency data.

6 Concluding Remarks

This paper considers estimation of first order VAR's using data sampled at mixed frequencies. The methodology uses Gibbs sampling the unobserved data at the high frequency to generate estimates with generally smaller standard errors. The simulation experiments demonstrate that BMF produces more accurate estimates of model parameters than the basic approach of sub-sampling at the coarse data frequency, and the two example applications show that using higher frequency data may produce sizable gains.

Improved accuracy is not the only advantage of the BMF estimator. Another benefit is the ability to update forecasts of a coarsely observed variable in response to new arrival of data measured at high frequencies. Along the lines of the applications presented above, examples include updating forecasts of next quarter GDP in response to monthly measurements of data or using weekly or even daily financial data to forecast aspects of the real economy. The BMF framework allows for a natural approach to incorporate high frequency observations to the low frequency forecasts, which would avoid the use of ad-hoc forecast revisions.

The approach is implemented here using a first order vector auto-regression for simplicity, but there are many possible extensions that generalize this approach. First, higher order VAR's can be implemented straightforwardly as the conditional distribution for the unobserved data is still

conditionally normal, although the moments of these distributions involve some tedious algebra.

One potential advantage of the Bayesian simulation approach is that it easily generalizes to more complicated models. For example, both financial market and macro variables are known to exhibit time-varying volatility. Following Jacquier *et al.* (1994), incorporating stochastic volatility into a VAR setting simply requires adding an additional Gibbs sampling draw for the unobserved volatility path. In a similar vein, the Bayesian approach can easily incorporate heavy tailed error distributions in the form of mixture of normals or t distributions, Markov mixtures as in Albert & Chib (1993) or jumps as in Eraker *et al.* (2003).

It is useful to consider how BMF compares to a Kalman filtering approach. The linearity of the VAR coupled with Gaussian errors implies that it is possible to write the unobserved high frequency data in a state space form, enabling the use of a Kalman filter approach to a basic model. One advantage of the Kalman filtering approach is that it does not require posterior simulation of unobserved data. On the other hand, the resulting likelihood function needs to be analyzed numerically. Frequentist analysis through maximum likelihood is possible, although this requires a numerical search in a typically high dimensional parameter space. For Bayesian inference, BMF avoids using Metropolis Hastings or other methods that would be required in to deal with the non-Gaussian nature of the posterior computed through Kalman filtering.

The application to weekly and monthly data in Section 5 highlights a second advantage of BMF over the Kalman filter and MIDAS. In the case of irregularly spaced data, where there are an unequal number of observations per unit of coarsely observed data, the specification of the Kalman filter and MIDAS equations can be awkward and tedious, whereas BMF requires no such changes. For the Kalman filter approach, the state equation is the VAR equation

$$\begin{bmatrix} x_t \\ z_t \end{bmatrix} = \begin{bmatrix} A_x \\ A_z \end{bmatrix} + \begin{bmatrix} B_{xx} & B_{xz} \\ B_{zx} & B_{zz} \end{bmatrix} \begin{bmatrix} x_{t-1} \\ z_{t-1} \end{bmatrix} + \begin{bmatrix} u_t \\ v_t \end{bmatrix} \quad (22)$$

and the observation equation changes from

$$Y_t^{obs} = \begin{bmatrix} I & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} x_t \\ z_t \end{bmatrix} \quad (23)$$

if z_t is observed, to

$$Y_t^{obs} = \begin{bmatrix} I & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} x_t \\ z_t \end{bmatrix} \quad (24)$$

if z_t is not observed. In the case of mixed frequency data where the missing data occur at regular frequencies, such as the monthly and quarterly application, the observation equation switches between (24) and (23) systematically. With irregularly spaced data, such as the weekly and monthly application, the switching observation equation changes at different intervals, making implementation more difficult. With even more irregular observations, such as if z_t has variables that are observed not at the same frequency, the complexity in changing the observation equation can grow substantially. The BMF approach does not require this constant changing of observation equations, since it simply draws all the missing data conditional on the parameters regardless of the observation interval.

A potential disadvantage of the VAR approach is that, especially as the number of lags grows, the number of parameters to estimate grows considerably. The MIDAS approach, by comparison, typically focuses on one equation to forecast one coarsely observed variable with more frequently observed variables. Considering this setup for the application of weekly and monthly data from Section 5, the MIDAS approach forecasts monthly industrial production using the weekly interest rate and oil data without putting restrictions on how industrial production affects the weekly data. However, BMF can incorporate this flexibility by including certain variables as exogenous rather than in the endogenous vector in the VAR. Then, for example, interest rates and IP could be affected by the exogenous evolution of oil, thereby ignoring other feedback effects on oil. Of course, whether to incorporate variables as endogenous or exogenous depends upon the application and research question, but the important point is that BMF handles these restrictions with minimal modification to the VAR

model considered here.

Another useful generalization of the VAR approach is to consider the use of BMF in connection with linear state-space models. Consider the case with an observation equation $Y_{t+1} = C + DX_{t+1}$, and state-equation $X_{t+1} = A + BX_t + \epsilon_{t+1}$, where there is a mixed sampling frequency for Y_t . The VAR(1) framework discussed in this paper nests the state-space model by writing $Y^* = (Y, X)$, and that $Y_{t+1}^* = A^* + B^*Y_t^* + \epsilon_{t+1}^*$ and $A^* = (C, A)$,

$$B^* = \begin{pmatrix} 0 & D \\ 0 & B \end{pmatrix}. \quad (25)$$

BMF then proceeds to estimate this model by simulating, as before, the sparsely observed elements of Y , but in addition treats X as an unobserved variable – a variable observed with zero frequency. Importantly, the algorithm for drawing the missing data applies directly in this setting. To proceed to the second step of the Gibbs sampler which involves drawing the parameters, the algorithm needs only slight modifications to impose the zero-constraints on B^* . Note that the estimation of VARMA models can be implemented using this approach.

This paper has also not considered the out-of-sample forecasting ability of BMF estimators. Of course, given the applications, forecasting is a natural extension given mixed frequency data. As with any VAR-based method, forecasting given BMF estimates involves iterated forecasting rather than direct forecasting. Marcellino *et al.* (2006) and Chevillon & Hendry (2005) discuss the advantages of both types of forecasting, and De Mol *et al.* (2008) show Bayesian VARs can have good forecasting performance.

Finally, while the BMF algorithm applies in general, identification considerations must, as usual, be investigated on a case by case basis depending upon the application. Consequently, the BMF framework developed in this paper represents an interesting starting point for a number of different extensions.

References

- ALBERT, J.H., & CHIB, S. 1993. Bayes Inference via Gibbs Sampling of Autoregressive Time Series Subject to Markov Mean and Variance Shifts. *Journal of the American Statistical Association*, **11**, 1–15.
- ANDREOU, E., GHYSELS, E., & KOURTELLOS, A. 2009. Should Macroeconomic Forecasters Use Daily Financial Data and How? *UNC Working Paper*.
- ANDREOU, E., GHYSELS, E., & KOURTELLOS, A. 2010. Regression Models with Mixed Sampling Frequencies. *Journal of Econometrics*, **158**, 246–261.
- ANGELINI, E., CAMBA-MENDEZ, G., GIANNONE, D., REICHLIN, L., & RUNSTLER, G. 2008. Short-Term Forecasts of Euro Area GDP Growth. *ECB Working Paper*, **949**.
- ARUOBA, S.B., DIEBOLD, F.X., & SCOTTI, C. 2009. Real-Time Measurement of Business Conditions. *Journal of Business and Economic Statistics*, **27**(4), 417–427.
- BAFFIGI, A., GOLINELLI, R., & PARIGI, G. 2004. Bridge Models to Forecast the Euro Area GDP. *International Journal of Forecasting*, 447–460.
- BAI, J., GHYSELS, E., & WRIGHT, J.H. 2010. State Space Models and MIDAS Regressions. *UNC Working Paper*.
- BANBURA, M., GIANNONE, D., & REICHLIN, L. 2010. Large Bayesian VARs. *Journal of Applied Econometrics*, 71–92.
- BARHOUMI, K., BENK, S., CRISTADORO, R., DEN REIJER, A., JAKAITIENE, A., JELONEK, P., RUA, A., RUNSTLER, G., RUTH, K., & NIEUWENHUYZE, C.V. 2008. Short-term Forecasting of GDP Using Large Monthly Datasets - A Pseudo Real-time Forecast Evaluation Exercise. *National Bank of Belgium Working Paper*.
- BLOOM, N. 2009. The Impact of Uncertainty Shocks. *Econometrica*, **77**(3), 623–685.

- CHEVILLON, G., & HENDRY, D. F. 2005. Non-Parametric Direct Multi-Step Estimation for Forecasting Economic Processes. *International Journal of Forecasting*, **21**, 201–218.
- CHOW, G.C., & LIN, A. 1971. Best Linear Unbiased Interpolation, Distribution, and Extrapolation of Time Series by Related Series. *The Review of Economics and Statistics*, **53**(4), 372–375.
- CHRISTIANO, L., EICHENBAUM, M., & EVANS, C. 1996. The Effects of Monetary Policy Shocks: Evidence from the Flow of Funds. *The Review of Economics and Statistics*, 16–34.
- CHRISTIANO, L., EICHENBAUM, M., & EVANS, C. 1999. Monetary Policy Shocks: What Have We Learned, and To What End. *Handbook of Monetary Economics*.
- CORRADO, C., & GREENE, M. 1988. Reducing Uncertainty in Short-Term Projections: Linkage of Monthly and Quarterly Models. *Journal of Forecasting*, **7**, 77–102.
- DE MOL, C., GIANNONE, D., & REICHLIN, L. 2008. Forecasting Using a Large Number of Predictors: Is Bayesian Shrinkage a Valid Alternative to Principal Components? *Journal of Econometrics*, **146**, 318–328.
- DIRON, M. 2008. Short-Term Forecasts of Euro Area Real GDP Growth: An Assessment of Real-Time Performance Based on Vintage Data. *Journal of Forecasting*, **27**(5), 371–390.
- ERAKER, B., JOHANNES, M. J., & POLSON, N. G. 2003. The Impact of Jumps in Returns and Volatility. *Journal of Finance*, **53**, 1269–1300.
- FRANCIS, N.R., GHYSELS, E., & OWYANG, M.T. 2011. The Low-Frequency Impact of Daily Monetary Policy Shocks. *Federal Reserve Bank of St. Louis Working Paper*, **2011-009B**.
- GHYSELS, E., SANTA-CLARA, P., & VALANKOV, R. 2004. There is a Risk-Return Tradeoff After All. *Journal of Financial Economics*.
- GHYSELS, E., SINKO, A., & VALANKOV, R. 2007. MIDAS Regressions: Further Results and New Directions. *UNC Working Paper*.

- GIANNONE, D., REICHLIN, L., & SMALL, D. 2008. Nowcasting: The Real-Time Informational Content of Macroeconomic Data. *Journal of Monetary Economics*, **55**(4), 665–676.
- GURKAYNAK, R.S., SACK, B., & WRIGHT, J. 2006. The U.S. Treasury Yield Curve: 1961 to the Present. *Federal Reserve Board Working Paper*.
- HARVEY, A.C., & PIERSE, R.G. 1984. Missing Observations in Economic Time Series. *Journal of the American Statistical Association*, **79**, 1–131.
- JACQUIER, E., POLSON, N. G., & ROSSI, P. 1994. Bayesian Analysis of Stochastic Volatility Models. *Journal of Business and Economic Statistics*, **12**, 371–389.
- KILIAN, L., & PARK, C. 2009. The Impact of Oil Price Shocks on the US Stock Market. *International Economic Review*, **50**(4), 1267–1287.
- KILIAN, L., & VIGFUSSON, R.J. 2011. Nonlinearities in the Oil Price-Output Relationship. *Macroeconomic Dynamics*.
- KILIAN, L., REBUCCI, A., & SPATAFORA, N. 2009. Oil Shocks and External Balances. *Journal of International Economics*, **77**(2), 181–194.
- KIM, S., SHEPHARD, N., & CHIB, S. 1998. Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models. *Review of Economic Studies*, **65**(3), 361–393.
- KUZIN, V., MARCELLINO, M., & SCHUMACHER, C. 2010. MIDAS vs. mixed-frequency VAR: Nowcasting GDP in the Euro Area. *International Journal of Forecasting*.
- LEEPER, E., & ZHA, T. 2003. Modest Policy Interventions. *Journal of Monetary Economics*, **50**(8), 1673–1700.
- MARCELLINO, M., & MUSSO, A. 2010. Real Time Estimates of the Euro Area Output Gap: Reliability and Forecasting Performance. *ECB Working Paper*, **1157**.

- MARCELLINO, M., STOCK, J. H., & WATSON, M. W. 2006. A Comparison of Direct and Iterated Multistep AR Methods for Forecasting Macroeconomic Time Series. *Journal of Econometrics*, **135**, 499–526.
- MARIANO, R., & MURASAWA, Y. 2010. A Coincident Index, Common Factors, and Monthly Real GDP. *Oxford Bulletin of Economics and Statistics*, **72**, 27–46.
- MILLER, P., & CHIN, D. 1996. Using Monthly Data to Improve Quarterly Model Forecasts. *Federal Reserve Bank of Minneapolis Quarterly Review*, 16–33.
- MITTNIK, S., & ZADROZNY, P. 2005. Forecasting Quarterly German GDP at Monthly Intervals Using Monthly Ifo Business Conditions Data. *Ifo Survey Data in Business Cycle and Monetary Policy Analysis*, 19–48.
- SIMS, C.A., & ZHA, T. 2006a. Does Monetary Policy Generate Recessions? *Macroeconomic Dynamics*, **10**(2), 231.
- SIMS, C.A., & ZHA, T. 2006b. Were There Regime Switches in US Monetary Policy? *American Economic Review*, **96**(1), 54–81.
- ZADROZNY, P. 1988. Gaussian Likelihood of Continuous-Time Armax Models When Data are Stocks and Flows at Different Frequencies. *Econometric Theory*, 108–124.

7 Appendix

This appendix shows the derivations for the analytical conditional distribution of the latent data

$$p\left(z_t|x_t, x_{t-1}, x_{t+1}, z_{t+1}, z_{t-1}, \Theta^{(i-1)}\right) \quad (26)$$

as shown in (14). First, the conditional distribution is proportional to the product of the densities of the data at $t + 1$ given data at time t , $p(x_{t+1}, z_{t+1}|x_t, z_t, \Theta)$ and the data at t given the data at $t - 1$,

$p(x_t, z_t | x_{t-1}, z_{t-1}, \Theta)$:

$$p(z_t | x_t, x_{t-1}, x_{t+1}, z_{t+1}, z_{t-1}, \Theta) \propto p(x_{t+1}, z_{t+1} | x_t, z_t, \Theta) p(x_t, z_t | x_{t-1}, z_{t-1}, \Theta). \quad (27)$$

Recall from the VAR structure that the second of these distributions in (27) is conditionally normal:

$$\begin{aligned} p(z_t, x_t | x_{t-1}, z_{t-1}, \Theta) &\propto l(y_t | y_{t-1}, \Theta) \\ &= \exp \left\{ -\frac{1}{2} \begin{bmatrix} v_1 \\ v_2 \end{bmatrix}' \begin{bmatrix} \Sigma_{xx} & \Sigma_{xz} \\ \Sigma_{zx} & \Sigma_{zz} \end{bmatrix}^{-1} \begin{bmatrix} v_1 \\ v_2 \end{bmatrix} \right\} \end{aligned} \quad (28)$$

where the inverse covariance matrix is

$$\Sigma^{-1} \equiv \begin{bmatrix} \Sigma^{xx} & \Sigma^{xz} \\ \Sigma^{zx} & \Sigma^{zz} \end{bmatrix} \quad (29)$$

and

$$v_1 \equiv x_t - A_x - B_{xx}x_{t-1} - B_{xz}z_{t-1} \quad (30)$$

$$v_2 \equiv z_t - A_z - B_{zx}x_{t-1} - B_{zz}z_{t-1} \quad (31)$$

Similarly, the first of the distributions in (27) is also conditionally normal:

$$p(x_{t+1}, z_{t+1} | x_t, z_t, \Theta) \propto \exp \left\{ -\frac{1}{2} \begin{bmatrix} w_1 \\ w_2 \end{bmatrix}' \begin{bmatrix} \Sigma_{xx} & \Sigma_{xz} \\ \Sigma_{zx} & \Sigma_{zz} \end{bmatrix}^{-1} \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} \right\} \quad (32)$$

where

$$w_1 = x_{t+1} - A_x - B_{xx}x_t - B_{xz}z_t \quad (33)$$

$$w_2 = z_{t+1} - A_z - B_{zx}x_t - B_{zz}z_t \quad (34)$$

Combining these implies

$$p(z_t | x_t, x_{t-1}, x_{t+1}, z_{t+1}, z_{t-1}, \Theta) \propto \exp \left\{ -\frac{1}{2} (z_t - M)' W_1 (z_t - M) \right\} \quad (35)$$

where M , W_1 , and W_2 are defined as in (15), (16), and (17). Thus, as asserted in (14), the latent data is conditionally normal:

$$\hat{z}_t | \hat{z}_{\setminus t}, x, \Theta \sim N(M, W_1). \quad (36)$$

Table 1: Accuracy of BMF Using Simulated Data

Root mean squared errors of parameters estimated using BMF and coarse data relative to true values over different sample length and pseudo true parameter values. The simulated model is $y_t = A + By_{t-1} + \varepsilon_t$, where $y_t = [x_t \ z_t]'$, $A = [A_x \ A_z]'$, $B = [B_{xx} \ B_{xz}; B_{zx} \ B_{zz}]$, $\varepsilon_t = [\varepsilon_t^x \ \varepsilon_t^z]'$, $\varepsilon_t \sim (\mathbf{0}, [1 \ \rho; \rho \ 1])$. T is sample size in quarters. The row labeled '*%diff*' gives the percentage difference in BMF relative to using the coarsest data.

Simulation 1.1	$A = [0; 0]; B = [0.9 \ -0.04; 0.04 \ 0.9]; \rho = 0, T = 20$					
	\hat{A}_x	\hat{A}_z	\hat{B}_{xx}	\hat{B}_{xz}	\hat{B}_{zx}	\hat{B}_{zz}
Quarterly	0.000478	0.000766	0.0231	0.0322	0.0701	0.133
BMF:Monthly	0.000395	0.000414	0.0188	0.0184	0.0332	0.0362
<i>%diff</i>	-17.5	-45.9	-18.9	-43	-52.7	-72.8
Simulation 1.2	$T = 80$					
Quarterly	0.00066	0.000748	0.0295	0.0435	0.0935	0.155
BMF:Monthly	0.000502	0.000418	0.0247	0.021	0.0426	0.0387
<i>%diff</i>	-23.9	-44.1	-16.2	-51.7	-54.4	-75.1
Simulation 2.1	$A = [0; 0]; B = [0.9 \ -0.04; 0.04 \ 0.9]; \rho = 0.95, T = 20$					
	\hat{A}_x	\hat{A}_z	\hat{B}_{xx}	\hat{B}_{xz}	\hat{B}_{zx}	\hat{B}_{zz}
Quarterly	0.00117	0.00103	0.0645	0.0661	0.0481	0.0499
BMF:Monthly	0.000862	0.000859	0.0552	0.0549	0.044	0.0452
<i>%diff</i>	-26.4	-16.4	-14.4	-16.9	-8.5	-9.43
Simulation 2.2	$T = 80$					
Quarterly	0.000946	0.00084	0.052	0.0517	0.0392	0.0384
BMF:Monthly	0.000717	0.000724	0.0467	0.0466	0.0374	0.0369
<i>%diff</i>	-24.2	-13.9	-10.1	-9.84	-4.51	-3.98
Simulation 3.1	$A = [0; 0]; B = [0.99 \ -0.04; 0.04 \ 0.99]; \rho = 0, T = 20$					
	\hat{A}_x	\hat{A}_z	\hat{B}_{xx}	\hat{B}_{xz}	\hat{B}_{zx}	\hat{B}_{zz}
Quarterly	0.0029	0.00238	0.0354	0.035	0.0249	0.0258
BMF:Monthly	0.00194	0.00198	0.0324	0.0316	0.0239	0.0245
<i>%diff</i>	-33.2	-16.6	-8.3	-9.83	-4.18	-4.85
Simulation 3.2	$T = 80$					
Quarterly	0.00223	0.00176	0.0259	0.0258	0.0188	0.019
BMF:Monthly	0.00151	0.00152	0.0255	0.0248	0.0189	0.0189
<i>%diff</i>	-32	-13.4	-1.65	-3.89	0.611	-0.349
Simulation 4.1	$A = [0; 0]; B = [0.99 \ 0.004; 0.004 \ 0.99]; \rho = 0.95, T = 20$					
	\hat{A}_x	\hat{A}_z	\hat{B}_{xx}	\hat{B}_{xz}	\hat{B}_{zx}	\hat{B}_{zz}
Quarterly	0.0029	0.00299	0.0589	0.0554	0.057	0.0576
BMF:Monthly	0.00263	0.00269	0.0576	0.0542	0.0554	0.0561
<i>%diff</i>	-9.23	-10.1	-2.16	-2.15	-2.82	-2.69
Simulation 4.2	$T = 80$					
Quarterly	0.00204	0.00194	0.0418	0.0402	0.0406	0.0416
BMF:Monthly	0.00189	0.0018	0.0406	0.0388	0.0396	0.0405
<i>%diff</i>	-7.47	-7.43	-2.96	-3.43	-2.39	-2.48

Table 2: Descriptive Statistics, Monthly and Quarterly Application

Basic statistics for the change in industrial production, the inflation rate, the unemployment rate, and the change in real GDP.

	Mean	Std.Dev.	Autocorrelation
$\Delta \ln(\text{IP})$	3.207	5.961	0.965
Inflation	3.719	3.005	0.988
Unempl	5.744	1.641	0.991
$\Delta \ln(\text{GDP})$	3.249	2.755	0.845

Table 3: Parameter Estimates A , Monthly and Quarterly Application

Posterior means and standard deviations of the VAR constant terms, A . The variables are, in order of appearance, the change in industrial production, the inflation rate, the unemployment rate, and the change in GDP. All parameters are at the monthly frequency.

	BMF:Monthly	Quarterly
A_1 ($\ln(\Delta \text{IP})$)	-0.2348 (0.2697)	-0.3961 (0.3585)
A_2 (Inflation)	-0.05023 (0.0804)	0.02557 (0.107)
A_3 (Unempl)	0.1849 (0.03613)	0.1868 (0.04355)
A_4 ($\ln(\Delta \text{GDP})$)	0.1393 (0.1406)	0.01557 (0.1657)

Table 4: Parameter Estimates B , Monthly and Quarterly Application

Posterior means and standard deviations of the VAR coefficients, B . The variables are, in order of appearance, the change in industrial production, the inflation rate, the unemployment rate, and the change in GDP. All parameters are at the monthly frequency.

		BMF:Monthly	Quarterly
B_{11}	(IP on IP)	1.041 (0.02255)	0.9757 (0.03406)
B_{12}	(Inflation on IP)	-0.09668 (0.01854)	-0.1123 (0.02605)
B_{13}	(Unempl on IP)	0.1662 (0.03687)	0.1793 (0.04953)
B_{14}	(GDP on IP)	-0.1514 (0.05043)	-0.04255 (0.07216)
B_{21}	(IP on Inflation)	0.02818 (0.006634)	0.04136 (0.009629)
B_{22}	(Inflation on Inflation)	0.989 (0.005464)	0.9878 (0.00752)
B_{23}	(Unempl on Inflation)	0.009744 (0.01101)	0.005253 (0.01449)
B_{24}	(GDP on Inflation)	-0.01966 (0.01459)	-0.04549 (0.02092)
B_{31}	(IP on Unempl)	-0.01269 (0.003044)	-0.00731 (0.003926)
B_{32}	(Inflation on Unempl)	0.008113 (0.002467)	0.008263 (0.003029)
B_{33}	(Unempl on Unempl)	0.972 (0.004923)	0.974 (0.005923)
B_{34}	(GDP on Unempl)	-0.001899 (0.006804)	-0.01156 (0.008528)
B_{41}	(IP on GDP)	0.1042 (0.0125)	0.07226 (0.01543)
B_{42}	(Inflation on GDP)	-0.02476 (0.009531)	-0.0313 (0.0118)
B_{43}	(Unempl on GDP)	0.0644 (0.01903)	0.07402 (0.02276)
B_{44}	(GDP on GDP)	0.7696 (0.02761)	0.8286 (0.03299)

Figure 1: Posterior Densities for Monthly and Quarterly Application

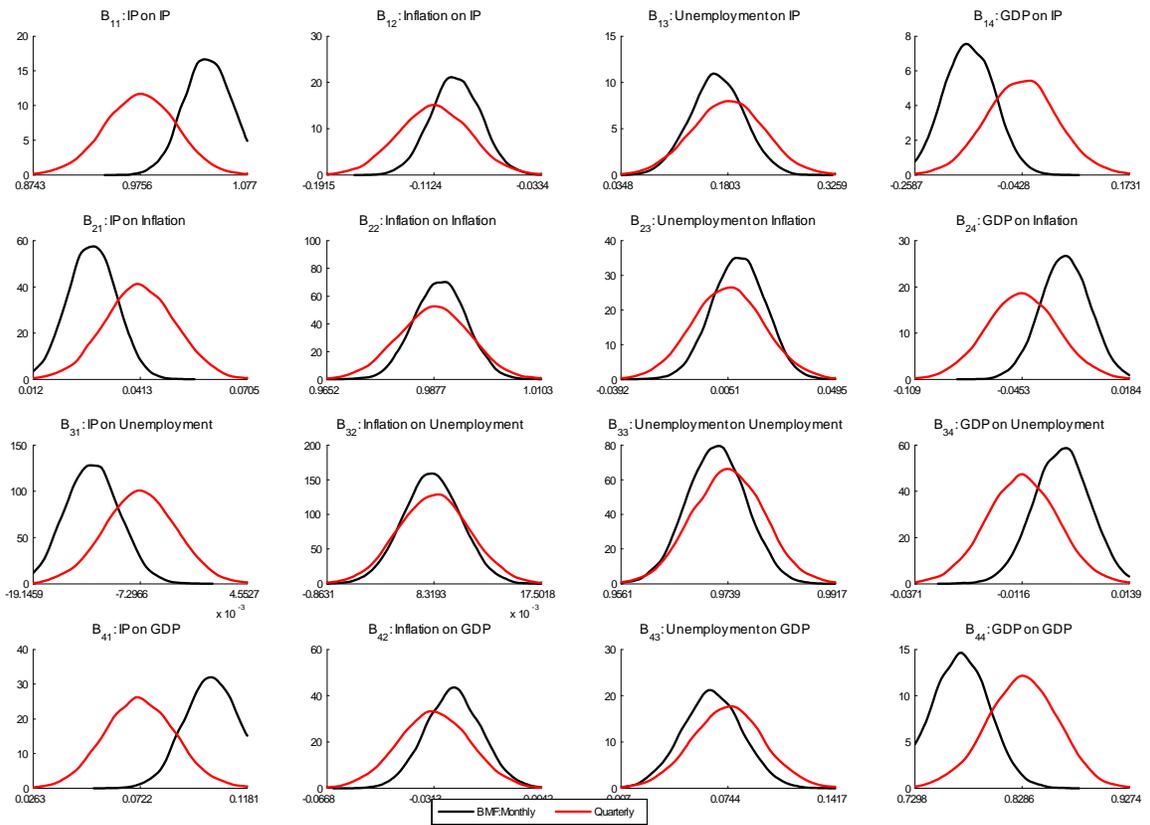


Table 5: Parameter Estimates, Σ , Monthly and Quarterly Application

Posterior means and standard deviations of the covariance matrix Σ . The variables are, in order of appearance, the change in industrial production, the inflation rate, the unemployment rate, and the change in GDP. All parameters are at the monthly frequency.

BMF: Monthly			
2.268 (0.1188)	0.054 (0.02499)	-0.1251 (0.01215)	0.6295 (0.0619)
0.054 (0.02499)	0.2033 (0.01048)	-0.006702 (0.003321)	-0.008757 (0.01574)
-0.1251 (0.01215)	-0.006702 (0.003321)	0.04074 (0.002091)	-0.05687 (0.009803)
0.6295 (0.0619)	-0.008757 (0.01574)	-0.05687 (0.009803)	0.5795 (0.04879)
Quarterly			
3.6802 (0.3647)	0.0435 (0.0727)	-0.2581 (0.0326)	1.0575 (0.1306)
0.0435 (0.0727)	0.3071 (0.0291)	-0.0083 (0.0078)	0.0042 (0.0312)
-0.2581 (0.0326)	-0.0083 (0.0078)	0.0425 (0.0041)	-0.0941 (0.0133)
1.0575 (0.1306)	0.0042 (0.0312)	-0.0941 (0.0133)	0.6850 (0.0669)

Table 6: Descriptive Statistics, Weekly and Monthly Application

Basic statistics for the one-year interest rate, slope of the yield curve, monthly growth rate in the real price of oil, and industrial production.

	Mean	Std.Dev.	Autocorrelation
Level	4.467	2.365	0.999
Slope	5.699	1.862	0.997
$\ln(\Delta\text{Oil})$	0.232	10.039	0.706
$\ln(\Delta\text{IP})$	2.165	4.145	0.997

Figure 2: Impulse Responses for Monthly and Quarterly Application

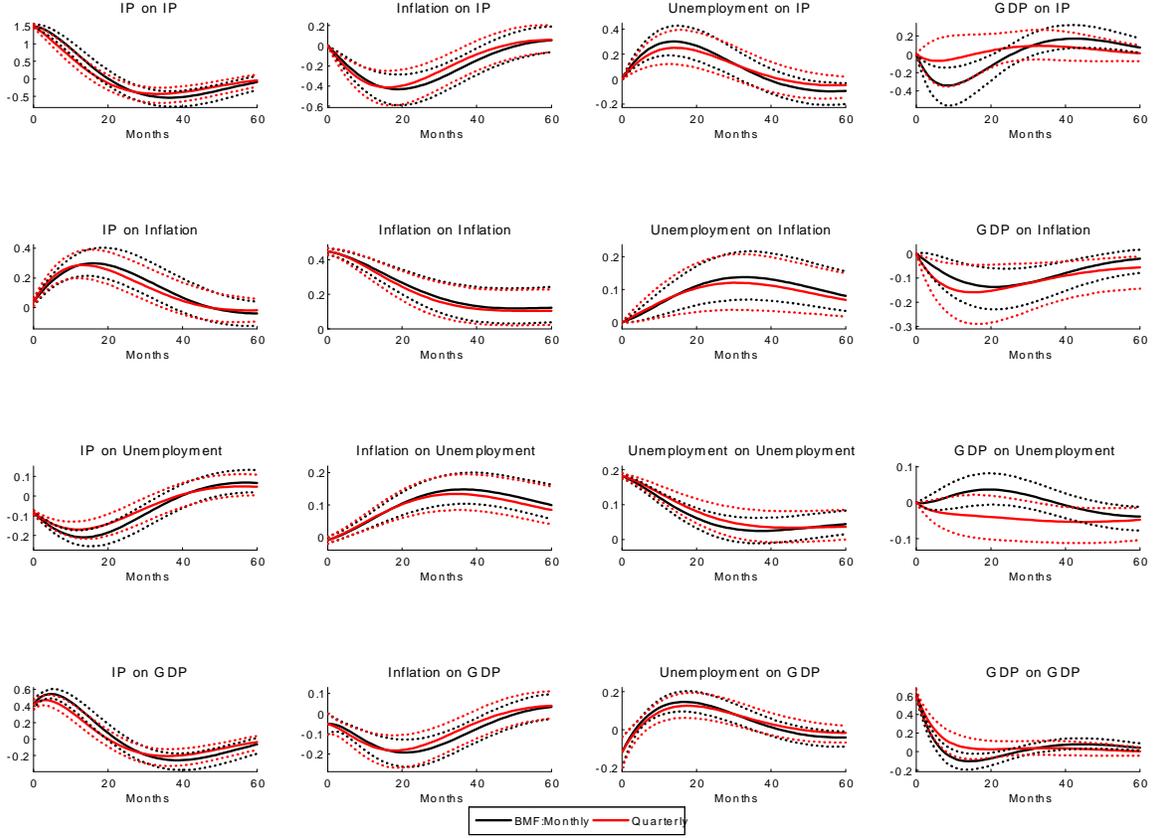


Table 7: Parameter Estimates A , Weekly and Monthly Application

Posterior means and standard deviations of the VAR constant terms, A . The variables are, in order of appearance, the yield curve level, yield curve slope, change in real price of oil, and industrial production growth. All parameters are at the weekly frequency.

		BMF:Weekly	Monthly
A_1	Level	-0.01043 (0.01292)	-0.009877 (0.01478)
A_2	Slope	0.01737 (0.01549)	0.01502 (0.01633)
A_3	$\ln(\Delta\text{Oil})$	0.7281 (0.7478)	0.184 (1.216)
A_4	$\ln(\Delta\text{IP})$	-0.1376 (0.05292)	-0.1097 (0.04865)

Table 8: Parameter Estimates B , Weekly and Monthly Application

Posterior means and standard deviations of the VAR coefficients, B . The variables are, in order of appearance, the yield curve level, slope of the yield curve, change in real price of oil, and industrial production growth. All parameters are at the weekly frequency.

		BMF:Weekly	Monthly
B_{11}	(Level on Level)	0.9948 (0.003563)	0.9949 (0.004036)
B_{12}	(Slope on Level)	0.003914 (0.004396)	0.003801 (0.005023)
B_{13}	(Oil on Level)	0.0004109 (0.0003272)	0.001773 (0.0009227)
B_{14}	(IP on Level)	0.002741 (0.0008671)	0.002648 (0.0009658)
B_{21}	(Level on Slope)	0.004161 (0.004294)	0.004179 (0.004494)
B_{22}	(Slope on Slope)	0.993 (0.005288)	0.9935 (0.005588)
B_{23}	(Oil on Slope)	0.0005827 (0.0003895)	0.002013 (0.001035)
B_{24}	(IP on Slope)	-0.000564 (0.001044)	-0.0006573 (0.001086)
B_{31}	(Level on Oil)	0.1082 (0.2101)	-0.03949 (0.4097)
B_{32}	(Slope on Oil)	-0.2079 (0.2555)	-0.01346 (0.4528)
B_{33}	(Oil on Oil)	0.7065 (0.01933)	0.4741 (0.1265)
B_{34}	(IP on Oil)	0.01108 (0.05147)	0.04647 (0.1123)
B_{41}	(Level on IP)	-0.04702 (0.01466)	-0.04703 (0.01352)
B_{42}	(Slope on IP)	0.06293 (0.01804)	0.05783 (0.01664)
B_{43}	(Oil on IP)	0.003177 (0.001624)	0.007024 (0.003202)
B_{44}	(IP on IP)	0.9945 (0.003479)	0.9965 (0.003295)

Figure 3: Posterior Densities for Weekly and Monthly Application

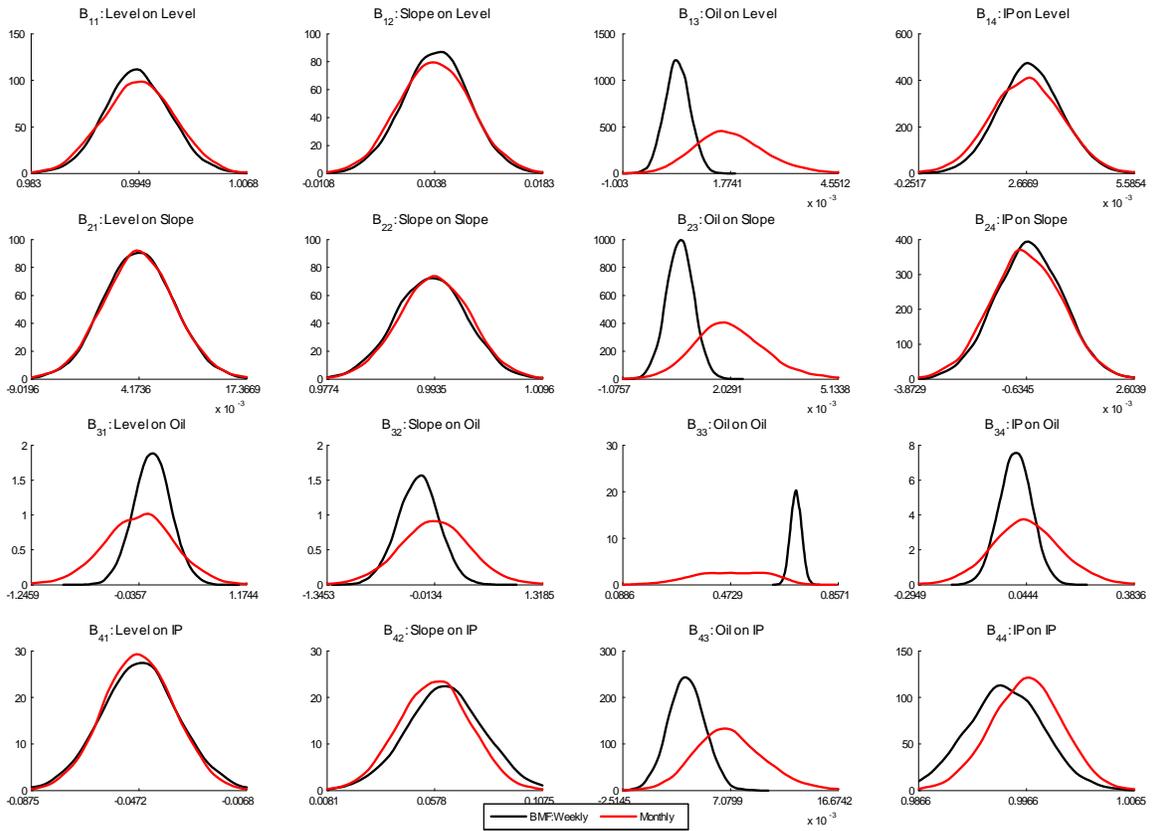


Figure 4: Impulse Responses for Weekly and Monthly Application

