

Real-Time Forecasting With a Mixed-Frequency VAR

Frank SCHORFHEIDE

Department of Economics, University of Pennsylvania, 3718 Locust Walk, Philadelphia, PA 19104
(schorf@ssc.upenn.edu)

Dongho SONG

Department of Economics, Boston College, 140 Commonwealth Avenue, Chestnut Hill, MA 02467
(dongho.song@bc.edu)

This article develops a vector autoregression (VAR) for time series which are observed at mixed frequencies—quarterly and monthly. The model is cast in state-space form and estimated with Bayesian methods under a Minnesota-style prior. We show how to evaluate the marginal data density to implement a data-driven hyperparameter selection. Using a real-time dataset, we evaluate forecasts from the mixed-frequency VAR and compare them to standard quarterly frequency VAR and to forecasts from MIDAS regressions. We document the extent to which information that becomes available within the quarter improves the forecasts in real time. This article has online supplementary materials.

KEY WORDS: Bayesian methods; Hyperparameter selection; Macroeconomic forecasting; Marginal data density; Minnesota prior; Real-time data; Vector autoregressions.

1. INTRODUCTION

Vector autoregressions (VARs) are widely used in empirical macroeconomics. A VAR is a multivariate time series model that can be used, for instance, to forecast individual time series, to predict comovements of macroeconomic or financial variables, to analyze the sources of business cycle fluctuations, or to assess the effects of monetary or fiscal policy interventions on the macroeconomy. VARs are typically estimated either exclusively based on quarterly observations or exclusively based on monthly observations. The advantage of using quarterly observations is that the set of macroeconomic series that could potentially be included in the VAR is larger. For instance, gross domestic product (GDP), as well as many other series that are published as part of the national income and product accounts (NIPA), are only available at quarterly frequency. The advantage of using monthly information, on the other hand, is that the VAR is able to track the economy more closely in real time.

To exploit the respective advantages of both monthly and quarterly VARs, this article develops a mixed-frequency VAR (MF-VAR) that allows some series to be observed at monthly and others at quarterly frequency. The MF-VAR can be conveniently represented as a state-space model, in which the state-transition equations are given by a VAR at monthly frequency and the measurement equations relate the observed series to the underlying, potentially unobserved, monthly variables that are stacked in the state vector. The MF-VAR is meant to be an attractive alternative to a standard VAR in which all series are time-aggregated to quarterly frequency (QF-VAR). To cope with the high dimensionality of the parameter space, the MF-VAR is equipped with a Minnesota prior and estimated using Bayesian methods. The Minnesota prior is indexed by a vector of hyperparameters that determine the relative weight of a priori and sample information.

This article makes contributions in two areas. On the methodological front we show how to numerically approximate the

marginal data density (MDD) of a linear Gaussian MF-VAR. The MDD can be used for a data-based selection of hyperparameters, which is essential to achieve a good forecasting performance with a densely parameterized VAR. The second set of contributions is empirical. We compile a real-time dataset for an 11-variable VAR that includes observations on real aggregate activity, prices, and financial variables, including GDP, unemployment, inflation, and the federal funds rate. Using this dataset, we recursively estimate the MF-VAR and assess its forecasting performance. The comparison to a QF-VAR is the main focus of the empirical analysis. This article does not provide a detailed comparison to other classes of forecasting models, for example, univariate time series models or dynamic factor models. We essentially operate under the assumption that the researcher has adopted a VAR framework, in part, for other reasons than pure forecasting, for example, to capture comovements, to study the effect of shocks on the macroeconomy, or to produce forecasts conditional on, say, an interest rate path.

First, we ask the following very basic question: What is the gain, if any, from using within-quarter monthly information in a VAR framework? To answer this question, we group our end-of-month forecast origins in three bins. Given the release schedule of macroeconomic data in the U.S., at the end of the first month of the quarter, no additional monthly observations of nonfinancial variables for the current quarter are available. At the end of the second and third month either one or two within-quarter monthly observations are available. We find that during the third month of the quarter the switch from a QF-VAR to a MF-VAR improves the one-step-ahead forecast (nowcast)

© 2015 American Statistical Association
Journal of Business & Economic Statistics
July 2015, Vol. 33, No. 3

DOI: 10.1080/07350015.2014.954707

Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/r/jbes.

accuracy on average by 60% for nonfinancial variables observed at the monthly frequency and by 11% for variables observed at quarterly frequency. In the first month of the quarter the improvements are about 6%. The improvement in forecast accuracy is most pronounced for short-horizon forecasts and tapers off in the medium and long run. Thus, if the goal is to generate VAR nowcasts or forecasts of one- or two-quarters ahead, it is well worth switching to from a QF-VAR to a MF-VAR. If the focus is on a one- to two-year horizon, the QF-VAR is likely to suffice.

Second, we generate real-time forecasts of macroeconomic aggregates for the 2008–2009 (Great) recession period. This episode is of great interest to macroeconomists, because the large drop in aggregate real activity poses a challenge for existing structural and nonstructural models. We document that the monthly information helped the MF-VAR track the economic downturn more closely in real time than the QF-VAR supporting the view that the MF-VAR is an attractive alternative to a standard QF-VAR. Third, as a byproduct of the MF-VAR estimation, we generate an estimate of monthly GDP growth rates, which may be of independent interest to business cycle researchers. Finally, we also provide a comparison of bivariate MF-VARs to mixed data sampling (MIDAS) regressions. We find for GDP forecasts that the percentage differential in forecast accuracy is the same, regardless whether the forecast is made at the end of the first, second, or third month of the quarter. We are interpreting this finding as both approaches being able to exploit the information differentials between the three months of the quarter in relative terms equally well. In absolute terms, the MF-VARs tend to outperform the MIDAS regressions in our particular implementation.

This article focuses on VARs which are time series models that generate multivariate predictive distributions. VARs have been an important forecasting tool in practice (see Litterman 1986 for an early assessment and Giannone, Lenza, and Primiceri 2014 for a recent assessment) and there is strong evidence that they perform well in high-dimensional environments if estimated with shrinkage estimation techniques (see, e.g., De Mol, Giannone, and Reichlin 2008; Banbura, Giannone, and Reichlin 2010). Moreover, in addition to generating unconditional forecasts, they are widely used to produce conditional forecasts, e.g., conditional on an interest rate path (see Doan, Litterman, and Sims 1984; Waggoner and Zha 1999), which do require a multivariate framework. In our comparison between MF-VARs and QF-VARs we mostly study univariate root-mean-squared errors (RMSEs), though we also consider log determinants of (multivariate) forecast error covariance matrices. To the extent that we are considering univariate RMSEs one could conduct comparisons with univariate predictive regressions. However, comparisons of VAR forecasts to forecasts from other classes of time series models are not the focus of this article and can be found elsewhere in the literature (see, e.g., Chauvet and Potter (2013) for forecasting output and Faust and Wright (2013) for forecasting inflation).

To cope with the high dimensionality of the parameter space, the MF-VAR is equipped with a Minnesota prior and estimated with Bayesian methods. Our version of the Minnesota prior is based on Sims and Zha (1998). This prior is also used, for instance, in Banbura, Giannone, and Reichlin (2010) and Giannone, Lenza, and Primiceri (2014) and the authors document

that the forecasting performance of the Bayesian VAR dominates that of an unrestricted VAR by a large margin. Alternative prior specifications for Bayesian VARs are surveyed in Karlsson (2013) and the effect of specification choice on forecast accuracy is studied in Carriero, Clark, and Marcellino (2013). To generate accurate forecasts it is important that the prior covariance matrix is properly configured. A set of hyperparameters controls the degree of shrinkage toward the prior mean and we choose the hyperparameter to maximize the log MDD. MDD-based hyperparameter selection has been discussed, for instance, in Phillips (1996), in Del Negro and Schorfheide (2004), Del Negro et al. (2007), and, most recently, studied in Giannone, Lenza, and Primiceri (2014). Del Negro and Schorfheide (2004) and Del Negro et al. (2007) documented in their empirical analysis (based on priors constructed from dynamic stochastic general equilibrium models) that the hyperparameter that maximizes the log MDD approximately minimizes the pseudo-out-of-sample forecast errors.

We are building on existing approaches of treating missing observations in state-space models (see, e.g., the books by Harvey 1989 and Durbin and Koopman 2001). We are employing modern Bayesian computational tools, in particular the method of data augmentation. We construct a Gibbs sampler along the lines of Carter and Kohn (1994) that alternates between the conditional distribution of the VAR parameters given the unobserved monthly series, and the conditional distribution of the missing monthly observations given the VAR parameters. Draws from the former distribution are generated by direct sampling from a Normal-Inverted Wishart distribution, whereas draws from the latter are obtained by applying a simulation smoother to the state-space representation of the MF-VAR. Our numerical approximation of the log MDD is based on the modified harmonic mean estimator proposed by Geweke (1999).

An alternative Gibbs sampling approach for the coefficients in an MF-VAR was explored by Chiu et al. (2012). Their algorithm also iterates over the conditional posterior distributions of the VAR parameters and the missing monthly observations, but uses a different procedure to draw the missing observations. The focus of their article is on parameter estimation rather than forecasting. The authors linked the coefficients of the MF-VAR to the coefficients of a QF-VAR via a transformation. Chiu et al. (2012) then compared the posterior distributions of parameters and impulse response functions obtained from the estimation of the two models to document the value of the monthly observations.

Mixed-frequency observations have also been used in the estimation of dynamic factor models (DFMs). Mariano and Murasawa (2003) applied maximum-likelihood factor analysis to a mixed-frequency series of quarterly real GDP and monthly business cycle indicators to construct an index that is related to monthly real GDP. Aruoba, Diebold, and Scotti (2009) developed a DFM to construct a broad index of economic activity in real time using a variety of data observed at different frequencies. Giannone, Reichlin, and Small (2008) used a mixed-frequency DFM to evaluate the marginal impact that intra-monthly data releases have on current-quarter forecasts (nowcasts) of real GDP growth.

When using our MF-VAR to forecast quarterly GDP growth, we are essentially predicting a quarterly variable based on a

mixture of quarterly and monthly regressors. Ghysels, Sinko, and Valkanov (2007) proposed a simple univariate regression model, called a mixed data sampling (MIDAS) regression, to exploit high-frequency information without having to estimate a state-space model. In a MIDAS regression a low-frequency dependent variable, for example, quarterly GDP growth, is regressed on its lags as well as lags of high-frequency variables, for example, monthly observations on unemployment and/or daily observations on interest rates. To cope with potentially large numbers of regressors, the coefficients for the high-frequency regressors are tightly restricted through distributed lag polynomials that are indexed by a small number of hyperparameters.

Bayesian versions of the univariate MIDAS approach were developed by Rodriguez and Puggioni (2010) and Carriero, Clark, and Marcellino (2012). The key difference compared to the frequentist approach in Ghysels, Sinko, and Valkanov (2007) is that prior distributions can be used to induce parsimony. Rodriguez and Puggioni (2010) used Bernoulli-type prior distributions that exclude regressors with a certain probability. The larger this probability, the sparser the estimated model. The authors modify the priors to reflect the belief that if a certain lag l of a high-frequency regressor is included, then adjacent lags are more likely to be included. Carriero, Clark, and Marcellino (2012) used a prior distribution that is very similar to the prior used in our article: a set of hyperparameters controls the prior variance and thereby the degree of shrinkage of the coefficients toward the prior mean.

Bai, Ghysels, and Wright (2013) examined the relationship between MIDAS regressions and state-space models applied to mixed-frequency data. They considered dynamic factor models and characterize conditions under which the MIDAS regression exactly replicates the steady state Kalman filter weights on lagged observables. They concluded that Kalman filter forecasts are typically a little better, but MIDAS regressions can be more accurate if the state-space model is misspecified or over-parameterized. Kuzin, Marcellino, and Schumacher (2011) compared the accuracy of Euro Area GDP growth forecasts from MIDAS regressions and MF-VARs estimated by maximum likelihood. The authors found that the relative performances of MIDAS and MF-VAR forecasts differ depending on the predictors and forecast horizons. Overall, the authors did not find a clear winner in terms of forecasting performance.

Ghysels (2012) generalized the MIDAS approach to a VAR setting. Unlike our MF-VAR, his MIDAS VAR is an observation-driven model that does not require numerical techniques to integrate out unobserved monthly variables. As in Chiu et al. (2012), the empirical analysis focuses on impulse responses but not on real-time forecasting. In our view, the state-space setup pursued in this article is more transparent and flexible because the underlying state-transition equation resembles a standard (high-frequency) VAR. The computational advances of the last decade make it feasible to estimate the parameters of Bayesian state-space models and to conduct inference on the unobserved monthly states with code written in high-level languages such as MATLAB in a short amount of time.

The remainder of this article is organized as follows. Section 2 presents the state-space representation of the MF-VAR and discusses Bayesian inference and forecasting. The real-time datasets used for the forecast comparison of MF-VAR and

QF-VAR, as well as the timing of within-quarter monthly information, are discussed in Section 3. The empirical results are presented in Section 4. Finally, Section 5 concludes. The online appendix provides detailed information about the Bayesian computations, the construction of the dataset, as well as additional empirical results.

2. A MIXED-FREQUENCY VECTOR AUTOREGRESSION

The MF-VAR considered in this article is based on a standard constant-parameter VAR in which the length of the time period is one month. Since some macroeconomic time series, for example, GDP, are measured only at quarterly frequency, we treat the corresponding monthly values as unobserved. To cope with the missing observations, the MF-VAR is represented as a state-space model in Section 2.1. To ease the exposition, we use a representation with a state vector that includes even those variables that are observable at monthly frequency, for example, the aggregate price level, the unemployment rate, and the interest rate. A computationally more efficient representation in which variables observed at monthly frequency are dropped from the state vector is presented in the online appendix. Bayesian inference and forecasting are discussed in Section 2.2.

Throughout this article, we use $Y_{t_0:t_1}$ to denote the sequence of observations or random variables $\{y_{t_0}, \dots, y_{t_1}\}$. If no ambiguity arises, we sometimes drop the time subscripts and abbreviate $Y_{1:T}$ by Y . If θ is the parameter vector, then we use $p(\theta)$ to denote the prior density, $p(Y|\theta)$ is the likelihood function, and $p(\theta|Y)$ the posterior density. We use iid to abbreviate independently and identically distributed, and $N(\mu, \Sigma)$ denotes a multivariate normal distribution with mean μ and covariance matrix Σ . Let \otimes be the Kronecker product. If $X|\Sigma \sim MN_{p \times q}(M, \Sigma \otimes P)$ is matricvariate Normal and $\Sigma \sim IW_q(S, \nu)$ has an inverted Wishart distribution, we say that (X, Σ) has a Normal-inverted Wishart distribution: $(X, \Sigma) \sim MNIW(M, P, S, \nu)$.

2.1 State-Transitions and Measurement

We assume that the economy evolves at monthly frequency according to the following VAR(p) dynamics:

$$x_t = \Phi_1 x_{t-1} + \dots + \Phi_p x_{t-p} + \Phi_c + u_t, \quad u_t \sim \text{iid}N(0, \Sigma). \quad (1)$$

The $n \times 1$ vector of macroeconomic variables x_t can be composed into $x_t = [x'_{m,t}, x'_{q,t}]'$, where the $n_m \times 1$ vector $x_{m,t}$ collects variables that are observed at monthly frequency, for example, the consumer price index and the unemployment rate, and the $n_q \times 1$ vector $x_{q,t}$ comprises the unobserved monthly variables that are only published quarterly, for example, GDP. Define $z_t = [x'_t, \dots, x'_{t-p+1}]'$ and $\Phi = [\Phi_1, \dots, \Phi_p, \Phi_c]'$. Write the VAR in (1) in companion form as

$$z_t = F_1(\Phi)z_{t-1} + F_c(\Phi) + v_t, \quad v_t \sim \text{iid}N(0, \Omega(\Sigma)), \quad (2)$$

where the first n rows of $F_1(\Phi)$, $F_c(\Phi)$, and v_t are defined to reproduce (1) and the remaining rows are defined to deliver the identities $x_{q,t-l} = x_{q,t-l}$ for $l = 1, \dots, p-1$. The $n \times n$ upper-left submatrix of Ω equals Σ and all other elements are

zero. Equation (2) is the state-transition equation of the MF-VAR.

We proceed by describing the measurement equation. One can handle the unobserved variables in several ways: by imputing zeros and modifying the measurement equation by setting the loadings on the state variables to zero (e.g., Mariano and Murasawa 2003); by setting the measurement error variance to infinity (e.g., Giannone, Reichlin, and Small 2008); or by varying the dimension of the vector of observables as a function of time t (e.g., Durbin and Koopman 2001). We employ the latter approach. To do so, some additional notation is useful. Let T denote the forecast origin and let $T_b \leq T$ be the last period that corresponds to the last month of the quarter for which all quarterly observations are available. The subscript b stands for *balanced* sample. Up until period T_b the vector of monthly series $x_{m,t}$ is observed every month. We denote the actual observations by $y_{m,t}$ and write

$$y_{m,t} = x_{m,t}, \quad t = 1, \dots, T_b. \quad (3)$$

Assuming that the underlying monthly VAR has at least three lags, that is, $p \geq 3$, we express the three-month average of $x_{q,t}$ as

$$\tilde{y}_{q,t} = \frac{1}{3}(x_{q,t} + x_{q,t-1} + x_{q,t-2}) = \Lambda_{qz}z_t. \quad (4)$$

For variables measured in logs, for example, \ln GDP, the formula can be interpreted as a log-linear approximation to an arithmetic average of GDP that preserves the linear structure of the state-space model. For flow variables such as GDP, we adopt the NIPA convention and annualize high-frequency flows. As a consequence, quarterly flows are the average and not the sum of monthly flows. This three-month average, however, is only observed for every third month, which is why we use a tilde superscript. Let $M_{q,t}$ be a selection matrix that equals the identity matrix if t corresponds to the last month of a quarter and is empty otherwise. Adopting the convention that the dimension of the vector $y_{q,t}$ is n_q in periods in which quarterly averages are observed and zero otherwise, we write

$$y_{q,t} = M_{q,t}\tilde{y}_{q,t} = M_{q,t}\Lambda_{qz}z_t, \quad t = 1, \dots, T_b. \quad (5)$$

For periods $t = T_b + 1, \dots, T$ no additional observations of the quarterly time series are available. Thus, for these periods the dimension of $y_{q,t}$ is zero and the selection matrix $M_{q,t}$ in (5) is empty. However, the forecaster might observe additional monthly variables. Let $y_{m,t}$ denote the subset of monthly variables for which period t observations are reported by the statistical agency after period T , and let $M_{m,t}$ be a deterministic sequence of selection matrices such that (3) can be extended to

$$y_{m,t} = M_{m,t}x_{m,t}, \quad t = T_b + 1, \dots, T. \quad (6)$$

Notice that the dimension of the vector $y_{m,t}$ is potentially time varying and less than n_m . The measurement Equations (3) to (6) can be written more compactly as

$$y_t = M_t\Lambda_z z_t, \quad t = 1, \dots, T. \quad (7)$$

Here, M_t is a sequence of selection matrices that selects the time t variables that have been observed by period T and are part of the forecaster's information set. In sum, the state-space representation of the MF-VAR is given by (2) and (7).

2.2 Bayesian Inference

The starting point of Bayesian inference for the MF-VAR is a joint distribution of observables $Y_{1:T}$, latent states $Z_{0:T}$, and parameters (Φ, Σ) , conditional on a presample $Y_{-p+1:0}$ to initialize lags. Using a Gibbs sampler, we generate draws from the posterior distributions of $(\Phi, \Sigma)|(Z_{0:T}, Y_{-p+1:T})$ and $Z_{0:T}|(\Phi, \Sigma, Y_{-p+1:T})$. Based on these draws, we are able to simulate future trajectories of y_t to characterize the predictive distribution associated with the MF-VAR and to calculate point and density forecasts.

Prior Distribution. An important challenge in practical work with VARs is to cope with the dimensionality of the coefficient matrix Φ . Informative prior distributions can often mitigate the curse of dimensionality. A widely used prior in the VAR literature is the so-called Minnesota prior. This prior dates back to Litterman (1980) and Doan, Litterman, and Sims (1984). We use the version of the Minnesota prior described in Del Negro and Schorfheide (2011)'s handbook chapter, which in turn is based on Sims and Zha (1998). The main idea of the Minnesota prior is to center the distribution of Φ at a value that implies a random-walk behavior for each of the components of x_t in (1). Our version of the Minnesota prior for (Φ, Σ) is proper and belongs to the family of MNIW distributions. We implement the Minnesota prior by mixing artificial (or *dummy*) observations into the estimation sample. The artificial observations are computationally convenient and allow us to generate plausible a priori correlations between VAR parameters. The variance of the prior distribution is controlled by a low-dimensional vector of hyperparameters λ . Details of the prior are relegated to the online appendix, and the choice of hyperparameters is discussed below.

Posterior Inference. The joint distribution of data, latent variables, and parameters conditional on some observations to initialize lags can be factorized as follows:

$$p(Y_{1:T}, Z_{0:T}, \Phi, \Sigma | Y_{-p+1:0}, \lambda) = p(Y_{1:T} | Z_{0:T})p(Z_{1:T} | z_0, \Phi, \Sigma)p(z_0 | Y_{-p+1:0})p(\Phi, \Sigma | \lambda). \quad (8)$$

The distribution of $Y_{1:T} | Z_{1:T}$ is given by a point mass at the value of $Y_{1:T}$ that satisfies (7). The density $p(Z_{1:T} | z_0, \Phi, \Sigma)$ is obtained from the linear Gaussian regression (2). The conditional density $p(z_0 | Y_{-p+1:0})$ is chosen to be Gaussian and specified in the online appendix. Finally, $p(\Phi, \Sigma | \lambda)$ represents the prior density of the VAR parameters. The factorization (8) implies that the conditional posterior densities of the VAR parameters and the latent states of the MF-VAR take the form

$$p(\Phi, \Sigma | Z_{0:T}, Y_{-p+1:T}) \propto p(Z_{1:T} | z_0, \Phi, \Sigma)p(\Phi, \Sigma | \lambda) \\ p(Z_{0:T} | \Phi, \Sigma, Y_{-p+1:T}) \propto p(Y_{1:T} | Z_{1:T}) \\ p(Z_{1:T} | z_0, \Phi, \Sigma)p(z_0 | Y_{-p+1:0}). \quad (9)$$

We follow Carter and Kohn (1994) and use a Gibbs sampler that iterates over the two conditional posterior distributions in (9). Conditional on $Z_{0:T}$, the companion-form state transition (2) is a multivariate linear Gaussian regression. Since our prior for (Φ, Σ) belongs to the MNIW family, so does the posterior and draws from this posterior can be obtained by direct Monte Carlo sampling. Likewise, since the MF-VAR is set up

as a linear Gaussian state-space model, a standard simulation smoother can be used to draw the sequence $Z_{0:T}$ conditional on the VAR parameters. The distribution $p(z_0|Y_{-p+1})$ provides the initialization for the Kalman-filtering step of the simulation smoother. A detailed discussion of these computations can be found in textbook treatments of the Bayesian analysis of state-space models, for example, the handbook chapters by Del Negro and Schorfheide (2011) and Giordani, Pitt, and Kohn (2011).

Computational Considerations. For expositional purposes, it has been convenient to define the vector of state variables as $z_t = [x'_t, \dots, x'_{t-p+1}]'$, which includes the variables observed at monthly frequency. From a computational perspective, this definition is inefficient because it enlarges the state space of the model unnecessarily. We show in the online appendix how to rewrite the state-space representation of the MF-VAR in terms of a lower-dimensional state vector $s_t = [x'_{q,t}, \dots, x'_{q,t-p}]'$ that only includes the variables (and their lags) observed at quarterly frequency. Our simulation smoother uses the small state vector s_t for $t = 1, \dots, T_b$ and then switches to the larger state vector z_t for $t = T_b + 1, \dots, T$ to accommodate missing monthly observations toward the end of the sample.

Forecasting. For each draw $(\Phi, \Sigma, Z_{0:T})$ from the posterior distribution we simulate a trajectory $Z_{T+1:T+H}$ based on the state-transition Equation (2). Since we evaluate forecasts of quarterly averages in our empirical analysis, we time-aggregate the simulated trajectories accordingly. Based on the simulated trajectories (approximate) point forecasts can be obtained by computing means or medians. Interval forecasts and probability integral transformations (see Section C.3 of the online appendix) can be computed from the empirical distribution of the simulated trajectories.

2.3 Marginal Likelihood Function and Hyperparameter Selection

The empirical performance of the MF-VAR is sensitive to the choice of hyperparameters. The prior is parameterized such that $\lambda = 0$ corresponds to a flat (and therefore improper) prior for (Φ, Σ) . As $\lambda \rightarrow \infty$, the MF-VAR is estimated subject to the random-walk restriction implied by the Minnesota prior. The best forecasting performance of the MF-VAR is likely to be achieved for values of λ that are in between the two extremes. In a Bayesian framework the hyperparameter, λ can be interpreted as a model index (since a Bayesian model is the product of likelihood function and prior distribution). We consider a grid $\lambda \in \Lambda$ and assign equal prior probability to each value on the grid. Thus, the posterior probability of λ is proportional to the MDD

$$p(Y_{1:T}|Y_{-p+1:0}, \lambda) = \int p(Y_{1:T}, Z_{0:T}, \Phi, \Sigma|Y_{-p+1:0}, \lambda) \times d(Z_{0:T}, \Phi, \Sigma). \quad (10)$$

The log MDD can be interpreted as the sum of one-step-ahead predictive scores:

$$\ln p(Y_{1:T}|Y_{-p+1:0}, \lambda) = \sum_{t=1}^T \ln \int p(y_t|Y_{-p+1:t-1}, \Phi, \Sigma) \times p(\Phi, \Sigma|Y_{-p+1:t-1}, \lambda) d(\Phi, \Sigma). \quad (11)$$

The terms on the right-hand side of (11) provide a decomposition of the one-step-ahead predictive densities $p(y_t|Y_{1-p:t-1}, \lambda)$. This decomposition highlights the fact that inference about the parameter is based on time $t - 1$ information, when making a one-step-ahead prediction for y_t .

Hyperparameter Selection. To generate the MF-VAR forecasts, for each forecast origin we condition on the value $\hat{\lambda}_T$ that maximizes the log MDD. This procedure can be viewed as an approximation to a model averaging procedure that integrates out λ based on the posterior $p(\lambda|Y_{-p+1:T})$. The MDD-based selection of VAR hyperparameters has a fairly long history and tends to work well for forecasting purposes (see Giannone, Lenza, and Primiceri 2014 for a recent study).

Marginal Data Density Approximation. From (10) we see that the computation of the MDD involves integrating out the latent states. In the remainder of this section we describe how we compute the integral. To simplify the exposition we consider the special case of $n = 2$, $p = 1$, and $T = 3$. We assume that one of the variables is observed at monthly frequency and the other as a quarterly average. Thus, we can write $z_t = [x_{m,t}, x_{q,t}]'$. The observations $Y_{1:3}$ are related to the states $Z_{1:3}$ as follows:

$$y_1 = x_{m,1}, \quad y_2 = x_{m,2}, \quad y_3 = \begin{bmatrix} x_{m,3} \\ \frac{1}{3}(x_{q,1} + x_{q,2} + x_{q,3}) \end{bmatrix}. \quad (12)$$

Using a change of variable of the form

$$Z_{1:3} = J \begin{bmatrix} Y_{1:3} \\ W_{1:3} \end{bmatrix}, \quad (13)$$

where $Z_{1:3} = [z'_1, z'_2, z'_3]'$, $Y_{1:3} = [y_1, y_2, y'_3]'$, $W_{1:3} = [x_{q,1}, x_{q,2}]'$ (note that despite the 1:3 subscript, $W_{1:3}$ is a 2×1 vector in this example), and J is a 6×6 nonsingular matrix of constants. Thus, we can replace $p(Z_{1:3}|\lambda)$ by $p(Y_{1:3}, W_{1:3}|\lambda) = p(Y_{1:3}|W_{1:3})p(W_{1:3}|\lambda)$. Using Bayes' theorem, we can write (abstracting from the initialization of the VAR)

$$\frac{1}{p(Y_{1:3}|\lambda)} = \frac{p(W_{1:3}|Y_{1:3}, \lambda)}{p(Y_{1:3}, W_{1:3}|\lambda)}. \quad (14)$$

Suppose that $f(W_{1:3})$ has the property that $\int f(W_{1:3})dW_{1:3} = 1$, and let $\{W_{1:3}^{(i)}\}_{i=1}^N$ denote a sequence of draws from the posterior distribution of $W_{1:3}|(Y_{1:3}, \lambda)$. Then the MDD can be approximated using Geweke's (1999) harmonic mean estimator, which is widely used in the DSGE model literature to approximate MDDs in high-dimensional settings:

$$\hat{p}(Y_{1:3}|\lambda) = \left[\frac{1}{N} \sum_{i=1}^N \frac{f(W_{1:3}^{(i)})}{p(Y_{1:3}, W_{1:3}^{(i)}|\lambda)} \right]^{-1}. \quad (15)$$

The draws from the distribution of $W_{1:3}|(Y_{1:3}, \lambda)$ can be obtained by transforming the draws from $Z_{1:3}|(Y_{1:3}, \lambda)$, which are generated as a by-product of the posterior sampler described in Section 2.2. Using the properties of the MNIW distribution, it is straightforward to compute

$$p(Z_{1:3}|\lambda) = \int p(Z_{1:3}|\Phi, \Sigma)p(\Phi, \Sigma|\lambda)d(\Phi, \Sigma) \quad (16)$$

analytically. A straightforward change of variables based on (13) leads from $p(Z_{1:3}|\lambda)$ to $p(Y_{1:3}, W_{1:3}^{(i)}|\lambda)$. Note that the Jacobian of this transformation is simply a constant term.

Generalization. Taking the initialization of the VAR into account, the identity provided in (14) can be generalized as follows:

$$\frac{1}{p(Y_{1:T}|Y_{-p+1:0}, \lambda)} = \frac{p(W_{1:T}, w_0|Y_{1:T}, Y_{-p+1:0}, \lambda)}{p(W_{1:T}, Y_{1:T}, w_0|Y_{-p+1:0}, \lambda)}, \quad (17)$$

with the understanding that $W_{1:T}$ stacks the unobserved values of $x_{q,t}$ for the first and second month of each quarter of the estimation sample and w_0 contains the corresponding values for the initialization period $t = -p + 1, \dots, 0$. The approximation of the MDD becomes:

$$\hat{p}(Y_{1:T}|Y_{-p+1:0}, \lambda) = c \left[\frac{1}{N} \sum_{i=1}^N \frac{f_0(w_0^{(i)})f(W_{1:T}^{(i)})}{p(Z_{1:T}^{(i)}|z_0^{(i)}, \lambda)p(z_0^{(i)}|Y_{-p+1:0}, \lambda)} \right]^{-1}. \quad (18)$$

The constant c in (18) captures the Jacobian term associated with the change-of-variables from $(w_0, W_{1:T}, Y_{1:T})$ to $(z_0, Z_{1:T})$. For the function $f(\cdot)$ we follow Geweke (1999) and use a trimmed multivariate normal distribution with mean $\hat{\mu}_{W_{1:T}} = \frac{1}{N} \sum_{i=1}^N W_{1:T}^{(i)}$ and variance $\hat{\Sigma}_{W_{1:T}} = \frac{1}{N} \sum_{i=1}^N W_{1:T}^{(i)} W_{1:T}^{(i)'} - \hat{\mu}_{W_{1:T}} \hat{\mu}_{W_{1:T}}'$. This normal distribution approximates $p(W_{1:T}|Y_{-p+1:T})$ and stabilizes the ratio in (18). We set $f_0(w_0^{(i)}) = p(z_0^{(i)}|Y_{-p+1:0}, \lambda)$ such that the two terms cancel. To evaluate the denominator, we use the analytical expression for $p(Z_{1:T}^{(i)}|z_0^{(i)}, Y_{-p+1:0}, \lambda)$, which is obtained from the the normalization constants for the MNIW distribution and is provided, for instance, in Del Negro and Schorfheide (2011, sec 2).

3. REAL-TIME DATASETS AND INFORMATION STRUCTURE

We subsequently conduct a pseudo-out-of-sample forecast experiment with real-time data to study the extent to which the incorporation of monthly observations via an MF-VAR model improves upon forecasts generated with a VAR that is based on time-aggregated quarterly data (QF-VAR). We consider VARs for 11 macroeconomic variables, which are summarized in Section 3.1. The construction of the real-time datasets and the classification of forecast origins based on within-quarter monthly information are described in Section 3.2. Section 3.3 explains our choice of actual values that are used to compute forecast errors.

3.1 Macroeconomic Variables

We consider VARs for eleven macroeconomic variables, of which three are observed at quarterly frequency and eight are observed at monthly frequency. The quarterly series are GDP, fixed investment (INVFIX), and government expenditures (GOV). The monthly series are the unemployment rate (UNR), hours worked (HRS), Consumer Price Index (CPI), Industrial Production Index (IP), Personal Consumption Expenditure (PCE), Federal Funds Rate (FF), Treasury Bond Yield (TB), and S&P 500 Index (SP500). Precise data definitions are provided in the online appendix. Series that are observed at a higher than monthly frequency are time-aggregated to monthly

frequency. The variables enter the VARs in log levels with the exception of UNR, FF, and TB, which are divided by 100 to make them commensurable in scale to the other log-transformed variables.

3.2 Real-Time Data for End-of-Month Forecasts

We consider an increasing sequence of estimation samples $Y_{-p+1:T}, T = T_{\min}, \dots, T_{\max}$, and generate forecasts for periods $T + 1, \dots, T + H$. The maximum forecast horizon H is chosen to be 24 months. The period $t = 1$ corresponds to 1968:M1, T_{\min} is 1997:M7, and T_{\max} is 2010:M1, which yields 151 estimation samples. We eliminated 4 of the 151 samples because the real-time data for PCE and INVFIX were incomplete. The estimation samples are constructed from real-time datasets, assuming that the forecasts are generated on the last day of each month. Due to data revisions by statistical agencies, observations of $Y_{1:T-1}$ published in period T are potentially different from the observations that had been published in period $T - 1$. For this reason, real-time data are often indexed by a superscript, say $\tau \geq T$, which indicates the vintage or data release date. Using this notation, a forecaster at time T potentially has access to a triangular array of data $Y_{-p+1:1}^1, Y_{-p+1:2}^2, \dots, Y_{-p+1:T}^T$. Rather than using the entire triangular array and trying to exploit the information content in data revisions, we estimate the MF-VAR and QF-VAR for each forecast origin T based on the information set $Y_{-p+1:T}^T = \{y_{-p+1}^T, \dots, y_T^T\}$. As in Section 2, we are using the convention that the vector y_t^T contains only the subset of the 11 variables listed above for which observations are available at the end of month T .

To assess the usefulness of within-quarter information from monthly variables, we sort the forecast origins $T_{\min}, \dots, T_{\max}$ into three groups that reflect different within-quarter information sets. Forecast error statistics will be computed for each group separately. The grouping of forecast origins is best explained in a concrete example. Consider the January 31, 1998, forecast origin. By the end of January, the Bureau of Economic Analysis (BEA) has just published an *advance* estimate of 1997:Q4 GDP. In addition, the forecaster has access to nonfinancial monthly indicators from December 1997 and earlier. A similar situation arises at the end of April, July, and October. We refer to this group of forecast origins as “+0 months,” because the current-quarter forecasts do not use any additional nonfinancial monthly variables.

At the end of February 1998, the forecaster has access to an *preliminary* estimate of 1997:Q4 GDP and to observations for unemployment, industrial production, and so forth, for January 1998. Thus, we group February, May, August, and November forecasts and refer to them as “+1 month.” Following the same logic, the last subgroup of forecast origins has two additional monthly indicators (“+2 months”) and the *final* release of GDP for 1997:Q4 in the information set. Unlike the nonfinancial variables, which are released with a lag, financial variables are essentially available instantaneously. In particular, at the end of each month, the forecaster has access to average interest rates (FF and TB) and stock prices (SP500). The typical information sets for the three subgroups of forecast origins are summarized in Table 1.

Table 1. Illustration of information sets

		UNR	HRS	CPI	IP	PCE	FF	TB	SP500	GDP	INVFIX	GOV
January (+0 Months)												
Q4	M12	X	X	X	X	X	X	X	X	QAv	QAv	QAv
Q1	M1	∅	∅	∅	∅	∅	X	X	X	∅	∅	∅
February (+1 Month)												
Q4	M12	X	X	X	X	X	X	X	X	QAv	QAv	QAv
Q1	M1	X	X	X	X	X	X	X	X	∅	∅	∅
Q1	M2	∅	∅	∅	∅	∅	X	X	X	∅	∅	∅
March (+2 Months)												
Q4	M12	X	X	X	X	X	X	X	X	QAv	QAv	QAv
Q1	M1	X	X	X	X	X	X	X	X	∅	∅	∅
Q1	M2	X	X	X	X	X	X	X	X	∅	∅	∅
Q1	M3	∅	∅	∅	∅	∅	X	X	X	∅	∅	∅

NOTE: ∅ indicates that the observation is missing. X denotes monthly observation and QAv denotes quarterly average. “+0 Months” group: January, April, July, October; “+1 Month” group: February, May, August, November; “+2 Months” group: March, June, September, December.

Unfortunately, due to variation in release dates, not all 151 estimation samples mimic the information structure in Table 1. For 47 samples the last PCE figure is released with a two-period (approximately five weeks) instead of one-period (approximately four weeks) lag. This exception occurs for 28 samples of the “+0 months” group. For these samples a late release of PCE implies the quarterly consumption for the last completed quarter is not available. In turn, the QF-VAR could only be estimated based on information up to $T - 4$ instead of $T - 1$ and would be at a severe disadvantage compared to the MF-VAR. Since PCE is released only a few days after the period T forecasts are made, we predate its release. Thus, for the 28 samples of the “+0 months” group that are subject to the irregular timing, we use PCE_{T-1} in the estimation of both the QF-VAR and MF-VAR. No adjustments are made for the “+1 month” and “+2 months” groups. Further details about these exceptions are provided in the online appendix.

3.3 Actuals for Forecast Evaluation

The real-time-forecasting literature is divided as to whether forecast errors should be computed based on the first release following the forecast date, say y_{T+h}^{T+h} , or based on the most recent vintage, say y_{t+h}^* . The former might do a better job of capturing the forecaster’s loss, whereas the latter is presumably closer to the underlying “true” value of the time series. We decided to follow the second approach and evaluate the forecasts based on actual values from the $T_* = 2012:M1$ data vintage. While the MF-VAR in principle generates predictions at the monthly frequency, we focus on the forecasts of quarterly averages, which can be easily compared to forecasts from the QF-VAR.

4. EMPIRICAL RESULTS

The empirical analysis proceeds in four parts. The hyperparameter selection is discussed in Section 4.1. Section 4.2 compares root mean squared error (RMSE) statistics from the MF-VAR to a QF-VAR and a set of MIDAS regressions. Section 4.3 contrasts MF-VAR density forecasts during the 2008–9 (Great) recession with QF-VAR forecasts. Finally, in Section 4.4 we

present a monthly GDP series that arises as a by-product of the MF-VAR estimation. Based on some preliminary exploration of the MDDs, we set the number of lags in the (monthly) state transition of the MF-VAR to $p_{(m)} = 6$ and the number of lags in the QF-VAR to $p_{(q)} = 2$.

Unless otherwise noted, for each estimation sample we generate 20,000 draws from the posterior distribution of the VAR parameters using the MCMC algorithm described in Section 2.2. We discard the first 10,000 draws and use the remaining 10,000 to calculate Monte Carlo approximations of posterior moments. The online appendix provides some information on the accuracy of the MCMC. The Monte Carlo standard deviation of the posterior mean forecasts (output, inflation, interest rates, and unemployment), computed across independent runs of the MCMC, is generally less than 0.5 basis points. For comparison, the RMSE associated with these forecasts ranges from 10 to 200 basis points.

4.1 Hyperparameter Selection

We will subsequently compare MF-VAR and QF-VAR forecasts. Both VARs are equipped with a Minnesota prior that is represented in terms of dummy observations and indexed by a vector of hyperparameters λ . We use the same set of dummy observations for both types of VAR. However, the hyperparameters are chosen for each type of VAR separately. The careful choice of this hyperparameter vector is crucial for obtaining accurate forecasts. As explained in Section 2.3, we determine the hyperparameters by maximizing the log MDD. For the QF-VAR the MDD can be computed analytically (see, e.g., Del Negro and Schorfheide 2011) and the maximization is straightforward. Thus, we will focus on the hyperparameter selection for the MF-VAR.

The hyperparameter vector consists of five element, controlling: the overall tightness of the prior (λ_1); the rate at which the prior variance on higher-order lag coefficients decays (λ_2); the dispersion of the prior on the innovation covariance matrix (λ_3); the extent to which the sum-of-coefficient on the lags of a variable $x_{i,t}$ is tilted toward unity (λ_4); and the extent to which cointegration restrictions are imposed on the VAR coefficients

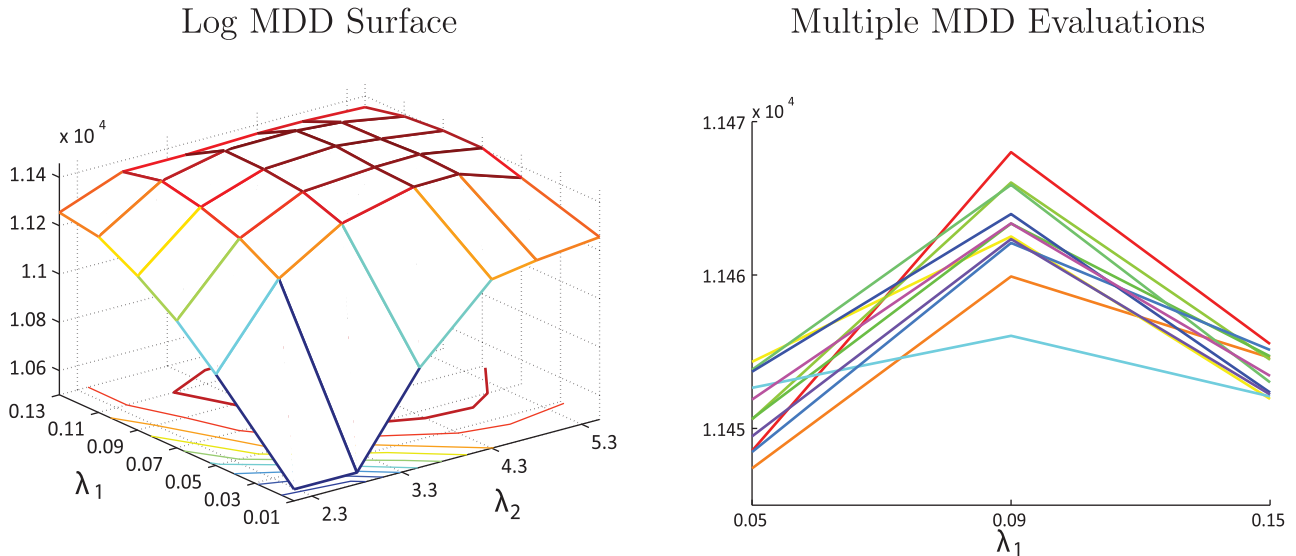


Figure 1. Log marginal data density for 11-Variable MF-VAR. *Notes:* The two plots depict $\ln \hat{p}(Y_{1:T} | Y_{-p+1:0,\lambda})$. In the left panel, we condition on $\lambda_3 = 1, \lambda_4 = 2.7,$ and $\lambda_5 = 4.3$. In the right panel we condition on $\lambda_2 = 4.3, \lambda_3 = 1, \lambda_4 = 2.7,$ and $\lambda_5 = 4.3$. Each “hair” corresponds to a separate run of the MCMC algorithm.

(λ_5). In general, the larger λ_i the smaller the prior variance and the more informative the prior. From a preliminary analysis based on the QF-VAR, we conclude that λ_3 is not particularly important for the forecasting performance and fix it as $\hat{\lambda}_3 = 1$. Based on a preliminary search over a grid $\Lambda^{(1)}$ we determine suitable values for λ_4 and λ_5 for the first recursive sample, which ranges from 1968:M1 to 1997:M7. These values are $\hat{\lambda}_4 = 2.7$ and $\hat{\lambda}_5 = 4.3$. Conditioning on $\hat{\lambda}_3$ to $\hat{\lambda}_5$, we use a second grid, $\hat{\Lambda}^{(2)}$ to refine the choice of λ_1 and λ_2 .

The log MDD surface is depicted in the left panel of Figure 1 as function of λ_1 and λ_2 , holding the remaining three hyperparameters fixed at $\lambda_3 = 1, \lambda_4 = \hat{\lambda}_4,$ and $\lambda_5 = \hat{\lambda}_5$. The surface has a convex shape and is maximized at $\hat{\lambda}_1 = 0.09$ and $\hat{\lambda}_2 = 4.3$. At its peak the value of the log MDD is approximately 11,460. While the surface is fairly flat near the peak, for example, for $\lambda_1 \in [0.05, 0.15]$ and $\lambda_2 \in [4, 4.5]$, the MDD values drop substantially for values of λ outside of these intervals. To assess the accuracy of the MDD evaluation, which involves the numerical evaluation of a high-dimensional integral, we display a hairplot of a slice of the MDD surface in the right panel of Figure 1, fixing λ_2 at 4.3. Each hairline corresponds to a separate run of the MCMC algorithm. We focus on the interval $\lambda_1 \in [0.05, 0.15]$. While there is some noticeable Monte Carlo variation with respect to the absolute magnitude of the log MDD, this variation does not affect inference with respect to the optimal value of λ on the grid. For each simulation, the log MDD peaks at 0.09. The accuracy of the approximation can be improved by increasing the number of MCMC draws.

The reoptimization of the hyperparameters for the MF-VAR is computationally costly. Because we expect the optimal hyperparameter choices to evolve smoothly over time, we are reoptimizing with respect to λ approximately every three years, namely for the 40th, the 75th, the 110th, and the 151th recursive sample. During this reoptimization we keep

$\hat{\lambda}_3, \hat{\lambda}_4,$ and $\hat{\lambda}_5$ fixed. The reoptimization essentially left the choice of hyperparameters unchanged. We obtained a similar result for the QF-VAR and decided to keep the MF-VAR and the QF-VAR hyperparameters constant for all recursive sample.

The hyperparameter estimates for the MF-VAR and the QF-VAR are summarized in Table 2. While the overall tightness of the prior, controlled by λ_1 , is larger for the QF-VAR than the MF-VAR, the MF-VAR strongly shrinks the coefficients on higher-order lags to zero. The QF-VAR only uses two lags which are associated with 22 regression coefficients for each endogenous variable. The MF-VAR, on the other hand, uses six lags which are associated with 66 regression coefficients. Roughly 30% of these coefficients are associated with regressors that are only observed a quarterly frequency. The hyperparameters for the QF-VAR are broadly in line with the results in Giannone, Lenza, and Primiceri (2014).

4.2 MF-VAR Point Forecasts

MF-VAR versus QF-VAR. We begin by comparing RMSEs for MF-VAR and QF-VAR forecasts of quarterly averages to assess the usefulness of monthly information. The RMSEs are computed separately for the “+0 months,” “+1 month,” and “+2 months” forecast origins defined in the previous section. Results for GDP growth (GDP), unemployment (UNR), inflation (INF), and the federal funds rate (FF) are reported in Figure 2. The

Table 2. Hyperparameters

	λ_1	λ_2	λ_3	λ_4	λ_5
MF-VAR(11)	0.09	4.30	1.0	2.70	4.30
QF-VAR(11)	3.08	0.01	1.0	1.12	1.62

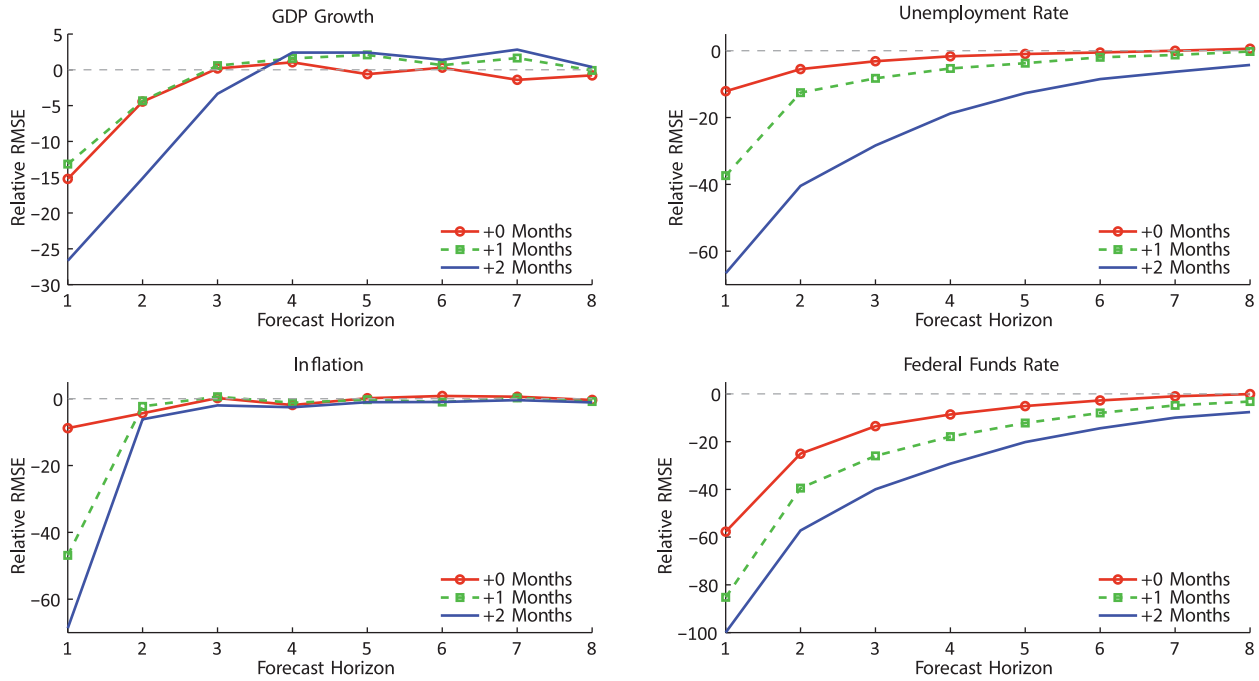


Figure 2. Relative RMSEs of 11-Variable MF-VAR versus QF-VAR.

figure depicts relative RMSEs defined as

$$\text{Relative RMSE}(i|h) = 100 \times \frac{\text{RMSE}(i|h) - \text{RMSE}_{\text{Benchmark}}(i|h)}{\text{RMSE}_{\text{Benchmark}}(i|h)}, \tag{19}$$

where i denotes the variable and we adopt the convention (in slight abuse of notation) that the forecast horizon h is measured in quarters. The QF-VAR serves as a benchmark model and $h = 1$ corresponds to the quarter in which the forecast is generated. The $h = 1$ forecast is often called a nowcast. Absolute RMSEs for the 11-variable MF-VAR are tabulated in the online appendix.

For all four series, the use of monthly information via the MF-VAR leads to a substantial RMSE reduction in the short run. Consider the GDP growth forecasts. The “+2” nowcasts have a 27% lower RMSE than the QF-VAR nowcasts. For the “+1 month” group and the “+0 months” group, the reductions are both 15%. While the “+2 months” group forecasts clearly dominate at the nowcast horizon $h = 1$, the relative ranking among the three sets of MF-VAR forecasts becomes ambiguous for $h \geq 2$. As the forecast horizon increases to $h = 4$, the QF-VAR catches up with the MF-VAR. For horizons $h \geq 4$, the RMSE differentials between QF-VAR and MF-VAR GDP growth forecasts are negligible.

For the monthly unemployment, inflation, and federal funds rate series, the short-run RMSE reductions attained by the MF-VAR for the monthly series are even stronger than for GDP growth, which is observed at quarterly frequency. This is, of course, not surprising. At the nowcast horizon, the MF-VAR is able to improve over the precision of the QF-VAR for the “+2 months” forecasts by 65% for unemployment, 70% for inflation, and 100% for the federal funds rate. Recall that “+2 months”

corresponds to the last month of the quarter, which means that at the end of the last month, the average quarterly interest rate is known. Thus, by construction the RMSE reduction for the federal funds rate is 100%. The RMSE reductions for the “+1 month” group range from 40% (unemployment) to 80% (federal funds rate). For the “+0 months” group the improvement of the nowcast from using the MF-VAR is about 10% for inflation and the unemployment rate and 60% for the federal funds rate. While the gains from using monthly information tend to persist for unemployment and interest rates as the forecast horizon h increases, for inflation, monthly observations generate no improvements of forecast performance beyond the nowcast horizon.

To summarize the multivariate forecast performance of the VARs and aggregate the univariate RMSE differentials across quarterly and monthly nonfinancial variables we consider the log-determinant of the forecast error covariance matrix, proposed by Doan, Litterman, and Sims (1984):

$$f(\hat{\epsilon}_t) = \ln \left(\left| \frac{1}{T_{\max} - T_{\min}} \sum_{t=T_{\min}}^{T_{\max}} \hat{\epsilon}_t \hat{\epsilon}_t' \right| \right), \tag{20}$$

where $\hat{\epsilon}_t$ is a vector of forecast errors. Log-determinant differentials of MF-VAR versus QF-VAR forecasts are depicted in Figure 3. We scale the log-determinant differentials by $100 \cdot 0.5/n_{\text{var}}$. The factor 0.5 converts mean-squared errors into RMSEs, the division by n_{var} yields an average across the variables included in $\hat{\epsilon}_t$, and the factor 100 converts the differential into percentages. This scaling makes the log-determinant differentials comparable to the RMSE differentials depicted in Figure 2. The results are qualitatively consistent with the comparison of univariate RMSEs. Not surprisingly, for the group of quarterly variables (GDP, INVMIX, GOV) the gain from including within-quarter monthly information is smaller than for

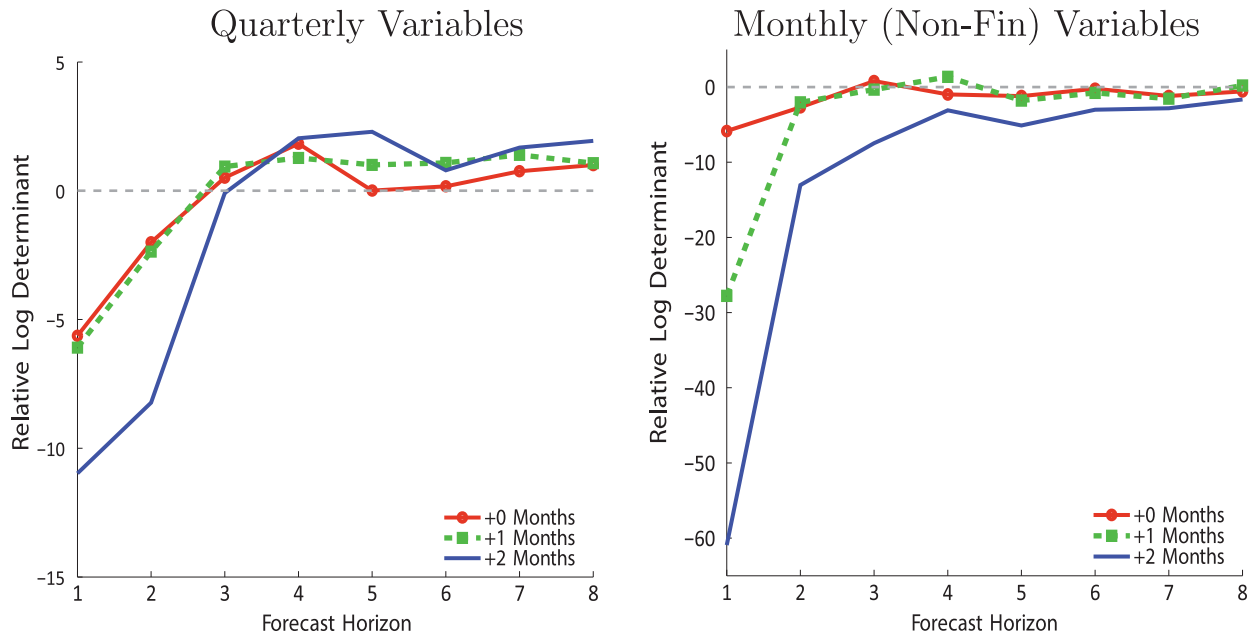


Figure 3. Log determinant of MF-VAR versus QF-VAR. *Notes:* The relative log determinant is defined as Relative Log Determinant = $(100 \cdot 0.5/n_{\text{var}})[f(\hat{\epsilon}_{t,\text{MF}}) - f(\hat{\epsilon}_{t,\text{QF}})]$, where $f(\cdot)$ is given in (20) and $n_{\text{var}} = 3$ for quarterly variables and $n_{\text{var}} = 5$ for monthly nonfinancial variables. The forecast horizon h is measured in quarters and $h = 1$ corresponds to the quarter in which the forecast is generated.

the group of monthly nonfinancial variables (UNR, HRS, CPI, IP, PCE). For quarterly variables the forecast accuracy gains relative to the QF-VAR range from 11% (“+2 months” group) to 6% (“+0 months” group). For monthly variables the gains for the three forecast origin groups are 60%, 30%, and 6%, respectively. For $h \geq 3$ the QF-VAR catches up with the MF-VAR and the benefit from using monthly information vanishes. The only exception are the “+2” months forecasts of the monthly variables. Here the within-quarter monthly information remains even for forecast horizons exceeding one year. We exclude the financial variables (FF, TB, SP500) from the group of monthly variables because the financial variables are essentially known at the end of each quarter (“+2 months” group) which creates a near-singularity in forecast error covariance matrices that include one or more financial variables.

MF-VAR versus MIDAS. A popular alternative to the multivariate state-space framework used in this article are MIDAS regressions. While there exist generalizations of the MIDAS approach to VAR settings, in most applications MIDAS regressions are used as univariate forecasting models. For a comparison of the two approaches we will focus on output growth. Our VAR models use 11 macroeconomic variables. If all of these variables are included in a MIDAS regression without any further restrictions, the MIDAS regression will perform very poorly. The distributed-lag restrictions on high-frequency regressors are designed to deal with many (high-frequency) observations of a single regressor but they are not designed to impose parsimony on a specification with many different right-hand-side variables. Thus, instead of comparing the 11-variable MF-VAR with MIDAS regressions, we will provide comparisons between bivariate MF-VARs and MIDAS regressions, estimated using the same set of variables.

Forni, Marcellino, and Schumacher (2013) proposed an unrestricted version of the MIDAS model (U-MIDAS) and show

that when the mismatch of the frequency is low, like in macroeconomic applications that typically involve monthly and quarterly data only, this unrestricted version performs better in Monte Carlo experiments and provides a better GDP nowcasting performance than a MIDAS regression with distributed-lag restrictions on the coefficients of the high-frequency variables. Thus, we consider U-MIDAS (instead of MIDAS) regressions in our comparison. The key aspect of our empirical analysis is the distinction between three groups of forecast origins, denoted by “+0,” “+1,” and “+2” (months). Each of these groups uses different within-quarter monthly information. Accordingly, we use three separate U-MIDAS regressions, which, using the notation of Section 2, can be written as

$$\begin{aligned}
 \text{“+0”} : \tilde{y}_{q,t+3h} &= \beta_0 + \beta_1 \tilde{y}_{q,t} + \beta_2 \tilde{y}_{q,t-3} + \sum_{s=1}^6 \gamma_s x_{m,t-s+1} \\
 &\quad + \text{resid}_{t+3h} \\
 \text{“+1”} : \tilde{y}_{q,t+3h} &= \beta_0 + \beta_1 \tilde{y}_{q,t} + \beta_2 \tilde{y}_{q,t-3} + \sum_{s=1}^6 \gamma_s x_{m,t-s+1} \\
 &\quad + \delta_1 x_{m,t+1} + \text{resid}_{t+3h} \\
 \text{“+2”} : \tilde{y}_{q,t+3h} &= \beta_0 + \beta_1 \tilde{y}_{q,t} + \beta_2 \tilde{y}_{q,t-3} + \sum_{s=1}^6 \gamma_s x_{m,t-s+1} \\
 &\quad + \delta_1 x_{m,t+1} + \delta_2 x_{m,t+2} + \text{resid}_{t+3h},
 \end{aligned} \tag{21}$$

where $t = 0, 3, 6, 9, \dots$. The quarterly variable $\tilde{y}_{q,t+3h}$ was defined as average of unobserved monthly variables in (4) and corresponds to log GDP. The monthly variable $x_{m,t}$ is assumed to be scalar and we consider all eight of our monthly variables individually. The regression (21) is estimated by OLS for each group of forecast origins and for each forecast horizon separately. Thus, as in Forni, Marcellino, and Schumacher (2013)

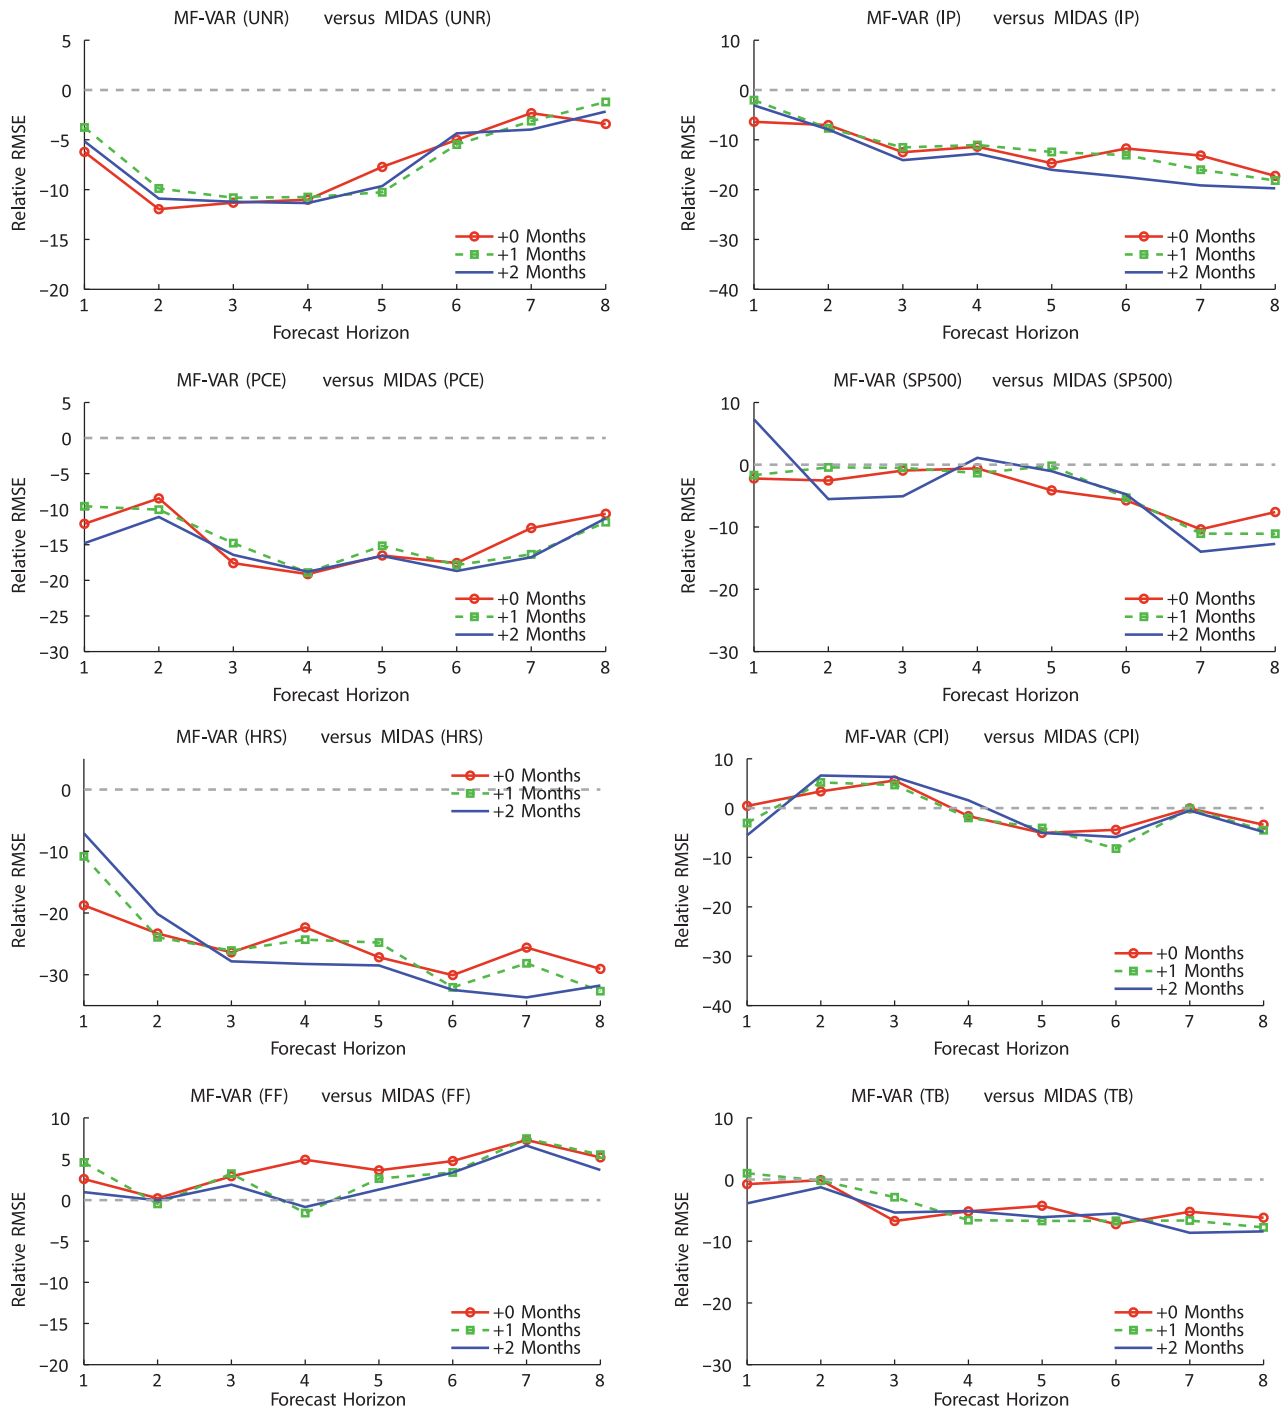


Figure 4. Relative RMSEs of Bivariate MF-VAR versus MIDAS.

we use direct estimation, that is, the projection of $\tilde{y}_{q,t+3h}$ on the predictors available at the forecast origin, to determine the coefficients for the multistep forecasting equation. Recall that under the Bayesian approach employed for the analysis of the MF-VAR multistep forecasts are generated by iterating the VAR forward and using the posterior distribution to integrate out the unknown parameters.

Figure 4 illustrates the log GDP forecast performance of the bivariate MF-VARs relative to the MIDAS regressions. Each panel corresponds to a different monthly variable. Two results

stand out. First, by and large, both the MF-VAR and MIDAS use the within-quarter monthly information equally well. The RMSE differentials are essentially the same for each of the three informational groups. For six out of the eight monthly variables the MF-VAR forecasts are more accurate than the MIDAS forecasts at some horizons, and no worse at the other horizons. For the unemployment rate, the gain from using the MF-VAR is highest for horizons of 2–5 quarters. For industrial production, the stock market index, hours, and the treasury bond rate the largest gain is realized at the long-horizon whereas for

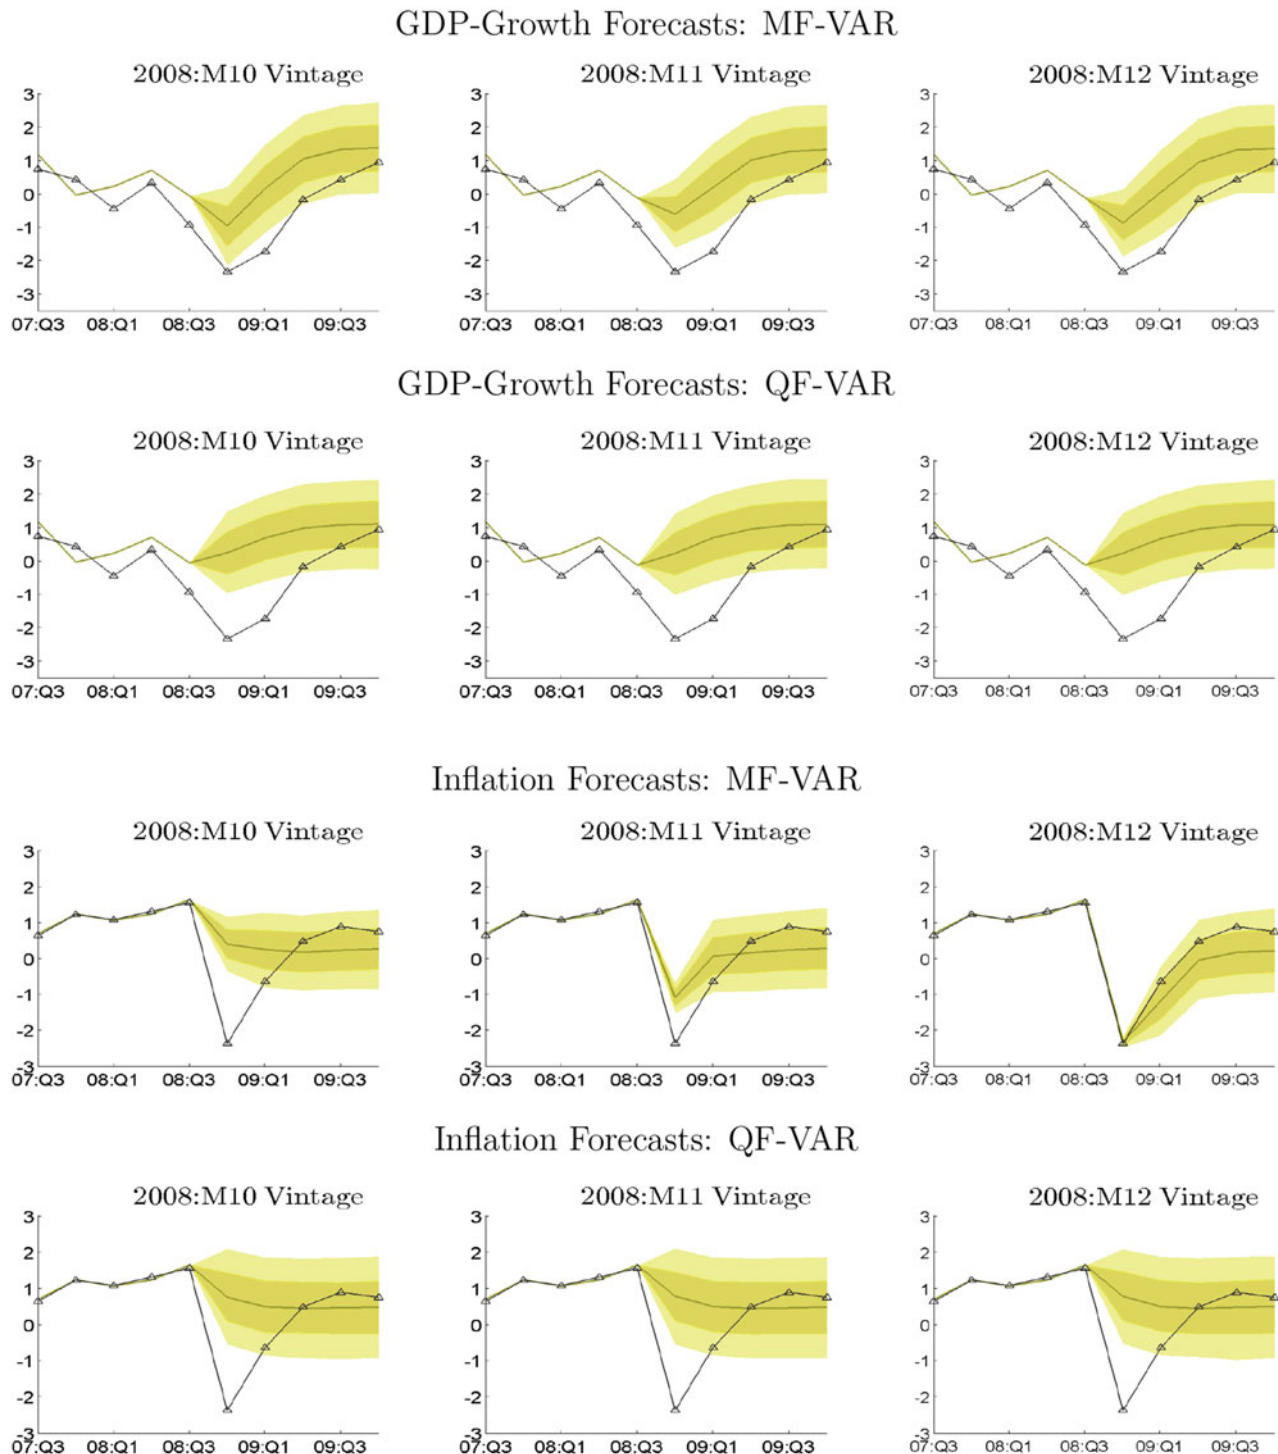


Figure 5. Real-time forecasts during the Great Recession. *Notes:* Actual values are from the $T_* = 2012 : M1$ data vintage and are denoted as the black line with triangles. The title in each subplot indicates the forecast origin and the data vintage that are used in the estimation. We show the median, 60% bands, and 90% bands constructed from the predictive distribution.

PCE the improvement is fairly uniform for one- to eight-quarter ahead forecasts. Only for the federal funds rate and CPI inflation MIDAS forecasts appear to be marginally more accurate than the MF-VAR forecasts.

Other Comparisons. In the online appendix we also provide RMSE comparisons between the 11-variable MF-VAR and univariate QF-AR(2) models; and between a 4-variable MF-VAR (GDP, CPI, UNR, FF) and a 4-variable QF-VAR. The results

are qualitatively very similar: there is a substantial gain from using the within-quarter-monthly information for nowcasting and short-horizon forecasting. This gain vanishes over one- to two-year horizons. Finally, the online appendix contains a careful comparison between MF-VAR forecasts and Greenbook (now Tealbook) forecasts, prepared by the staff of the Board of Governors for the meetings of the Federal Open Market Committee. At the nowcast horizon the unemployment forecasts

of the MF-VAR are at par with the Greenbook forecasts, whereas the GDP growth and inflation forecasts are less accurate than the Greenbook forecasts. Over a four- to five-quarter horizon the MF-VAR generates more accurate GDP forecasts, whereas the Greenbook contains more precise inflation and unemployment rate forecasts.

4.3 Forecasting During the Great Recession

The pseudo-out-of-sample forecast performance of the previous section documented that the use of within-quarter monthly observations increases the precision of short-run forecast. We now examine how the use of monthly real-time information sharpened the VAR forecasts during the recent recession. We focus on the period from October to December 2008. Figure 5 depicts real-time interval forecasts from the MF-VAR and the QF-VAR. Moreover, we plot actual values using the 2011:M7 data vintage. We focus on real GDP growth and CPI inflation. The figure is divided into subpanels that correspond to particular estimation samples and forecast horizons. The first column of panels depicts October 2008 forecasts (“+0 months” group), and the second and third columns show November (“+1 month”) and December (“+2 months”) forecasts, respectively. A comparison between the first and second (third and fourth) row of panels shows how monthly within-quarter information alters the density forecast for GDP (inflation).

The most striking feature of the top panels of Figure 5 is the -2% quarter-on-quarter growth rate of GDP in 2008:Q4. The magnitude of the drop in output growth in late 2008 is unexpected by the VAR models. It is, for all forecast origins, outside of the 90% predictive interval. The drop in GDP growth is equally unexpected by state-of-the-art dynamic stochastic general equilibrium (DSGE) models and the Blue Chip survey of professional forecasters as documented by Del Negro and Schorfheide (2013). A comparison of the MF-VAR and QF-VAR forecasts highlights how monthly information alters the within-quarter predictions. Notice that the QF-VAR forecasts do not stay constant within the quarter. The variation is caused by data revisions. As discussed in Section 3, each month new data releases for the previous quarter become available and change the lagged observations that determine the initial conditions for the VAR at the forecast origin. However, the within-quarter variation of the QF-VAR forecasts is fairly small.

By December 2008 the QF-VAR nowcasts and forecasts show still no evidence of a severe downturn, because the latest information that is used to generate the predictions stems from 2008:Q3. The MF-VAR forecasts, on the other hand, do get revised more substantially during each quarter. In addition to the presence of data revisions, the forecasts are updated based on the information that is available at monthly frequency. Compared to the QF-VAR forecasts, the MF-VAR nowcasts during the fourth quarter of 2008 are a lot more pessimistic, which is in line with the actual realization of output growth. Over a one-year horizon the discrepancy between the MF-VAR and QF-VAR forecasts vanishes, which is consistent with the forecast error statistics presented in the previous section. According to both VARs the GDP growth forecasts are mean reverting. The models predict a GDP growth rate of about 1% for the second half of 2009. This prediction turned out to be accurate.

The bottom panels of Figure 5 depict the evolution of inflation forecasts in the last quarter of 2008. Since the CPI is published at a monthly frequency, the differences between within-quarter inflation forecasts from the MF-VAR and QF-VAR are much more pronounced than for GDP. Throughout 2008:Q4 the inflation forecasts from the QF-VAR stay essentially constant and miss the -2% deflation rate in 2008:Q4. The MF-VAR, on the other hand, detects the deflation by November 2008 as it is unfolding. At the longer horizon, the MF-VAR correctly predicts that the deflation episode is short-lived and that inflation rate will, with about 50% probability, be positive by the end of 2009. To summarize, these real-time forecasts during the Great Recession illustrate that the MF-VAR can transform within-quarter monthly information into more accurate nowcasts and forecasts of quarterly averages.

4.4 Monthly GDP

The estimation of the MF-VAR generates a monthly GDP series as a byproduct. This series is implicitly extracted during the smoothing step of the Gibbs sampler (see Section 2.2) from the 11 macroeconomic time series that enter the MF-VAR. A time series plot of monthly GDP growth is depicted in Figure 6. For each trajectory of log GDP generated with the Gibbs sampler, we compute month-on-month growth rates (scaled by a factor of 3 to make them comparable to quarter-on-quarter rates). For each month we then plot the median growth rate across the




Figure 6. Monthly GDP growth (scaled to a quarterly rate).

simulated trajectories. We overlay monthly GDP growth rates published by Stock and Watson (2010), who combined monthly information about GDP components to distribute quarterly GDP across the three months of the quarter. Frale et al. (2011) used a similar approach to construct a monthly GDP series for the Euro Area. Moreover, we plot growth rates computed from NIPA's quarterly GDP, implicitly assuming that GDP growth is constant within a quarter. Two observations stand out. First, at a monthly frequency GDP growth is much more volatile than at a quarterly level. Second, the monthly GDP growth series obtained from the MF-VAR estimation is somewhat smoother than the Stock-Watson series. While the two monthly measures are positively correlated, they are not perfectly synchronized, which is consistent with these measures being constructed from very different source data.

5. CONCLUSION

We have specified a VAR for observations that are observed at different frequencies, namely, monthly and quarterly. A Gibbs sampler was used to conduct Bayesian inference for model parameters and unobserved monthly variables. To cope with the dimensionality of the MF-VAR, we used a Minnesota prior that shrinks the VAR coefficients toward univariate random-walk representations. The degree of shrinkage is determined in a data-driven way, by maximizing the log MDD with respect to a low-dimensional vector of hyperparameters and we show how to approximate the MDD of a MF-VAR. Finally, we used the model to generate forecasts. The main finding is that within-quarter monthly information leads to drastic improvements in the short-horizon forecasting performance. These improvements are increasing in the time that has passed since the beginning of the quarter. Over a one- to two-year horizon there are, however, no noticeable gains from using the monthly information.

SUPPLEMENTARY MATERIALS

File consisting of (i) details of the implementation of the Bayesian computations for the MF-VAR presented in the main text; (ii) the construction of the real-time dataset; and (iii) tables and figures with additional empirical results.

ACKNOWLEDGMENTS

We thank the Editor (Rong Chen), an Associate Editor, two anonymous referees, Frank Diebold, Jonathan Wright, Kei-Mu Yi, and seminar participants at the 2012 AEA Meetings, the Banco Central do Brasil, and the University of Pennsylvania for helpful comments and discussions. We greatly benefited from a MATLAB program written by Marco Del Negro and Dan Herbst to compile real-time datasets for the recursive estimation of forecasting models. Financial support from the Federal Reserve Bank of Minneapolis is gratefully acknowledged. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

[Received September 2012. Revised June 2014.]

REFERENCES

- Aruoba, S. B., Diebold, F. X., and Scotti, C. (2009), "Real-Time Measurement of Business Conditions," *Journal of Business and Economic Statistics*, 27, 417–427. [367]
- Bai, J., Ghysels, E., and J. Wright (2013), "State Space Models and MIDAS Regressions," *Econometric Reviews*, 32, 779–813. [368]
- Banbura, M., Giannone, D., and Reichlin, L. (2010), "Large Bayesian VARs," *Journal of Applied Econometrics*, 25, 71–92. [367]
- Carriero, A., Clark, T. E., and Marcellino, M. (2012), "Real-Time Nowcasting With a Bayesian Mixed Frequency Model With Stochastic Volatility," *FRB Cleveland Working Paper 12–27*. [368]
- (2013), "Bayesian VARs: Specification Choices and Forecast Accuracy," *Journal of Applied Econometrics*, doi: 10.1002/jae.2315. [367]
- Carter, C. K., and Kohn, R. (1994), "On Gibbs Sampling for State Space Models," *Biometrika*, 81, 541–553. [367,369]
- Chauvet, M., and Potter, S. (2013), "Forecasting Output," in *Handbook of Economic Forecasting* (vol. 2), eds. G. Elliott, and A. Timmermann, Amsterdam, Netherlands: Elsevier. [367]
- Chiu, C. W., Eraker, B., Foerster, A., Kim, T. B., and Seoane, H. (2012), "Estimating VAR's Sampled at Mixed or Irregular Spaced Frequencies: A Bayesian Approach," Research Working Paper No. RWP 11-11, Federal Reserve Bank of Kansas City. [367,368]
- De Mol, C., Giannone, D., and 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. [367]
- Del Negro, M., and Schorfheide, F. (2004), "Priors From General Equilibrium Models for VARs," *International Economic Review*, 45, 643–673. [367]
- (2011), "Bayesian Macroeconometrics," in *The Oxford Handbook of Bayesian Econometrics*, eds. J. Geweke, G. Koop, and H. van Dijk, Oxford, UK: Oxford University Press, pp. 293–389. [369,371,372]
- (2013), "DSGE Model-Based Forecasting," in *Handbook of Economic Forecasting* (vol. 2), ed. G. Elliott, and A. Timmermann, Amsterdam, Netherlands: Elsevier. [378]
- Del Negro, M., Schorfheide, F., Smets, F., and Wouters, R. (2007), "On the Fit of New Keynesian Models," *Journal of Business and Economic Statistics*, 25, 123–162. [367]
- Doan, T., Litterman, R., and Sims, C. A. (1984), "Forecasting and Conditional Projection Using Realistic Prior Distributions," *Econometric Reviews*, 3, 1–100. [367,369,374]
- Durbin, J., and Koopman, S. J. (2001), *Time Series Analysis by State Space Methods*, Oxford, UK: Oxford University Press. [367,369]
- Faust, J., and Wright, J. (2013), "Forecasting Inflation," in *Handbook of Economic Forecasting* (vol. 2), eds. G. Elliott, and A. Timmermann, Amsterdam, Netherlands: Elsevier. [367]
- Foroni, C., Marcellino, M., and Schumacher, C. (2013), "U-MIDAS: MIDAS Regressions With Unrestricted Lag Polynomials," *Journal of the Royal Statistical Society, Series A*, forthcoming. Available at <http://onlinelibrary.wiley.com/doi/10.1111/rssa.12043/pdf>. [375]
- Frale, C., Marcellino, M., Mazzi, G. L., and Proietti, T. (2011), "EUROMIND: A Monthly Indicator of the Euro Area Economic Conditions," *Journal of the Royal Statistical Society, Series A*, 174, 439–470. [379]
- Geweke, J. (1999), "Using Simulation Methods for Bayesian Econometric Models: Inference, Development, and Communication," *Econometric Reviews*, 18, 1–126. [367,370,371]
- Ghysels, E. (2012), "Macroeconomics and the Reality of Mixed Frequency Data," unpublished manuscript, University of North Carolina at Chapel Hill. [368]
- Ghysels, E., Sinko, A., and Valkanov, R. (2007), "MIDAS Regressions: Further Results and New Directions," *Econometric Reviews*, 26, 53–90. [368]
- Giannone, D., Lenza, M., and Primiceri, G. (2014), "Prior Selection for Vector Autoregressions," *Review of Economics and Statistics*, forthcoming. Available at <http://www.mitpressjournals.org/action/doSearch?AllField=primiceri>. [367,370,373]
- Giannone, D., Reichlin, L., and Small, D. (2008), "Nowcasting: The Real-Time Informational Content of Macroeconomic Data," *Journal of Monetary Economics*, 55, 665–676. [367,369]
- Giordani, P., Pitt, M. K., and Kohn, R. (2011), "Bayesian Inference for Time Series State Space Models," in *The Oxford Handbook of Bayesian Econometrics*, eds. J. Geweke, G. Koop, and H. K. van Dijk, Oxford, UK: Oxford University Press, pp. 61–124. [370]
- Harvey, A. C. (1989), *Forecasting, Structural Time Series Models, and the Kalman Filter*, Cambridge, UK: Cambridge University Press. [367]
- Karlsso, S. (2013), "Forecasting With Bayesian Vector Autoregressions," in *Handbook of Economic Forecasting* (vol. 2), eds. G. Elliott and A. Timmermann, Amsterdam, The Netherlands: Elsevier. [367]

- Kuzin, V., Marcellino, M., and Schumacher, C. (2011), "MIDAS vs. Mixed-Frequency VAR: Nowcasting GDP in the Euro Area," *International Journal of Forecasting*, 27, 529–542. [368]
- Litterman, R. (1986), "Forecasting With Bayesian Vector Autoregressions — Five Years of Experience," *Journal of Business and Economic Statistics*, 4, 25–38. [367]
- (1980), "Techniques for Forecasting With Vector Autoregressions," Ph.D. thesis, University of Minnesota. [369]
- Mariano, R. S., and Murasawa, Y. (2003), "A New Coincident Index of Business Cycles Based on Monthly and Quarterly Series," *Journal of Applied Econometrics*, 18, 427–443. [367,369]
- Phillips, P. C. B. (1996), "Econometric Model Determination," *Econometrica*, 64, 763–812. [367]
- Rodriguez, A., and Puggioni, G. (2010), "Mixed Frequency Models: Bayesian Approaches to Estimation and Prediction," *International Journal of Forecasting*, 26, 293–311. [368]
- Sims, C. A., and Zha, T. (1998), "Bayesian Methods for Dynamic Multivariate Models," *International Economic Review*, 39, 949–968. [367,369]
- Stock, J. H., and Watson, M. W. (2010), "Monthly Estimates of Real GDP/GDI: Research Memorandum," unpublished manuscript, Harvard University and Princeton University. [379]
- Waggoner, D. F., and Zha, T. (1999), "Conditional Forecasts in Dynamic Multivariate Models," *Review of Economics and Statistics*, 81, 639–651. [367]