Real-Time Forecasting with a Mixed Frequency VAR

Frank Schorfheide*
University of Pennsylvania, CEPR, and NBER

Dongho Song
University of Pennsylvania

December 5, 2011

Abstract

This paper develops a vector autoregression (VAR) for macroeconomic time series which are observed at mixed frequencies – quarterly and monthly. The mixed-frequency VAR is cast in state-space form and estimated with Bayesian methods under a Minnesota-style prior. Using a real-time data set, we generate and evaluate forecasts from the mixed-frequency VAR and compare them to forecasts from a VAR that is estimated based on data time-aggregated to quarterly frequency. We document how information that becomes available within the quarter alters the forecasts in real time. JEL CLASSIFICATION: C11, C32, C53.

KEY WORDS: Bayesian Methods, Real-Time Data, Macroeconomic Forecasting, Vector Autoregressions

*Correspondence: Department of Economics, 3718 Locust Walk, University of Pennsylvania, Philadelphia, PA 19104-6297. Email: schorf@ssc.upenn.edu (Frank Schorfheide) and donghos@sas.upenn.edu (Dongho Song). We thank Frank Diebold, Kei-Mu Yi and seminar participants at 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 data sets for the recursive estimation of forecasting models. Financial support from the Federal Reserve Bank of Minneapolis is gratefully acknowledged.

1 Introduction

In macroeconomic applications, vector autoregressions (VARs) are typically estimated either exclusively based on quarterly observations or exclusively based on monthly observations. In a forecasting setting the advantage of using quarterly observations is that the set of macroeconomic series that could potentially be included in the VAR is larger. In particular, 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, since many important indicators, e.g. unemployment, prices, and interest rates get updated by the statistical agencies within each quarter. To exploit the respective advantages of both monthly and quarterly VARs, this paper 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 main contribution of this paper is an empirical one. We compile a real-time data set for an eleven-variable VAR that includes observations on real aggregate activity, prices, and financial variables. Using this data set, we recursively estimate our MF-VAR and compare its forecasting performance to a standard VAR in which all series are time-aggregated to quarterly frequency (QF-VAR). We carefully document how within-quarter information from variables that are observed at monthly frequency, can drastically increase the precision of GDP growth, unemployment, inflation, and interest rate forecasts of quarterly averages in the short-run. Over a one- to two-year horizon, the gain from using monthly information tends to vanish. We also compare real-time MF-VAR forecasts to the Greenbook forecasts prepared by the Board of Governors prior to meetings of the Federal Open Market Committee (FOMC). We assess the MF-VAR density forecasts based on probability integral transformations and conduct a small case-study and compare how monthly information altered and improved density forecasts generated during the 2008-09 recession.

To cope with the high dimensionality of the parameter space, the MF-VAR is equipped with a Minnesota prior, e.g. Sims and Zha (1998) or Del Negro and Schorfheide (2011), and estimated using Bayesian methods. By and large we are building on existing approaches of treating missing observations in state-space models (see, for instance, the textbook treatment by Durbin and Koopman (2001)). More specifically, we use data-augmentation to 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.

To the extent that there exist very few studies that estimate MF-VARs, the implementation of the Bayesian inference contains several noteworthy aspects. In the filtering/smoothing step of the Gibbs sampler we are alternating between different state-space representations of the MF-VAR in order keep the vector of latent state variables as small as possible while at the same time being able to handle irregular patterns of missing monthly variables toward the end of the estimation sample (ragged edges). As is common for the estimation of Bayesian VARs on single-frequency data, the prior is indexed by hyperparameters that are selected in a data-driven way by maximizing the marginal likelihood function.

Our paper is related to several strands of the time series literature. An alternative Gibbs sampling approach for the coefficients in an MF-VAR is explored in Eraker, Chiu, Foerster, Kim, and Seoane (2011). Their algorithm also iterates over the conditional posterior distributions of the VAR parameters and the missing monthly observations, but utilizes a different procedure to draw the missing observations. The focus of their paper is on parameter estimation rather than forecasting. The authors link the coefficients of the MF-VAR to the coefficients of a QF-VAR via a transformation. Eraker, Chiu, Foerster, Kim, and Seoane (2011) then compare 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 utilized in the estimation of dynamic factor models (DFMs). Mariano and Murasawa (2003) apply 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) develop 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) use 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 mixed data sampling (MIDAS) regression, to exploit high-frequency information without having to estimate a state-space model. To cope with a 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.

Bai, Ghysels, and Wright (2011) examine the relationship between MIDAS regressions and state-space models applied to mixed-frequency data. They consider dynamic factor models and characterize conditions under which the MIDAS regression exactly replicates the steady state Kalman filter weights on lagged observables. They conclude 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) compare the accuracy of Euro Area GDP growth forecasts from MIDAS regressions and MF-VARs estimated by maximum likelihood. The authors find that the relative performances of MIDAS and MF-VAR forecasts differ depending on the predictors and forecast horizons. Overall, the authors do not find a clear winner in terms of forecasting performance.

The remainder of this paper is organized as follows. Section 2 presents the state-space representation of the MF-VAR and discusses Bayesian inference and forecasting. The empirical results are presented in Section 3. We begin with a description of the real-time data set used for the recursive forecast evaluation. We then document how the within-quarter monthly information reduces the root-mean-squared error (RMSE) of VAR forecasts. Finally we evaluate MF-VAR density forecasts based on probability integral transformations and illustrate how these forecasts evolved during the 2008-09 recession. We conclude in Section 4. The Online Appendix provides detailed information about the Bayesian computations, the construction of the data set, as well as additional empirical results.

2 A Mixed-Frequency Vector Autoregression

The MF-VAR considered in this paper is based on a standard constant-parameter VAR in which the length of the time period is one month. Since some macroeconomic time series, e.g. 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. In order to ease the exposition, we use a representation with a state vector that includes even those variables that are observable at monthly frequency, e.g. the aggregate price level, the unemployment rate, and the interest rate. A computationally more efficient representation is presented in the Appendix. Bayesian inference and forecasting are discussed in Section 2.2.

Throughout this paper we use $Y_{t_0:t_1}$ to denote the sequence of observations or random variables $\{y_{t_0}, \ldots, y_{t_1}\}$. 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 an 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:

$$x_t = \Phi_1 x_{t-1} + \ldots + \Phi_p x_{t-p} + \Phi_c + u_t, \quad u_t \sim iidN\big(0, \Sigma\big). \tag{1}$$

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, e.g. the consumer price index and the unemployment rate, and the $n_q \times 1$ vector $x_{q,t}$ is comprised of unobserved monthly variables that are only published at quarterly frequency, e.g. GDP. Define $z_t = [x'_t, \ldots, x'_{t-p+1}]'$ and $\Phi = [\Phi_1, \ldots, \Phi_p, \Phi_c]'$. Write the VAR in (1) in companion form as

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

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 $l = 1, \ldots, p-1$. The $n \times n$ upper-left submatrix of Ω equals Σ and all other elements are zero. (2) is the state-transition equation of the MF-VAR.

We proceed by describing the measurement equation. 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 end of a quarter and for which all quarterly observations are available.[1] The vector of monthly series $x_{m,t}$ is observed every month. If the actual observations are denoted by $y_{m,t}$ then we obtain

$$y_{m,t} = x_{m,t}, \quad t = 1, \ldots, T_b. \tag{3}$$

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

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

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

$$y_{q,t} = M_{q,t}\tilde{y}_{q,t} = M_{q,t}\Lambda_{qz} z_t. \tag{5}$$

For periods $t = T_b + 1, \ldots, T$ no additional observations of the quarterly time series are available. 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 prior to period T and let $M_{m,t}$ be a deterministic sequence of selection matrices such that (3) can be extended to

$$y_{m,t} = M_{m,t} x_{m,t}, \quad t = T_b + 1, \ldots, T. \tag{6}$$

Notice that 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

$$y_t = M_t \Lambda_z z_t, \quad t = 1, \ldots, T. \tag{7}$$

[1] The subscript b stands for *balanced* sample.

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

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 pre-sample $Y_{-p+1:0}$ 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 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

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

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 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

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

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, e.g., 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 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 3.3) can be computed from the empirical distribution of the simulated trajectories.

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 and Σ. As $\lambda \longrightarrow \infty$, 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. From a practitioner's view, choosing λ based on the marginal likelihood function

$$\begin{aligned} & p(Y_{1:T}|Y_{-p+1:0}, \lambda) \qquad (10) \\ &= \int p(Y_{1:T}, Z_{0:T}, \Phi, \Sigma | Y_{-p+1:0}, \lambda) d(\Phi, \Sigma, Z_{0:T}) \\ &= \int p(Y_{1:T}|Z_{0:T}) \left[\int p(Z_{1:T}|z_0, \Phi, \Sigma) p(\Phi, \Sigma|\lambda) d(\Phi, \Sigma) \right] p(z_0|Y_{-p+1:0}) dZ_{0:T} \end{aligned}$$

tends to work well for forecasting purposes (see Giannone, Lenza, and Primiceri (2010) for a recent study). The log marginal likelihood $p(Y_{1:T}|Y_{-p+1:0}, \lambda)$ can be interpreted as the sum of one-step ahead predictive scores:

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

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.

From (10) we see that the computation of the marginal likelihood involves integrating out the latent state which is very time-consuming. As a short-cut we use the posterior median values of the latent states $\hat{Z}_{0:T}$ and approximate the marginal likelihood as follows to reduce the computational burden:

$$p(Y_{1:T}|Y_{-p+1:0}, \lambda) \simeq \int p(\hat{Z}_{1:T}|\hat{z}_0, \Phi, \Sigma) p(\Phi, \Sigma|\lambda) d(\Phi, \Sigma). \qquad (12)$$

An analytical expression for the approximate marginal likelihood can be obtained by using the normalization constants for the MNIW distribution and is provided in Section 2 of

Del Negro and Schorfheide (2011). We maximize the right-hand side of (12) with respect to λ over a grid. More specifically, we start from an initial choice of λ to generate an initial sequence $\hat{Z}_{0:T}$. Subsequently, the marginal likelihood is maximized with respect to λ and we generate a new sequence $\hat{Z}_{0:T}$. These steps are repeated until the changes in $\hat{\lambda}$ are negligible. This method appears to be sufficiently robust and reliable in selecting the hyperparameters.

3 Empirical Analysis

The goal of the empirical analysis is to study the extent to which the incorporation of monthly observations via a MF-VAR model improves upon forecasts generated with a VAR that is based on time-aggregated quarterly data (QF-VAR). We consider a MF-VAR and a QF-VAR 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 Fund Rate (FF), Treasury Bond Yield (TB), S&P 500 Index (SP500). Precise data definitions are provided in the 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 in order to make them commensurable in scale to the other log transformed variables. Based on some preliminary exploration of marginal likelihood functions, we set the number of lags in the (monthly) state-transition of the MF-VAR to $p = 6$ and the number of lags in the QF-VAR to 2.

The remainder of this section is organized as follows. The construction of our real-time data set and the grouping of forecast origins according to within-quarter information is discussed in Section 3.1. We present RMSE statistics in Section 3.2 and assess VAR density forecasts based on probability integral transformations in Section 3.3. Finally, we compare forecasts from the MF-VAR and QF-VAR during the 2008-09 recession in Section 3.4.

3.1 Real-Time Data Set and Information Structure

To compare the empirical performance of our MF-VAR with a standard QF-VAR we conduct a pseudo-out-of-sample forecast experiment. In this forecast experiment we consider an increasing sequence of estimation samples $Y_{-p+1:T}$, $T = T_{min}, \ldots, T_{max}$, and generate forecasts for periods $T+1, \ldots, T+H$. The maximum forecast horizon H is chosen to be 24 months. The period $t = 1$ corresponds to 1968:M1. We align T_{min} with 1997:M7 and T_{max} with 2009:M3, which yields 141 estimation samples.[2] The estimation samples are constructed from real-time data sets, 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, time series 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 has potentially access to a triangular array of data $Y^1_{-p+1:1}, Y^2_{-p+1:2}, ..., Y^T_{-p+1:T}$. 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 $Y^T_{-p+1:T} = \{y^T_{-p+1}, \ldots, y^T_T\}$. As in Section 2 we are using the convention that the vector y^T_t contains only the subset of the eleven variables listed above for which observations are available at the end of month T.

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 $y^{T_*}_{t+h}$. The former might do a better job 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 $T_* =$ 2011:M7 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.

In order to assess the usefulness of within-quarter information from monthly variables we sort the forecast origins $T_{min}, \ldots, T_{max}$ into three groups that reflect different within-quarter information sets. Forecast error statistics are computed for each group separately. Before discussing the classification of forecast origins, we begin with a brief review of NIPA release

[2] We eliminated four of the 141 samples because the real-time data for PCE were incomplete.

Figure 1: Classification of the Information Set

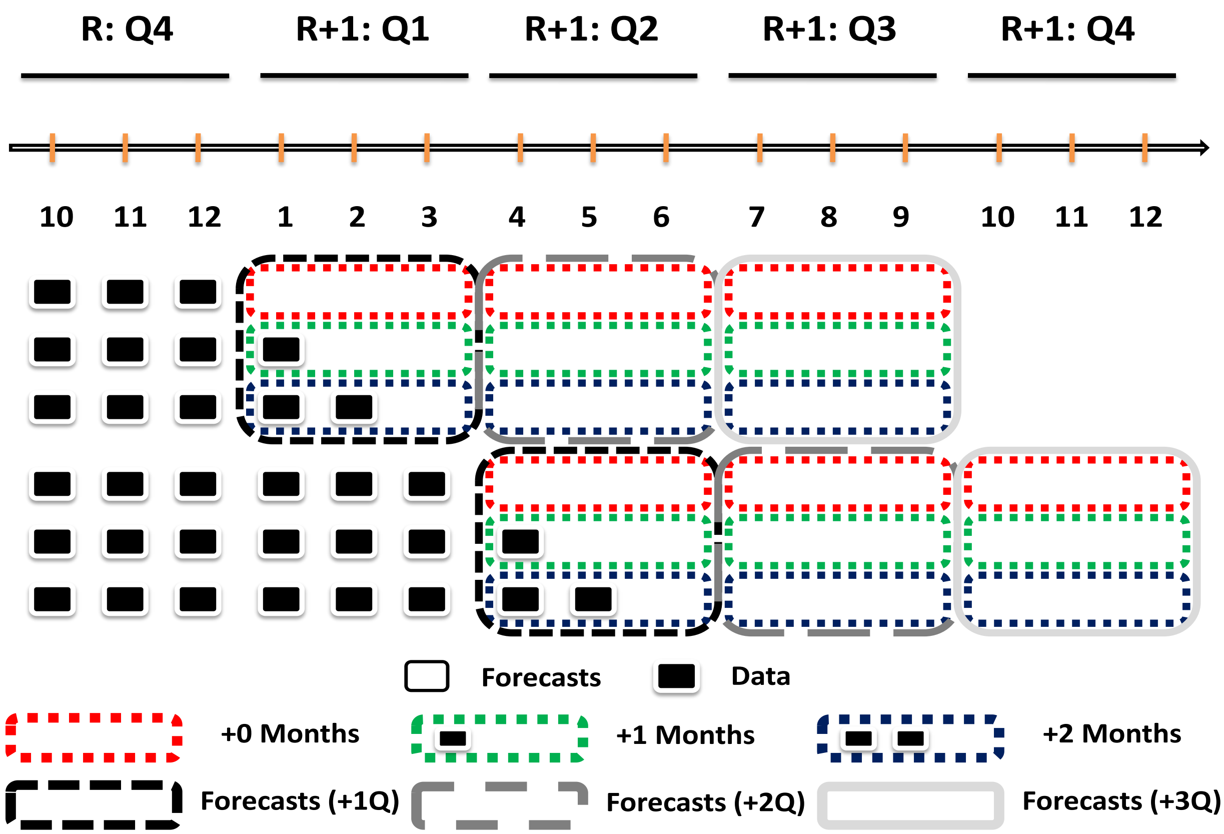

dates. For concreteness, consider GDP for 1994:Q4. The Bureau of Economic Analysis (BEA) published an *advance* estimate of 1997:Q4 GDP at the end of 1998:M1 (January). A *preliminary* estimate was subsequently published by the end of 1998:M2 (February). At last, a *final* release was available with a 3 month delay. Thus, a QF-VAR estimated in 1997:M12 cannot use any information about 1997:Q4 GDP. QF-VARs estimated at the end of January, February, and March, on the other hand, can be based, respectively, on the advance, preliminary, and final estimate of 1997:Q4 GDP.

While the QF-VAR forecasts do not use any within-quarter monthly information, the MF-VAR forecasts can exploit monthly observations that become available between 1998:M1 to 1998:M3. Figure 1 illustrates our classification of information sets with respect to non-financial variables for prototypical years R and $R+1$. For instance, collecting all available non-financial data at the end of January 1998 ($R+1$), leaves the forecaster with the advance estimate of 1997:Q4 GDP as well as values for the monthly macroeconomic indicators until 1997:M12. 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 non-financial monthly variables. At the end of February 1998, the forecaster has access to the observations of 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 third release of GDP in the information set. Unlike the non-financial 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 (S&P500). The typical information sets for the three subgroups of forecast origins are summarized in Table 1.

Unfortunately, due to variation in release dates, not all 140 estimation samples mimic the information structure in Table 1. For 44 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 26 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 pre-date its release. Thus, for the 26 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 Appendix.

Before presenting the pseudo-out-of-sample forecast results we briefly examine the monthly GDP series that is implicitly extracted during the smoothing step of the Gibbs sampler (see Section 2.2) from the eleven macroeconomic time series that enter the MF-VAR. A time series plot of monthly GDP growth is depicted in Figure 2. 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

Table 1: Illustration of Information Sets

January (+0 Months)

		UNR	HRS	CPI	IP	PCE	FF	TB	SP500	GDP	INVFIX	GOV
Q4	M10	X	X	X	X	X	X	X	X	QAv	QAv	QAv
Q4	M11	X	X	X	X	X	X	X	X	QAv	QAv	QAv
Q4	M12	X	X	X	X	X	X	X	X	QAv	QAv	QAv
Q1	M1	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	X	X	X	$\emptyset$	$\emptyset$	$\emptyset$

February (+1 Month)

		UNR	HRS	CPI	IP	PCE	FF	TB	SP500	GDP	INVFIX	GOV
Q4	M11	X	X	X	X	X	X	X	X	QAv	QAv	QAv
Q4	M12	X	X	X	X	X	X	X	X	QAv	QAv	QAv
Q1	M1	X	X	X	X	X	X	X	X	$\emptyset$	$\emptyset$	$\emptyset$
Q1	M2	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	X	X	X	$\emptyset$	$\emptyset$	$\emptyset$

March (+2 Month)

		UNR	HRS	CPI	IP	PCE	FF	TB	SP500	GDP	INVFIX	GOV
Q4	M12	X	X	X	X	X	X	X	X	QAv	QAv	QAv
Q1	M1	X	X	X	X	X	X	X	X	$\emptyset$	$\emptyset$	$\emptyset$
Q1	M2	X	X	X	X	X	X	X	X	$\emptyset$	$\emptyset$	$\emptyset$
Q1	M3	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	X	X	X	$\emptyset$	$\emptyset$	$\emptyset$

Notes: $\emptyset$ 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 Month” group: March, June, September, December.

Figure 2: Monthly GDP Growth (Scaled to a Quarterly Rate)

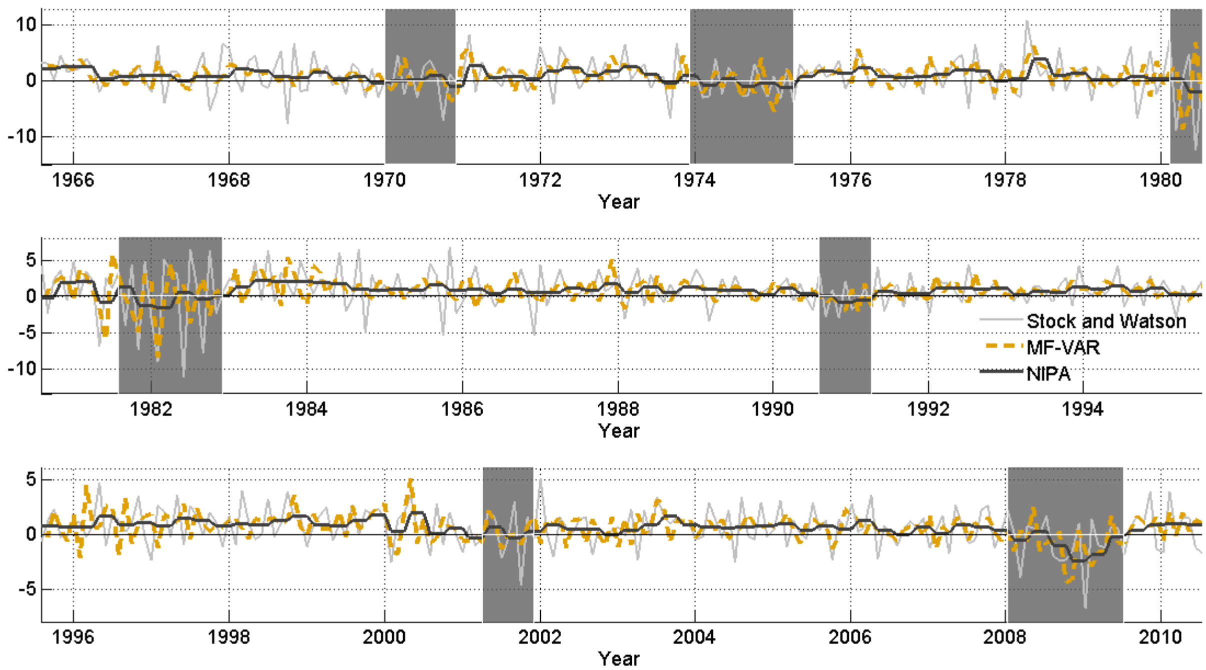

month we then plot the median growth rate across the simulated trajectories. We overlay monthly GDP growth rates published by Stock and Watson (2010), who combine monthly information about GDP components to distribute quarterly GDP across the three months of the quarter.[3] Moreover, we plot growth rates 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.

[3]Frale, Marcellino, Mazzi, and Proietti (2011) use a similar approach to construct a monthly GDP series for the Euro Area.

3.2 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. Hyperparameters λ for the two VARs are selected by maximizing the respective marginal likelihood functions $p(Y_{1:T}|Y_{-p+1:0},\lambda)$ for the first estimation sample $Y_{-p+1:T_{min}}$. The hyperparameters are held constant for the subsequent samples ($T > T_{min}$). Results for GDP growth (GDP), unemployment (UNR), inflation (INF), and the Federal Funds rate (FF) are reported in Figure 3. The figure depicts relative RMSEs defined as

$$RelativeRMSE(i|h) = 100 \times \frac{RMSE(i|h) - RMSE_{Benchmark}(i|h)}{RMSE_{Benchmark}(i|h)}, \tag{13}$$

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.

For all four series the use of monthly information via the MF-VAR leads to a substantial RMSE reduction in the short-run. The "+2" GDP growth nowcasts have a 25% lower RMSE than the QF-VAR nowcasts. For the "+1" group and the "+0" group the reductions are 17% and 12% respectively. Thus, the monthly series provide important information in the short run. As the forecast horizon increases to $h = 4$ the relative ranking between the "+0" and "+1" forecasts becomes ambiguous and the QF-VAR catches up with the MF-VAR. For horizons $h \geq 4$. The precision of QF-VAR and MF-VAR GDP growth forecasts is essentially identical.

Not surprisingly, 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 the quarterly frequency. At the nowcast horizon the MF-VAR is able to improve over the precision of the QF-VAR for the "+2" forecasts by 65% for unemployment, 70% for inflation, and 100% for the Federal Funds rate. Recall that "+2" 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" group range from 40% (unemployment) to 80% (Federal Funds rate).

Figure 3: Relative RMSEs of 11-Variable MF-VAR versus QF-VAR

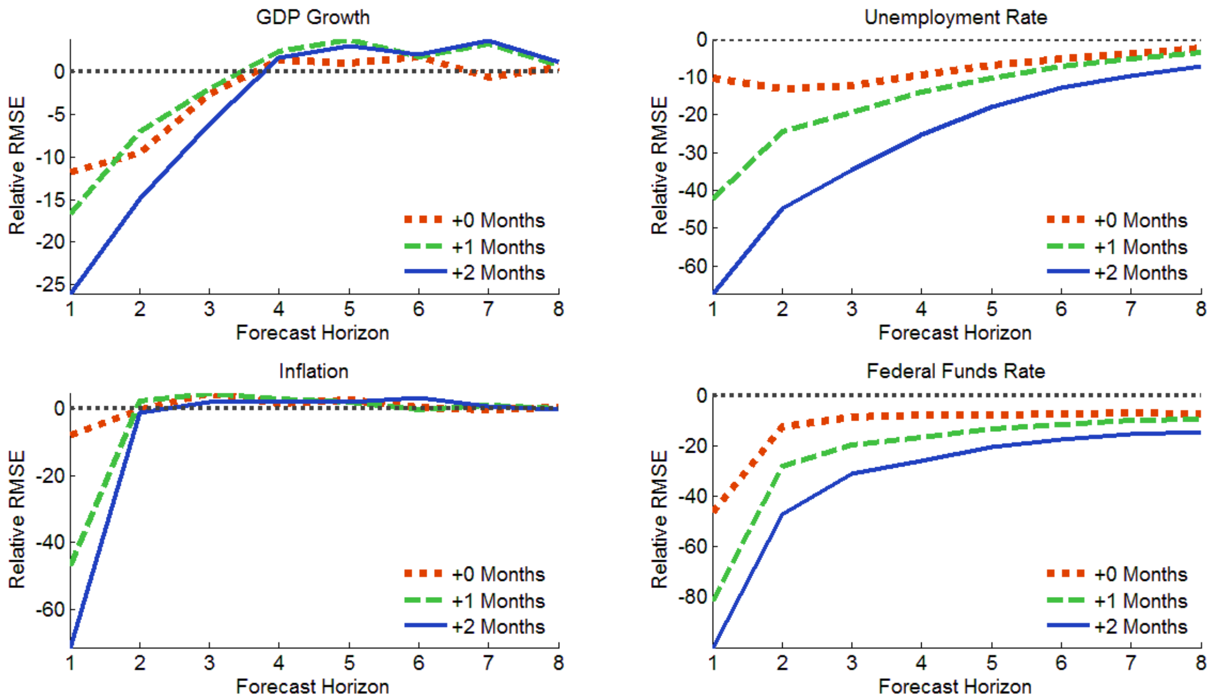

Interestingly the nowcast improvements of the "+0"group for unemployment and inflation are only 10%. The gains from using monthly information tend to persist for unemployment and interest rates as the forecast horizon h increases. Only for a two-year horizon, $h = 8$, the QF-VAR catches up with the MF-VAR. For inflation, monthly observations generate no improvements of forecast performance beyond the nowcast horizon.

MF-VAR versus QF-AR. We also compare the MF-VAR forecasts to univariate forecasts from AR(2) models estimated on quarterly-frequency data (QF-AR). Just as the QF-VAR, the QF-AR models are equipped with a Minnesota prior and estimated on time-aggregated quarterly data. The hyperparameters are chosen to maximize the marginal likelihood function for the first estimation sample and kept constant subsequently. Relative RMSEs, now with the QF-AR models as benchmarks, are plotted in Figure 4. The results are qualitatively similar to the ones depicted in Figure 3. At the nowcast horizon the monthly information used by the MF-VAR leads to a substantial improvement in forecast accuracy also in comparison to univariate quarterly models. At the medium term horizon a comparison

Figure 4: Relative RMSEs of 11-Variable MF-VAR versus QF-AR

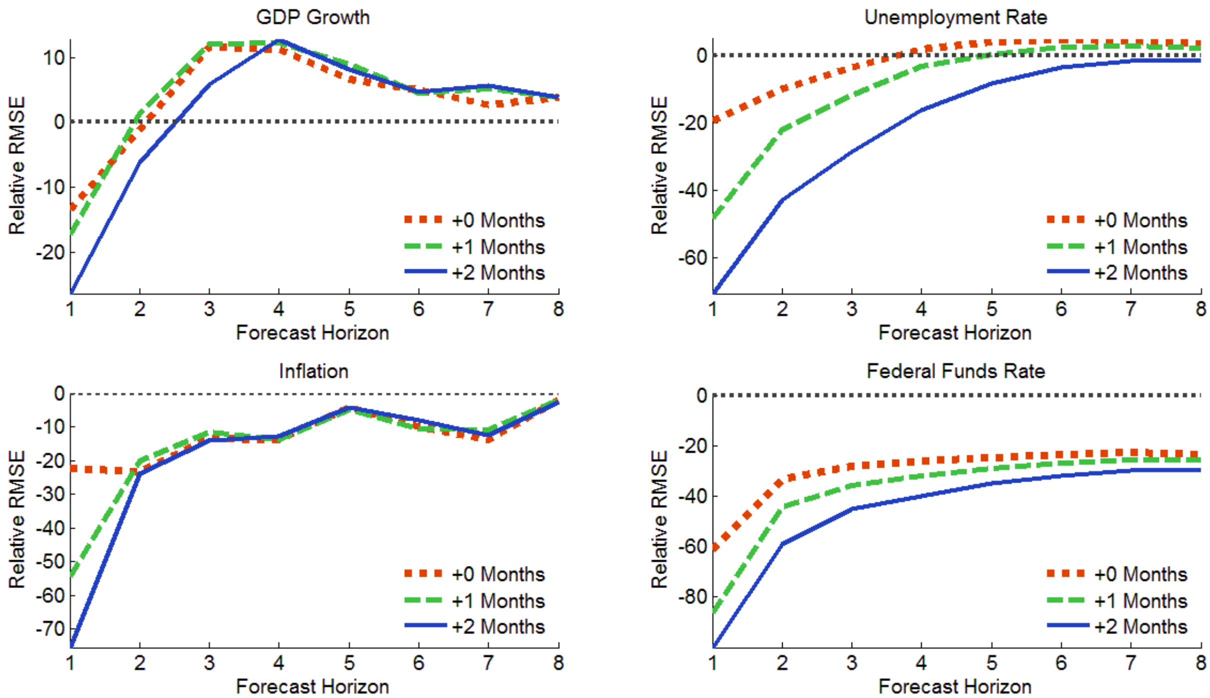

between Figures 4 and 3 reveals differences in the forecast performance of the QF-AR and the QF-VAR models. The univariate models tend to be more accurate for GDP growth and unemployment, whereas the multivariate QF-VAR tends to dominate for inflation and the Federal Funds rate. Overall, we find that the use of within quarter information on monthly indicators can result in marked reductions in RMSEs.

MF-VAR versus Greenbook Forecasts. We also compare the MF-VAR forecasts to Greenbook forecasts, prepared by the staff of the Board of Governors for the FOMC meetings. Greenbook forecasts are publicly available with a five-year delay. Our comparison involves 63 Greenbook forecasts from March 19, 1997 to December 8, 2004. We repeat the recursive estimation of the MF-VAR to align the information that is used for the MF-VAR forecasts with the information that was available to the staff of the Board of Governors. As in the previous analysis, period $t = 1$ corresponds to 1968:M1. One important difference is that financial data from the month in which the forecast is made are not included in the MF-VAR's information set.

Figure 5: RMSEs of 11-Variable MF-VAR versus Greenbook

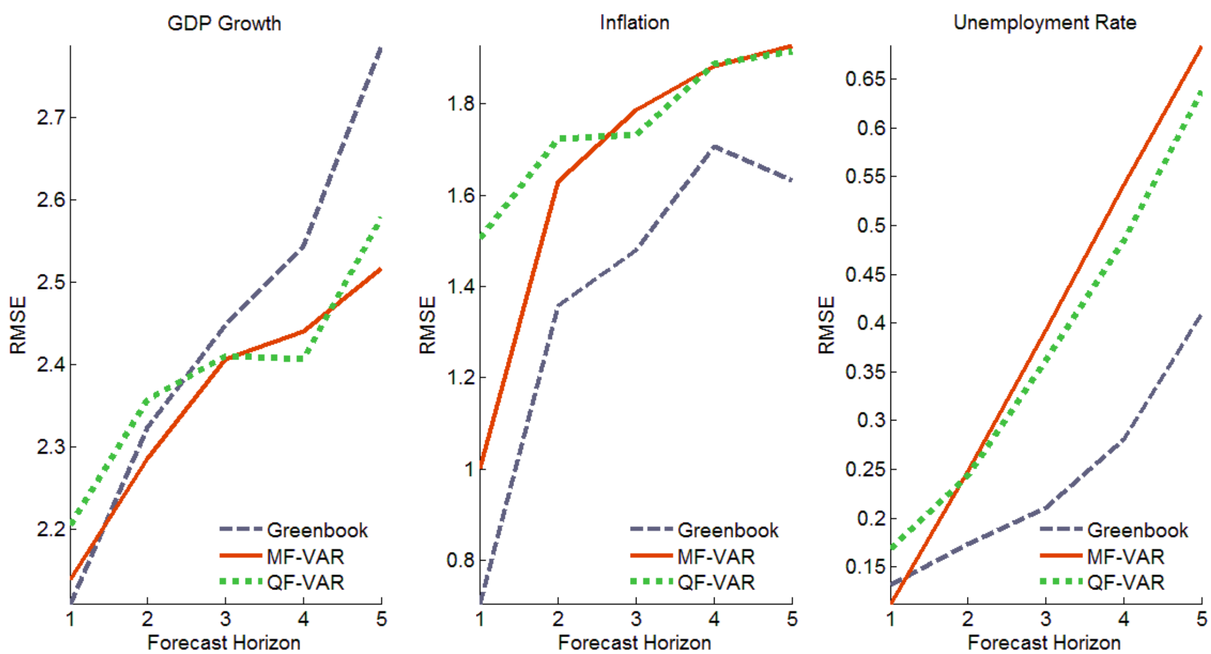

Results are plotted in Figure 5, which depicts absolute RMSEs for quarter-on-quarter GDP growth (annualized), CPI inflation (annualized), and the unemployment rate. The figure also shows RMSEs for the QF-VAR. Unlike in the previous figures, we are now pooling the forecast errors for all estimation samples. As before, the forecasts from the MF-VAR attain a smaller RMSE than the QF-VAR forecasts in the short-run. For horizons $h \geq 3$ the two VARs deliver forecasts that are similarly accurate. While the MF-VAR produces GDP growth predictions that dominate those published in the Greenbook. The Greenbook forecasts for inflation and unemployment, on the other hand, are more precise than the MF-VAR forecasts.

3.3 MF-VAR Density Forecasts

The MF-VAR generates an entire predictive distribution for the future trajectories of the eleven macroeconomic variables. While, strictly speaking, predictive distributions in a Bayesian framework are subjective, it is desirable that predicted probabilities are consistent with observed frequencies if the forecast procedure is applied in a sequential setting. To assess the MF-VAR density forecasts, we construct probability integral transformations

Figure 6: PIT Histograms for 11-Variable MF-VAR

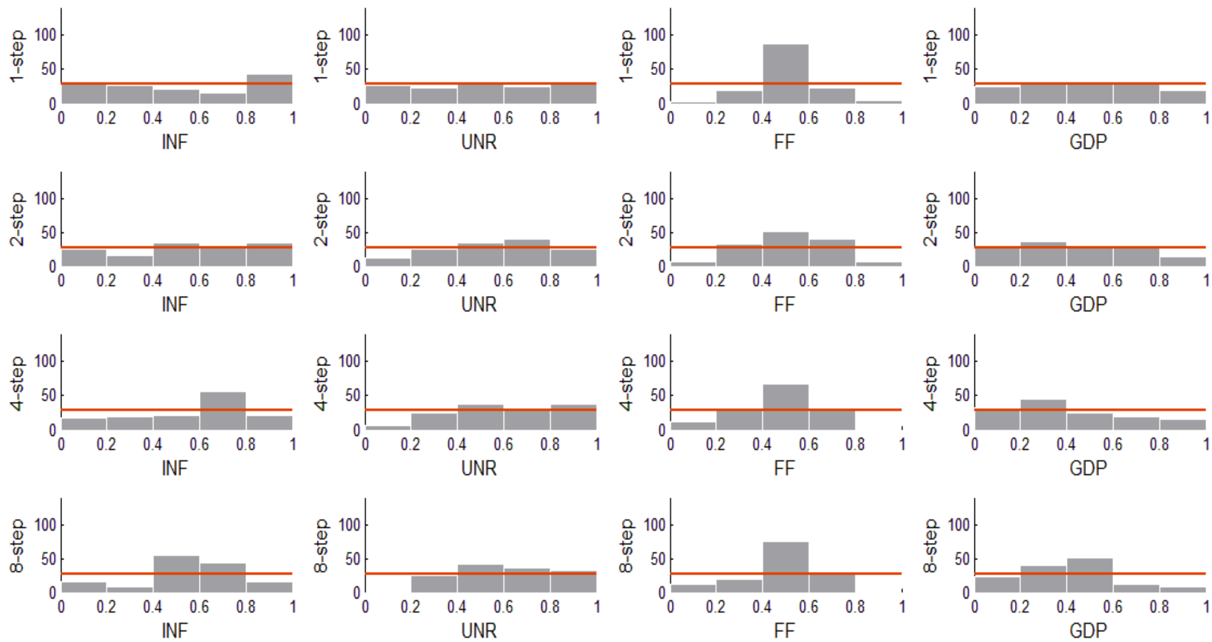

Notes: Probability integral transforms for forecasts of inflation (INF), unemployment rate (UNR), federal fund rate (FF), and GDP growth (GDP). The bars represent the frequency of PITs falling in each bin. The solid line marks 20 percent.

(PITs) from (univariate) marginal predictive densities. The probability integral transformation of an h-step ahead forecast of $y_{i,t+h}$ based on time t information is defined as

$$z_{i,h,t} = \int_{-\infty}^{y_{i,t+h}} p(\widetilde{y}_{i,t+h}|Y_{1:t}) d\widetilde{y}_{i,t+h}. \tag{14}$$

Starting with Dawid (1984) and Kling and Bessler (1989) the use of PITs has a fairly long tradition in the literature on density forecast evaluation. PITs, sometimes known as generalized residuals, are relatively easy to compute and facilitate comparisons among elements of a sequence of predictive distributions, each of which is distinct in that it conditions on the information available at the time of the prediction. It is shown in Rosenblatt (1952) and Diebold, Gunther, and Tay (1998) that for $h = 1$ the $z_{i,h,t}$'s are independent across time and uniformly distributed: $z_{i,h,t} \sim iid\mathcal{U}[0, 1]$. For $h > 1$ the PITs remain uniformly distributed but they are no longer independently distributed.

Figure 6 displays histograms for the PITs based on density forecasts from the MF-VAR. The PITs are computed from the empirical distribution of the simulated trajectories $Y_{T+1:T+H}$. To generate the histogram plots, the unit interval is divided into $J = 5$ equally sized subintervals and we depict the fraction of PITs (measured in percent) that fall in each bin. Since, under the predictive distribution, the PITs are uniformly distributed on the unit interval, we also plot the 20 percent line. For $h = 1$ (nowcast) and $h = 2$ (forecast for next quarter) the frequency of PITs falling in each of the five bins is close to 20% for inflation, unemployment, and output growth, indicating that the predictive densities are well calibrated.[4] The Federal Funds rate density forecasts, on the other hand, appear to be too diffuse, because of the small number of PITs falling into the 0-0.2 and 0.8-1 bins. Over longer horizons, specifically for $h = 4$ and $h = 8$, the deviations from uniformity become more pronounced. The Federal Funds rate density forecasts remain to diffuse and the MF-VAR tends to overpredict GDP growth and underpredict unemployment. We also computed PITs for the QF-VAR (reported in the Appendix) and found that deviations from uniformity tend to be larger than for the MF-VAR forecasts.

3.4 Predicting the Crisis: Interval Forecasts and Actuals

Finally, we examine how the use of monthly real-time information affected the VAR forecasts during the recent recession. We focus on the period from July to December 2008. Figures 7 to 9 depict real-time interval forecasts from the MF-VAR and the QF-VAR. Moreover, we plot actual values using the 2011:M7 data vintage. Each figure is divided into subpanels that correspond to particular estimation samples and forecast horizons. The first column of panels depicts forecasts from the "+0 Months" group and the second and third column correspond to "+1 Month" and "+2 Months" forecasts, respectively. Thus, each row indicates how monthly within-quarter information alters the density forecast.

The most striking feature of Figure 7 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 model. It is, at all times, outside of the 90% predictive interval. Interestingly, the MF-VAR does a reasonably good job forecasting output growth for 2009. After missing the large drop in the last quarter of 2008, the forecast is essentially back on track for the

[4] A Bayesian predictive check that formally assess the uniformity of PITs is developed in Herbst and Schorfheide (2011).

subsequent quarters. A comparison of the MF-VAR and QF-VAR forecast highlights how monthly information alters the within-quarter predictions. Notice from the bottom panels of Figure 7 that the QF-VAR forecasts do not stay constant within the quarter. The variation is caused by data revisions. As discussed in Section 3.1 each month new data releases for the previous quarter becomes 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. Even by December 2008 the QF-VAR nowcasts and forecasts show 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. For instance, throughout 2008:Q3 the GDP growth forecasts become more and more pessimistic.

Figure 8 depicts the evolution of inflation forecasts in the second half 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% inflation rate in the current quarter. The MF-VAR, on the other hand, picks up the deflation by November 2008 as it occurs. Finally, we consider the unemployment forecasts in Figure 9. Neither the QF-VAR nor the MF-VAR anticipate the large rise in unemployment between 2008:Q4 to 2009:Q3. However, due to the use of monthly data, the MF-VAR forecast adapts between January and March 2009 to the rising level of unemployment. In February and March 2009 the MF-VAR generates 90% predictive intervals for 2009:Q2 and Q3 that include unemployment rates near 10%. The QF-VAR, on the other hand, predicts that unemployment is unlikely to rise about 8.5%, which turned out to be incorrect.

Figure 7: GDP Growth Forecasts

MF-VAR

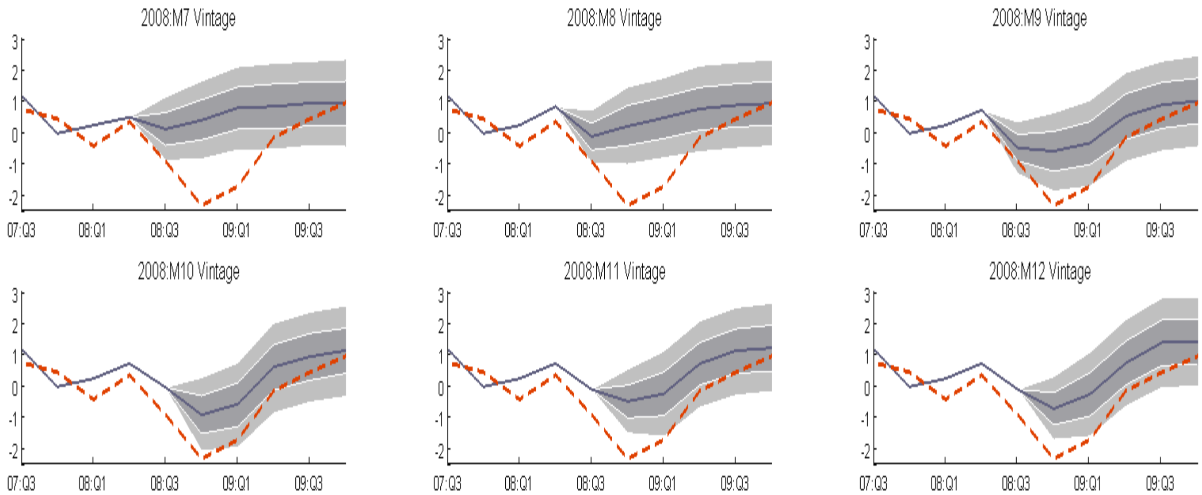

QF-VAR

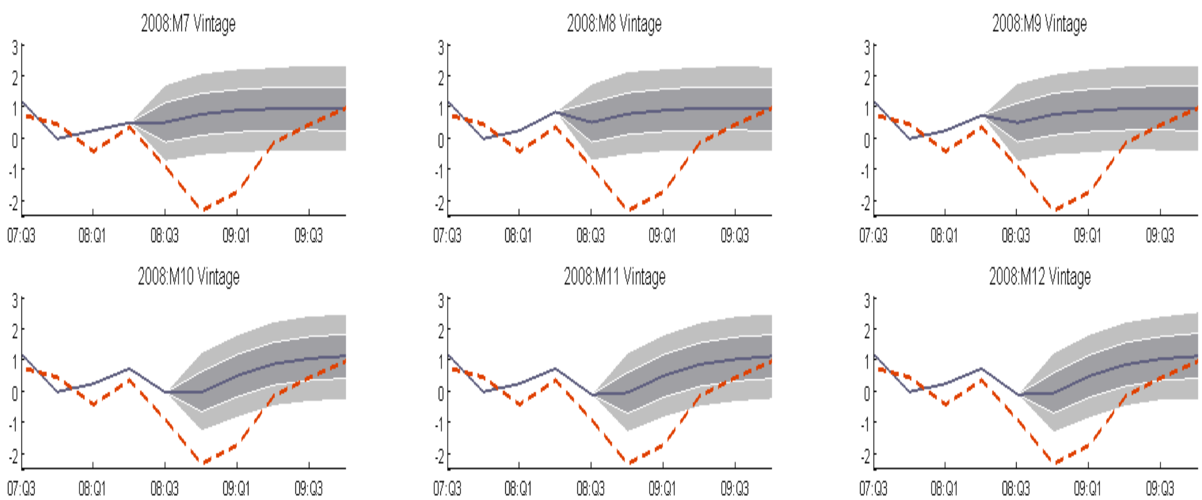

Notes: Actual values are from the $T_* = 2011 : M7$ data vintage and are denoted as the red dashed line. Starting from the leftmost column, we show the results of "+0 Months", "+1 Months", and "+2 Months" subgroups. The title in each subplot indicates the data vintage that are used in the estimation.

Figure 8: Inflation Forecasts

MF-VAR

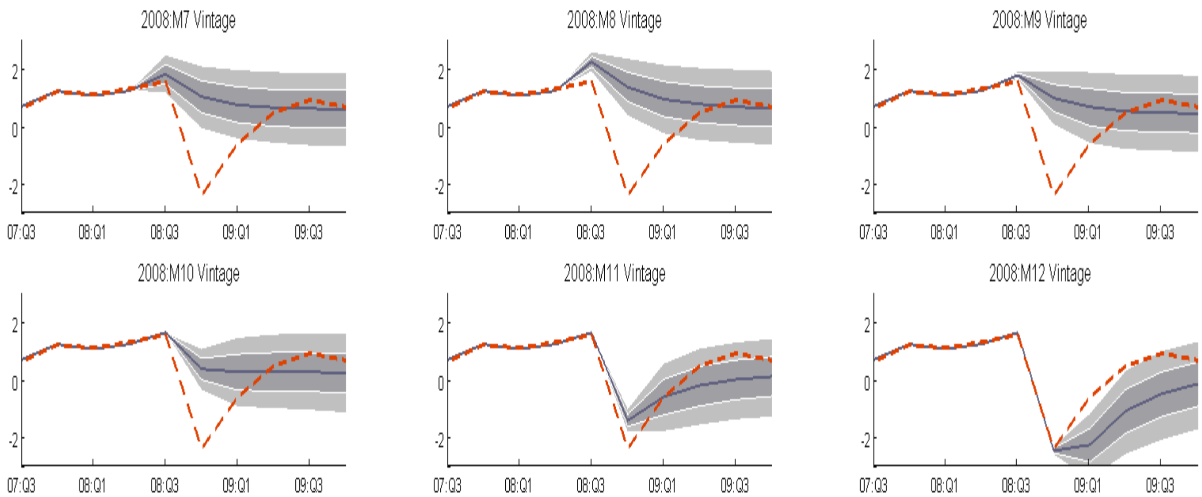

QF-VAR

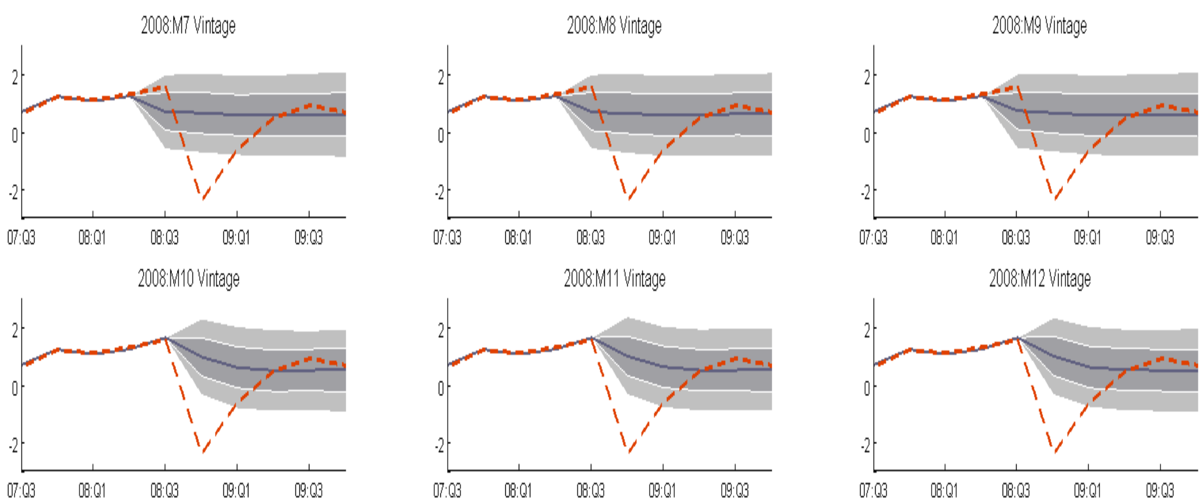

Notes: Actual values are from the $T_* = 2011 : M7$ data vintage and are denoted as the red dashed line. Starting from the leftmost column, we show the results of "+0 Months", "+1 Months", and "+2 Months" subgroups. The title in each subplot indicates the data vintage that are used in the estimation.

Figure 9: Unemployment Forecasts

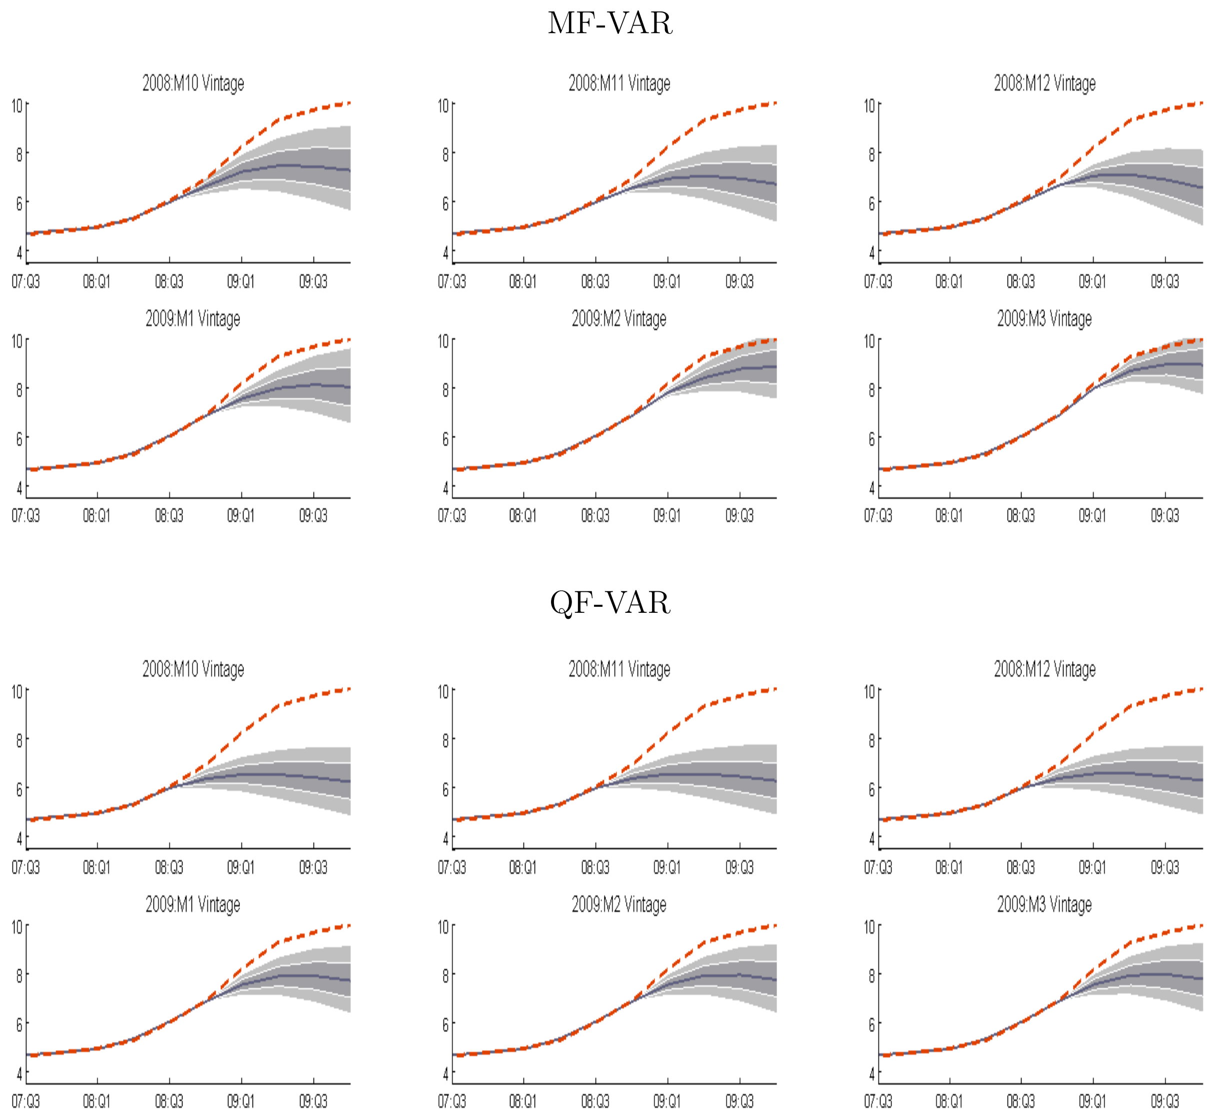

Notes: Actual values are from the $T_* = 2011 : M7$ data vintage and are denoted as the red dashed line. Starting from the leftmost column, we show the results of "+0 Months", "+1 Months", and "+2 Months" subgroups. The title in each subplot indicates the data vintage that are used in the estimation.

4 Conclusion

We have specified a VAR for observations that are observed at different frequencies, namely monthly and quarterly. Markov-Chain-Monte-Carlo methods were utilized 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 marginal likelihood function with respect to a low-dimensional vector of hyperparameters. 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. The short-term density forecasts appear to be well calibrated in the sense that the empirical distribution of probability integral transformations are nearly uniform. Over a longer horizon, on the other hand, there appear to be some deficiencies. Recent work by Clark (2011) suggests that real-time VAR density forecasts can be improved by adding stochastic volatility to the VAR. We plan to incorporate time-varying volatilities into our MF-VAR in future work.

References

Aruoba, B., F. Diebold, and C. Scotti (2009): "Real-Time Measurement of Business Conditions," *Journal of Business Economics & Statistics*, 27(4), 417–427.

Bai, J., E. Ghysels, and Wright (2011): "State Space Models and MIDAS Regressions," *Manuscript, Johns Hopkins University.*

Carter, C., and R. Kohn (1994): "On Gibbs Sampling for State Space Models," *Biometrika*, 81(3), 541–553.

Clark, T. E. (2011): "Real-Time Density Forecasts From Bayesian Vector Autoregressions With Stochastic Volatility," *Journal of Business Economics & Statistics*, 29(3), 327–341.

Dawid, A. (1984): "Statistical Theory: The Prequential Approach," *Journal of the Royal Statistical Society, Series A*, 147(2), 278–292.

Del Negro, M., and F. Schorfheide (2011): "Bayesian Macroeconometrics," in *The Oxford Handbook of Bayesian Econometrics*, ed. by J. Geweke, G. Koop, and H. van Dijk, pp. 293–389. Oxford University Press.

Diebold, F., T. Gunther, and A. Tay (1998): "Evaluating Density Forecasts with Applications to Financial Risk Management," *International Economic Review*, 39(4), 863–883.

Doan, T., R. Litterman, and C. A. Sims (1984): "Forecasting and Conditional Projections Using Realistic Prior Distributions," *Econometric Reviews*, 3(4), 1–100.

Durbin, J., and S. J. Koopman (2001): *Time Series Analysis by State Space Methods.* Oxford University Press.

Eraker, B., C. Chiu, A. Foerster, T. Kim, and H. Seoane (2011): "Bayesian Mixed Frequency VAR's," *Manuscript, University of Wisconsin.*

Frale, C., M. Marcellino, G. L. Mazzi, and T. Proietti (2011): "EUROMIND: A Monthly Indicator of the Euro Area Economic Conditions," *Journal of the Royal Statistical Society, Series A*, 174(2), 439–470.

Ghysels, E., A. Sinko, and R. Valkanov (2007): "The MIDAS Regressions: Further Results and New Directions," *Econometric Reviews*, 26(1), 53–90.

Giannone, D., M. Lenza, and G. Primiceri (2010): "Prior Selection for Vector Autoregressions," *Manuscript, European Centeral Bank and Northwestern University.*

Giannone, D., L. Reichlin, and D. Small (2008): "Nowcasting: The Real-Time Informational Content of Macroeconomic Data," *Journal of Monetary Economics*, 55, 665–676.

Giordani, P., M. K. Pitt, and R. Kohn (2011): "Bayesian Inference for Time Series State Space Models," in *The Oxford Handbook of Bayesian Econometrics*, ed. by J. Geweke, G. Koop, and H. K. van Dijk, pp. 61–124. Oxford University Press.

Herbst, E., and F. Schorfheide (2011): "Evaluating DSGE Model Forecasts of Comovements," *Manuscript, University of Pennsylvania.*

Kling, J., and D. Bessler (1989): "Calibration-Based Predictive Distributions: An Application of Prequential Analysis to Interest Rates, Money, Prices, and Output," *Journal of Business*, 62(4), 447–499.

Kuzin, V., M. Marcellino, and C. Schumacher (2011): "MIDAS versus Mixed-Frequency VAR: Nowcasting GDP in the Euro Area," *International Journal of Forecasting*, 27, 529–542.

Litterman, R. B. (1980): "Techniques for Forecasting with Vector Autoregressions," Ph.D. thesis, University of Minnesota.

Mariano, R. S., and Y. Murasawa (2003): "A New Coincident Index of Business Cycles Based on Monthly and Quarterly Series," *Journal of Applied Econometrics*, 18(4), 427–443.

Rosenblatt, M. (1952): "Remarks on a Multivariate Transformation," *Annals of Mathematical Statatistics*, 23(3), 470–472.

Sims, C. A., and T. Zha (1998): "Bayesian Methods for Dynamic Multivariate Models," *International Economic Review*, 39(4), 949–968.

Stock, J. H., and M. W. Watson (2010): "Monthly Estimates of Real GDP/GDI: Research Memorandum," *Manuscript, Harvard University and Princeton University.*

Online Appendix for *Real-Time Forecasting with a Mixed-Frequency VAR*

Frank Schorfheide and Dongho Song

Section A of this appendix provides details of the implementation of the Bayesian computations for the MF-VAR presented in the main text. Section B discusses the construction of the real-time data set. Finally, Section C of this appendix provides tables and figures with additional empirical results. References to equations, tables, and figures without an A, B, or C prefix refer to equations, tables, and figures in the main text.

A Implementation Details

Recall from the exposition in the main text (see Equation (9)) that the Bayesian computations are implemented with a Gibbs sampler that iterates over the conditional distributions:

$$p(\Phi, \Sigma | Z_{0:T}, Y_{-p+1:T}) \quad \text{and} \quad p(Z_{0:T} | \Phi, \Sigma, Y_{-p+1:T}).$$

Conditional on $Z_{0:T}$ the MF-VAR reduces to a standard linear Gaussian VAR with a conjugate prior. The reader is referred to Section 2 of the handbook chapter by Del Negro and Schorfheide (2011) for a detailed discussion of posterior inference for such a VAR.

We limit the exposition in this appendix to a brief presentation of the Minnesota prior and the hyperparameter selection (Section A.1). The sampling from the conditional posterior of $Z_{0:T}|(\Phi, \Sigma, Y_{-p+1:T})$ is implemented with a standard simulation smoother, discussed in detail, for instance in Carter and Kohn (1994), the state-space model textbook of Durbin and Koopman (2001), or the handbook chapter by Giordani, Pitt, and Kohn (2011). The only two aspects that of our implementation that deserve further discussion are the initialization (Section A.2) and the use of the more compact state-space representation for periods $t = 1, \ldots, T_b$ (Section A.3).

A.1 Minnesota Prior and Its Hyperparameters

To simplify the exposition, suppose that $n = 2$ and $p = 2$. A transposed version of (1) can be written as

$$x_t' = [x_{t-1}', x_{t-2}', 1]'\Phi + u_t' = w_t'\Phi + u_t', \quad u_t \sim iidN(0, \Sigma). \tag{A-1}$$

We generate the Minnesota prior by dummy observations (x_*, w_*) that are indexed by a 5×1 vector of hyperparameters λ with elements λ_i. Using a pre-sample, let $\underline{x}$ and $\underline{s}$ be $n \times 1$ vectors of means and standard deviations. For time series that are observed at monthly frequency the computation of pre-sample moments is straightforward. In order to obtain pre-sample means and standard deviations for those series that are observed at quarterly frequency, we simply equate $\underline{x_q}$ with the pre-sample mean of the observed quarterly values and set $\underline{s}$ equal to the pre-sample standard deviation of the observed quarterly series.

Dummy Observations for Φ_1.

$$\begin{bmatrix} \lambda_1 \underline{s}_1 & 0 \\ 0 & \lambda_1 \underline{s}_2 \end{bmatrix} = \begin{bmatrix} \lambda_1 \underline{s}_1 & 0 & 0 & 0 & 0 \\ 0 & \lambda_1 \underline{s}_2 & 0 & 0 & 0 \end{bmatrix} \Phi + \begin{bmatrix} u_{11} & u_{12} \\ u_{21} & u_{22} \end{bmatrix}. \tag{A-2}$$

We can rewrite the first row of (A-2) as

$$\lambda_1 \underline{s}_1 = \lambda_1 \underline{s}_1 \phi_{11} + u_{11}, \quad 0 = \lambda_1 \underline{s}_1 \phi_{21} + u_{12}.$$

Since, according to (A-1) the u_t's are normally distributed we can interpret the relationships as

$$\phi_{11} \sim \mathcal{N}(1, \Sigma_{11}/(\lambda_1^2 \underline{s}_1^2)), \quad \phi_{21} \sim \mathcal{N}(0, \Sigma_{22}/(\lambda_1^2 \underline{s}_1^2)).$$

ϕ_{ij} denotes the element i,j of the matrix Φ, and Σ_{ij} corresponds to element i,j of Σ. The hyperparameter λ_1 controls the tightness of the prior.

Dummy Observations for Φ_2.

$$\begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 0 & \lambda_1 \underline{s}_1 2^{\lambda_2} & 0 & 0 \\ 0 & 0 & 0 & \lambda_1 \underline{s}_2 2^{\lambda_2} & 0 \end{bmatrix} \Phi + U, \tag{A-3}$$

where the hyperparameter λ_2 is used to scale the prior standard deviations for coefficients associated with x_{t-l} according to $l^{-\lambda_2}$.

Dummy Observations for Σ. A prior for the covariance matrix Σ, centered at a matrix that is diagonal with elements equal to the presample variance of x_t, is obtained by stacking the observations

$$\begin{bmatrix} \underline{s}_1 & 0 \\ 0 & \underline{s}_2 \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \Phi + U \tag{A-4}$$

λ_3 times.

Sums-of-coefficients Dummy Observations. When lagged values of a variable $x_{i,t}$ are at the level $\underline{x}_i$, the same value $\underline{x}_i$ is *a priori* likely to be a good forecast of $x_{i,t}$, regardless of the value of other variables:

$$\begin{bmatrix} \lambda_4 \underline{x}_1 & 0 \\ 0 & \lambda_4 \underline{x}_2 \end{bmatrix} = \begin{bmatrix} \lambda_4 \underline{x}_1 & 0 & \lambda_4 \underline{x}_1 & 0 & 0 \\ 0 & \lambda_4 \underline{x}_2 & 0 & \lambda_4 \underline{x}_2 & 0 \end{bmatrix} \Phi + U. \tag{A-5}$$

Co-persistence Dummy Observations. When all lagged x_t's are at the level $\underline{x}$, *a priori* x_t tends to persist at that level:

$$\begin{bmatrix} \lambda_5 \underline{x}_1 & \lambda_5 \underline{x}_2 \end{bmatrix} = \begin{bmatrix} \lambda_5 \underline{x}_1 & \lambda_5 \underline{x}_2 & \lambda_5 \underline{x}_1 & \lambda_5 \underline{x}_2 & \lambda_5 \end{bmatrix} \Phi + U. \tag{A-6}$$

Prior Distribution. After collecting the T^* dummy observations in matrices X^* and W^*, the likelihood function associated with (A-1) can be used to relate the dummy observations to the parameters Φ and Σ. If we combine the likelihood function with the improper prior $p(\Phi, \Sigma) \propto |\Sigma|^{-(n+1)/2}$, we can deduce that the product $p(X^*|\Phi, \Sigma) \cdot |\Sigma|^{-(n+1)/2}$ can be interpreted as

$$(\Phi, \Sigma) \sim MNIW(\underline{\Phi}, (W^{*'} W^*)^{-1}, \underline{S}, T^* - k), \tag{A-7}$$

where $\underline{\Phi}$ and $\underline{S}$ are

$$\underline{\Phi} = (W^{*\prime} W^*)^{-1} W^{*\prime} W^*, \quad \underline{S} = (X^* - W^* \underline{\Phi})'(X^* - W^* \underline{\Phi}).$$

Provided that $T^* > k + n$ and $W^{*'} W^*$ is invertible, the prior distribution is proper.

Hyperparameter Choices. We use the same set of dummy observations for the MF-VAR, the QF-VAR, and the quarterly AR(2) models. The hyperparameters are selected based on the first sample that is used for parameter estimation in the pseudo-out-of-sample forecasting. The hyperparameter choices are summarized in Tables A-1 and A-2.

A.2 Initial Distribution $p(z_0|Y_{-p+1:0})$

Recall that $t = 1$ corresponds to 1968:M1. Let $T_- = -11$ such that $t = T_-$ corresponds to 1967:M1. We then initialize z_{T_-} using actual observations. This is straightforward for x_{m,T_-}, $x_{m,T_- -1}$, $x_{m,T_- -p}$ because they are observed. We set x_{q,T_-}, $x_{q,T_- -1}$, $x_{q,T_- -p}$ equal to the observed quarterly values, assuming that during these periods the monthly-within-quarter values simply equal the observed averages during the quarter. This provides us

Table A-1: Hyperparameters: VAR

	11-MF-VAR	4-MF-VAR	11-QF-VAR	4-QF-VAR
λ_1	0.81	0.29	2.85	0.41
λ_2	2.18	2.18	0.10	0.10
λ_3	1	1	1	1
λ_4	1.66	1.66	0.11	0.11
λ_5	2.14	2.14	10.64	10.64

Notes: Hyperparameters are selected based on the first recursive sample.

Table A-2: Hyperparameters: QF-AR

	UNR	HRS	CPI	IP	PCE	FF	TB	SP500	GDP	INVFIX	GOV
λ_1	0.32	0.39	0.09	0.59	0.57	11.40	0.00	0.00	0.48	0.51	80.23
λ_2	0.00	0.00	0.00	0.00	0.00	0.00	11.35	11.43	0.00	0.00	105.00
λ_3	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
λ_4	105.00	10.56	10.56	10.56	0.00	31.67	105.00	10.56	10.56	10.56	10.56
λ_5	10.56	11.67	105.00	105.00	10.56	0.00	10.56	10.56	10.56	105.00	0.00

Notes: Hyperparameters are selected based on the first recursive sample.

with a distribution for $p(z_{T_-})$ that is simply a point mass. We then set Φ and Σ equal to their respective prior means and apply the Kalman filter for $t = T_- + 1, \dots 0$ to the state space system described in (2) and (7), updating the beliefs about the latent state z_t with pre-sample observations $Y_{T_-:0}$. In slight abuse of notation, we denote the distribution of z_t obtained after the period 0 updating by $p(z_0|Y_{-p+1})$. Note that this distribution does not depend on the "unknown" parameters Φ and Σ, because the Kalman filter iterations were implemented based on the prior means of these matrices.

A.3 Compact State-Space Representation

As discussed in the main text, the computational efficiency of the simulation-smoother step in the Gibbs sampler can be improved by eliminating, for $t = 1, \dots, T_b$, the monthly observations $x_{m,t}$ from the state vector z_t that appears in the measurement equation (7). We begin by re-ordering the lags of x_t and the VAR coefficients in (1) to separate lags of $x_{m,t}$ from lags of $x_{q,t}$. Define the $pn_m \times 1$ vector $z_{m,t}$ and $pn_q \times 1$ vector $z_{q,t}$ as

$$z'_{m,t} = \left[x'_{m,t}, \dots, x'_{m,t-p+1}\right], \quad z'_{q,t} = \left[x'_{q,t}, \dots, x'_{q,t-p+1}\right].$$

In a similar manner, define the $n_m \times pn_m$ matrix Φ_{mm}, the $n_m \times pn_q$ matrix Φ_{mq}, the $n_q \times pn_m$ matrix Φ_{qm}, and the $n_q \times pn_q$ matrix Φ_{qq} such that (1) can be rewritten as

$$\begin{bmatrix} x_{m,t} \\ x_{q,t} \end{bmatrix} = \begin{bmatrix} \Phi_{mm} & \Phi_{mq} \\ \Phi_{qm} & \Phi_{qq} \end{bmatrix} \begin{bmatrix} z_{m,t-1} \\ z_{q,t-1} \end{bmatrix} + \begin{bmatrix} \Phi_{mc} \\ \Phi_{qc} \end{bmatrix} + \begin{bmatrix} u_{m,t} \\ u_{q,t} \end{bmatrix}. \tag{A-8}$$

Recall that for $t \leq T_b$ all the monthly series are observed. Thus, $y_{m,t} = x_{m,t}$ and, in slight abuse of notation, $z_{m,t-1} = y_{m,t-p:t-1}$. Now define $s_t = [x'_{q,t}, z'_{q,t-1}]'$ and notice that based on the second equation in (A-8) one can define matrices Γ_s, Γ_{zm}, Γ_c, and Γ_u such that we obtain a state-transition equation in companion form

$$s_t = \Gamma_s s_{t-1} + \Gamma_{zm} y_{m,t-p:t-1} + \Gamma_c + \Gamma_u u_{q,t}. \tag{A-9}$$

The measurement equation for the monthly series takes the form

$$y_{m,t} = \Lambda_{ms} s_t + \Phi_{mm} y_{m,t-p:t-1} + \Phi_{mc} + u_{m,t}. \tag{A-10}$$

Finally, the measurement equation for the quarterly series can be expressed as

$$y_{q,t} = M_{q,t} \Lambda_{qs} s_t, \tag{A-11}$$

where the matrix $\Lambda_{qs}s_t$ averages $x_{q,t}$, $x_{q,t-1}$, and $x_{q,t-2}$ and $M_{q,t}$ is a time-varying selection matrix that selects the elements of $\Lambda_{qs}s_t$ that are observed in period t. In sum, (A-9), (A-10), and (A-11) provide an alternative state-space representation of the MF-VAR that reduces the dimension of the state vector from np to $n_q(p+1)$. In this alternative representation the "measurement errors" $u_{m,t}$ in (A-10) are correlated with the innovations $u_{q,t}$ in the state-transition equation (A-9). Moreover, the lagged observables $y_{m,t-p:t-1}$ directly enter the state-transition and measurement equations. Since these observables are part of the $t-1$ information the modification of the Kalman filter and simulation smoother is straightforward.

At the end of period $t = T_b$ we switch from the state-space representation in terms of $s_t = [x'_{q,t}, \ldots, x'_{q,t-p}]'$ to a state-space representation in terms of $\tilde{z}_t = [z'_t, x'_{t-p}] = [x'_t, \ldots, x'_{t-p}]'$.[5] In the forward pass of the Kalman filter, let $\hat{s}_{t|t} = \mathbb{E}[s_t|Y_{-p+1:t}]$ and $P^s_{t|t} = \mathbb{V}[s_t|Y_{-p+1:t}]$ (omitting (Φ, Σ) from the conditioning set). Since $x_{m,t}, \ldots, x_{m,t-p+1}$ is known conditional on the $Y_{-p+1:t}$, we can easily obtain $\hat{\tilde{z}}_{t|t} = \mathbb{E}[\tilde{z}_t|Y_{-p+1:t}]$ by augmenting $\hat{s}_{t|t}$ with $y_{m,t}, \ldots, y_{m,t-p}$. Moreover, $P^{\tilde{z}}_{t|t} = \mathbb{V}[\tilde{z}_t|Y_{-p+1:t}]$ can be obtained by augmenting $P^s_{t|t}$ by zeros, to reflect that $x_{m,t}, \ldots, x_{m,t-p}$ are known with certainty. In the backward pass of the simulation smoother we start out with a sequence of draws from $\tilde{z}_T|Y_{-p+1:T}$ and $\tilde{z}_t|(\tilde{Z}_{t+1:T}, Y_{-p+1:T})$ for $t = T-1, \ldots T_b+1$. Let $\hat{\tilde{z}}_{t|T}$ and $P^{\tilde{z}}_{t|T}$ denote the mean and variance associated with this distribution. At $t = T_b$ we convert the conditional mean and variance of $\tilde{z}_{T_b}$ into a conditional mean and variance for s_{T_b}. This is done by eliminating all elements associated with $x_{m,t}, \ldots, x_{m,t-p}$.

B Construction of Real-Time Data Set

The eleven real-time macroeconomic data series are obtained from the ALFRED database maintained by the Federal Reserve Bank of St. Louis. Table B summarizes how the series used in this paper are linked to the series provided by ALFRED.

We construct two sequences of dates that contain the set of forecast origins $(T_{min}, \ldots, T_{max})$. One sequence contains the last day of each month and the other sequence will be comprised of the Greenbook forecast dates. ALFRED provides a publication date for each data vintage. We wrote a computer program that for every forecast origin, selects the most recent

[5]We augment the state vector z_t in (2) and (7) by an additional lag of x_t to ensure that s_t is a subvector of the resulting $\tilde{z}_t$. This augmentation requires a straightforward modification of the state-transition equation (2) and the measurement equations (7).

Table B-1: ALFRED Series Used in Analysis

Time Series	ALFRED Name
Gross Domestic Product (GDP)	GDPC1
Fixed Investment (INVFIX)	FPIC1
Government Expenditures (GOV)	GCEC1
Unemployment Rate (UNR)	UNRATE
Hours Worked (HRS)	AWHI
Consumer Price Index (CPI)	CPIAUCSL
Industrial Production Index (IP)	INDPRO
Personal Consumption Expenditure (PCE)	PCEC96
Federal Fund Rate (FF)	FEDFUNDS
Treasury Bond Yield (TB)	GS10
SP 500 (SP500)	SP500

ALFRED vintage for each of the eleven variables and combines the series into a single data set. This leaves us with a real-time data set for each forecast origin. Based on the missing values in each real-time data set, we construct the selection matrices M_t, $t = T_b + 1, \dots, T$, that appear in (7). The patterns of missing values are summarized in Tables 1 and C-1. Greenbook forecasts are also obtained from the ALFRED database.

Some of the vintages of PCE and INVFIX extracted from ALFRED were incomplete. The recent vintages of PCE and INVFIX from ALFRED do not include data prior to 1990 or 1995 (depending on the vintages). However, the most recent data for PCE and INVFIX can be obtained from BEA or NIPA, say from 1/1/1967 to 7/1/2011. Let us consider PCE for illustration. For the vintages between 12/10/2003 and 6/25/2009, data start from 1/1/1990 and for the vintages between 7/31/2009 and present, data start from 1/1/1995. First, we compute the growth rates from the most recent data. Based on the computed growth rates, we can backcast historical series up to 1/1/1967 using the 1/1/1990 (1/1/1995) data points as initializations. We think this is a reasonable way to construct the missing points. We eliminated 4 of the 141 samples (28, 29, 33, 96) because the vintages for PCE and INVFIX were incomplete. In principle, we could backcast as for the other vintages but we took a short cut.

C Additional Tables and Figures

Table C-1 lists exceptions for the classification of information sets for specific forecast origins.

Table C-2 provides numerical values for the RMSEs attained by the eleven-variable MF-VAR.

Figure C-1 displays PITs for the eleven-variable QF-VAR.

We also consider a four-variable MF-VAR based on one quarterly series and three monthly series. The three monthly series are: Consumer Price Index (CPI), Unemployment Rate (UNR), Federal Fund Rate (FF). The quarterly series is Real GDP. Real GDP and CPI enter the MF-VAR in log levels, whereas UNR and FF are simply divided by 100 to make their scale comparable to the scale of the two other variables. As for the eleven-variable VAR, the number of lags is set to six.

Figure C-2 reports RMSE ratios for the four-variable MF-VAR versus a four-variable QF-VAR.

Figure C-3 reports RMSE ratios for the four-variable MF-VAR versus univariate QF-AR(2) models.

Figure C-4 to C-6 report PIT histograms for the four-variable MF-VAR, QF-VAR, and for univariate QF-AR(2) models.

Figures C-7 to C-10 report interval forecast and actual values for GDP growth, inflation, the unemployment rate, and the Federal Funds rate from the eleven-variable MF-VAR.

Table C-1: Illustration of Information Sets: Exceptions

Exceptions E_0: January (+0 Months)

		UNR	HRS	CPI	IP	PCE	FF	TB	SP500	GDP	INVFIX	GOV
Q4	M10	X	X	X	X	X	X	X	X	QAv	QAv	QAv
Q4	M11	X	X	X	X	X	X	X	X	QAv	QAv	QAv
Q4	M12	X	X	X	X	∅	X	X	X	QAv	QAv	QAv
Q1	M1	∅	∅	∅	∅	∅	X	X	X	∅	∅	∅

Exceptions E_1: February (+1 Month)

		UNR	HRS	CPI	IP	PCE	FF	TB	SP500	GDP	INVFIX	GOV
Q4	M11	X	X	X	X	X	X	X	X	QAv	QAv	QAv
Q4	M12	X	X	X	X	X	X	X	X	QAv	QAv	QAv
Q1	M1	X	X	X	X	∅	X	X	X	∅	∅	∅
Q1	M2	∅	∅	∅	∅	∅	X	X	X	∅	∅	∅

Exceptions E_2: March (+2 Months)

		UNR	HRS	CPI	IP	PCE	FF	TB	SP500	GDP	INVFIX	GOV
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	∅	∅	∅
Q1	M3	∅	∅	∅	∅	∅	X	X	X	∅	∅	∅

Notes: ∅ indicates that the variable 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 Month" group: March, June, September, December. The table illustrates exceptions that arise due to an occasional two-month publication lag for PCE. Exception E_0 occurs for 26 out of 140 recursive samples (1, 4, 7, 10, 13, 16, 19, 22, 28, 37, 43, 52, 61, 64, 73, 79, 85, 88, 96, 105, 108, 114, 123, 129, 132, 138). Exception E_1 occurs for 13 out of 140 recursive samples (8, 20, 44, 53, 56, 68, 89, 97, 100, 103, 115, 118, 139). Exception E_2 occurs for 5 out of 140 recursive samples (21, 27, 48, 51, 78).

Table C-2: RMSEs for 11-Variable MF-VAR

Horizon	UNR	HRS	CPI	IP	PCE	FF	TB	SP500	GDP	INVFIX	GOV
+0 Months											
1	0.21	0.48	0.59	0.97	0.52	0.22	0.17	3.12	0.61	1.69	0.80
2	0.45	0.73	0.64	1.37	0.69	0.72	0.43	8.19	0.77	2.36	0.76
3	0.75	0.93	0.68	1.68	0.75	1.09	0.60	8.34	0.86	2.78	0.75
4	1.07	1.00	0.66	1.73	0.75	1.41	0.68	8.31	0.89	2.86	0.75
5	1.35	0.98	0.66	1.68	0.72	1.69	0.78	8.12	0.86	2.81	0.70
6	1.60	0.92	0.64	1.61	0.70	1.94	0.86	8.13	0.83	2.70	0.67
7	1.81	0.86	0.64	1.57	0.66	2.14	0.87	8.26	0.79	2.61	0.67
8	1.99	0.85	0.63	1.57	0.67	2.27	0.88	8.19	0.79	2.64	0.71
+1 Month											
1	0.13	0.37	0.34	0.97	0.44	0.08	0.08	1.28	0.57	1.55	0.81
2	0.39	0.70	0.66	1.36	0.68	0.60	0.34	8.22	0.79	2.25	0.77
3	0.68	0.91	0.68	1.70	0.76	0.97	0.55	8.32	0.86	2.77	0.76
4	1.01	1.00	0.67	1.74	0.76	1.29	0.63	8.34	0.90	2.88	0.75
5	1.31	0.99	0.66	1.72	0.73	1.59	0.71	8.08	0.88	2.84	0.69
6	1.56	0.94	0.64	1.65	0.70	1.85	0.78	8.11	0.83	2.76	0.68
7	1.78	0.89	0.64	1.61	0.67	2.06	0.80	8.22	0.82	2.65	0.67
8	1.97	0.86	0.63	1.59	0.67	2.21	0.83	8.18	0.79	2.67	0.71
+2 Months											
1	0.08	0.32	0.19	0.75	0.38	0.00	0.00	0.00	0.52	1.44	0.80
2	0.29	0.56	0.62	1.09	0.66	0.44	0.39	7.29	0.72	1.96	0.81
3	0.56	0.84	0.67	1.62	0.75	0.83	0.62	8.49	0.85	2.68	0.76
4	0.89	0.98	0.68	1.74	0.77	1.14	0.72	8.27	0.90	2.88	0.77
5	1.21	1.00	0.66	1.71	0.74	1.46	0.78	8.21	0.87	2.86	0.70
6	1.48	0.95	0.67	1.65	0.71	1.72	0.87	8.27	0.84	2.77	0.67
7	1.71	0.91	0.63	1.60	0.69	1.94	0.85	8.12	0.82	2.69	0.70
8	1.91	0.87	0.63	1.62	0.67	2.07	0.83	8.39	0.81	2.65	0.69
All Forecasts											
1	0.15	0.39	0.41	0.91	0.45	0.14	0.11	1.95	0.57	1.56	0.80
2	0.38	0.67	0.64	1.28	0.68	0.60	0.39	7.92	0.76	2.20	0.78
3	0.67	0.89	0.68	1.67	0.75	0.97	0.59	8.38	0.85	2.74	0.76
4	0.99	0.99	0.67	1.74	0.76	1.29	0.68	8.31	0.90	2.88	0.76
5	1.29	0.99	0.66	1.70	0.73	1.58	0.76	8.14	0.87	2.84	0.70
6	1.55	0.94	0.65	1.64	0.70	1.84	0.84	8.17	0.83	2.74	0.67
7	1.77	0.89	0.64	1.59	0.68	2.05	0.84	8.20	0.81	2.65	0.68
8	1.95	0.86	0.63	1.59	0.67	2.19	0.85	8.25	0.80	2.65	0.70

Notes: RMSEs for UNR (%), FF (annualized %), and TB (annualized %) refer to forecasts of levels. The remaining RMSEs refer to forecasts of quarter-on-quarter growth rates in percentages.

Figure C-1: PIT Histograms for 11-Variable QF-VAR

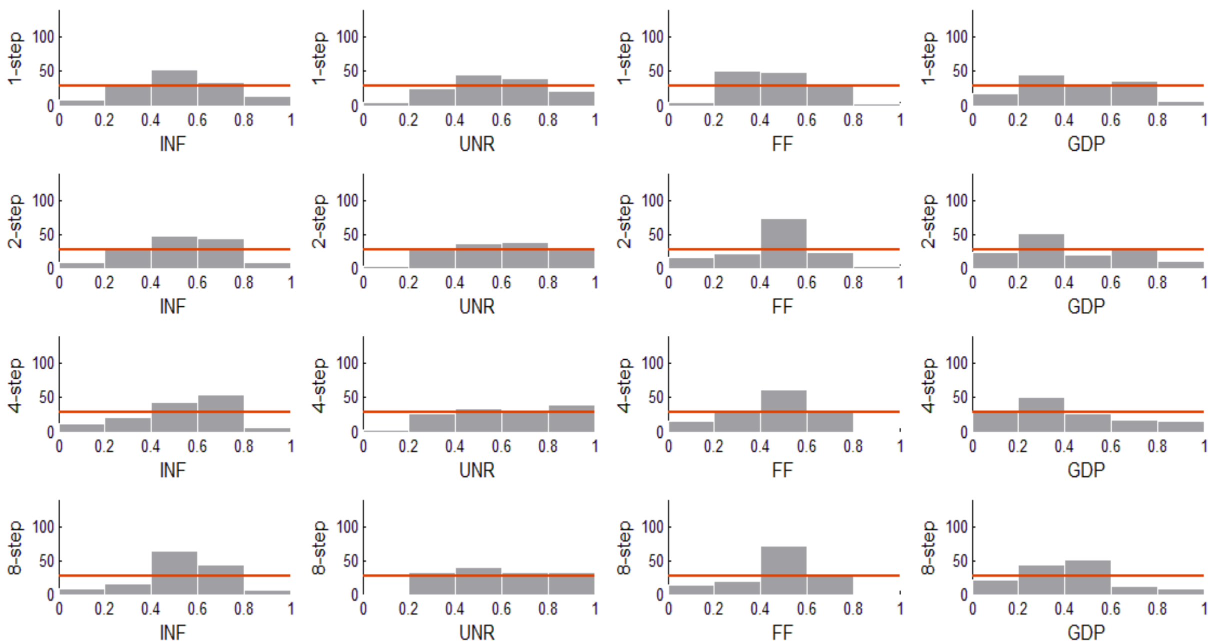

Notes: Probability integral transforms for forecasts of inflation (INF), unemployment rate (UNR), federal fund rate (FF), and GDP growth (GDP). The bars represent the frequency of PITs falling in each bin. The solid line marks 20 percent.

Figure C-2: Relative RMSEs of 4-Variable MF-VAR versus QF-VAR

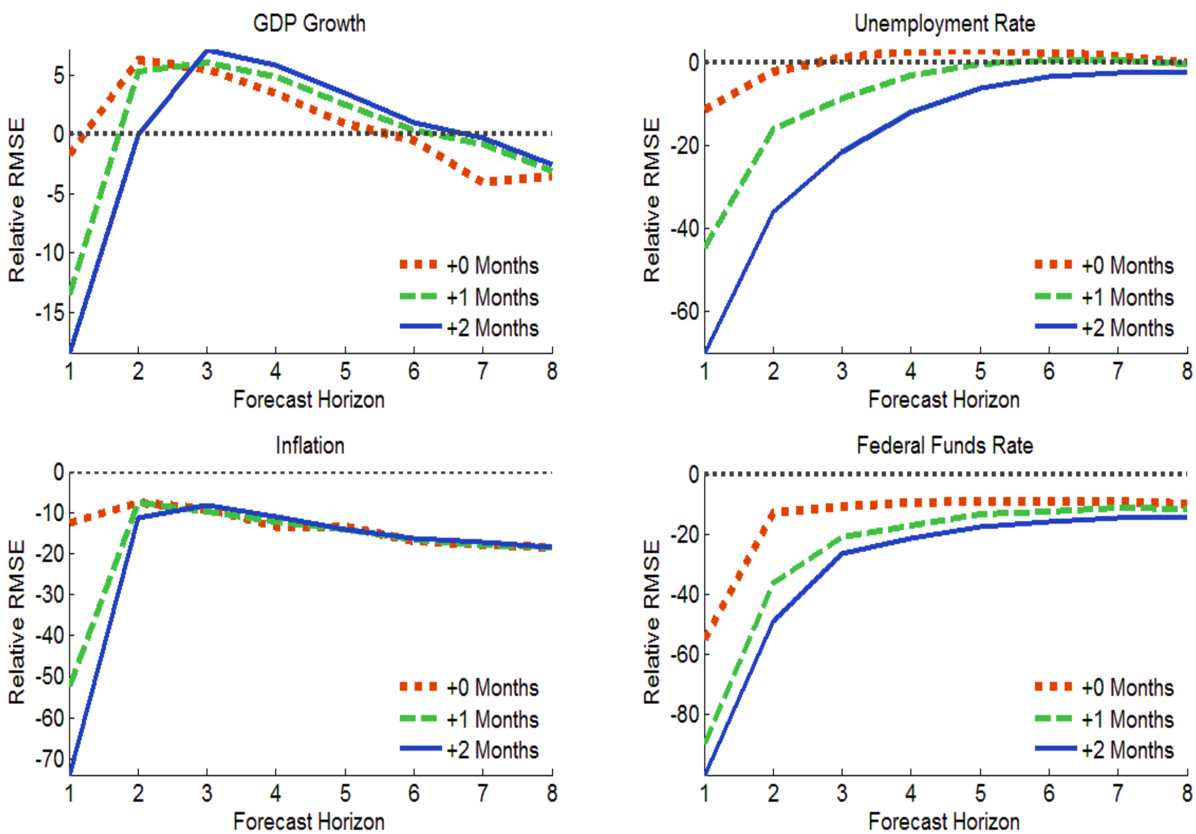

Figure C-3: Relative RMSEs of 4-Variable MF-VAR versus QF-AR

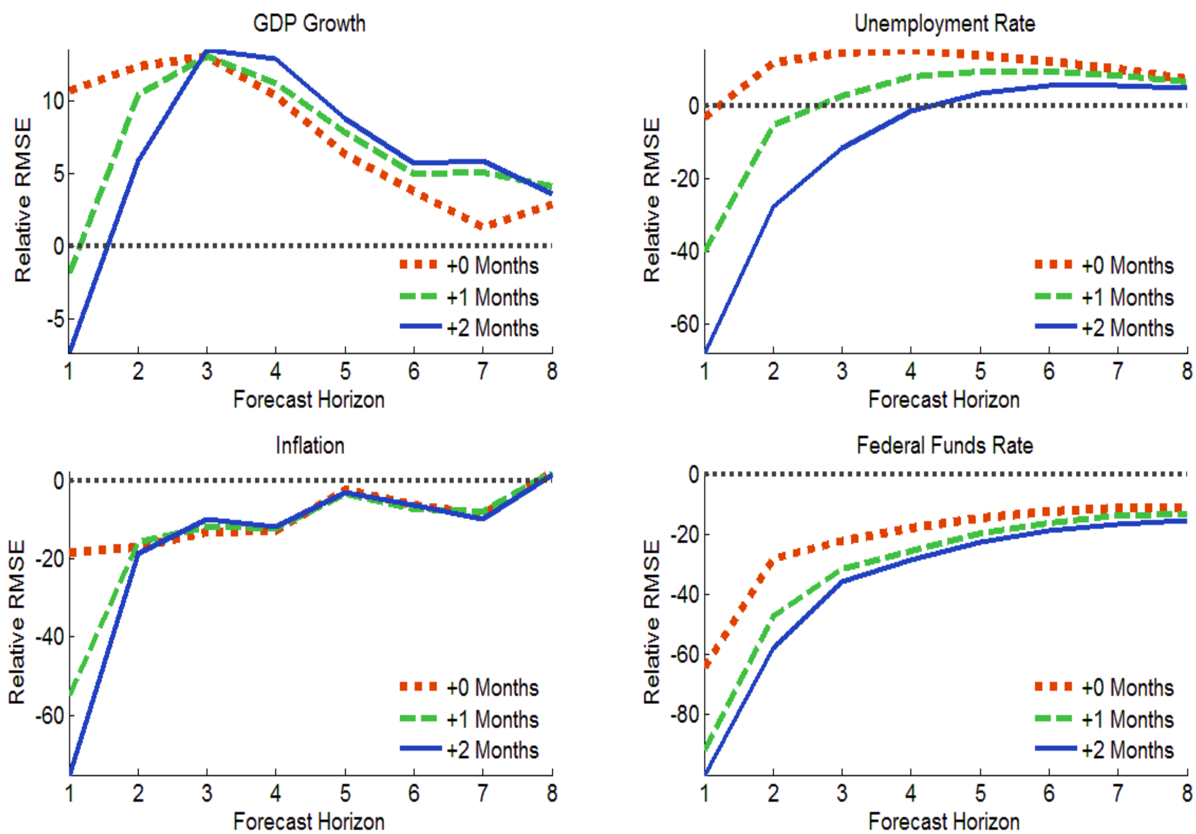

Figure C-4: PIT Histograms for 4-Variable MF-VAR

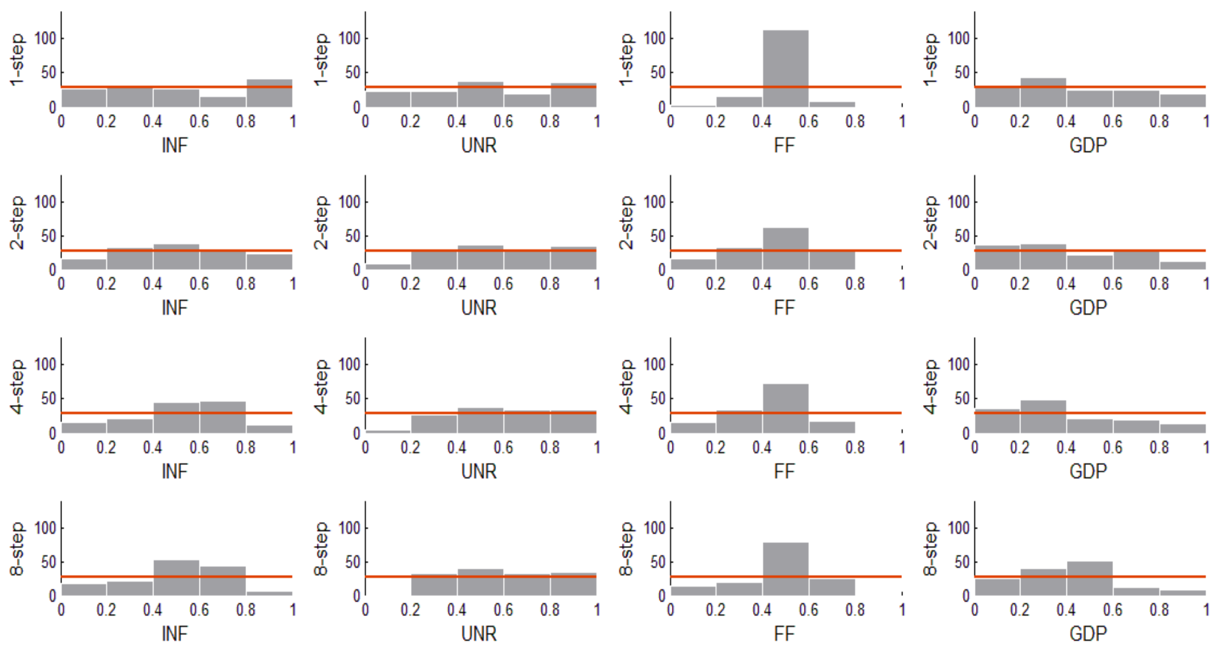

Notes: Probability integral transforms for forecasts of inflation (INF), unemployment rate (UNR), federal fund rate (FF), and GDP growth (GDP) using the MF-VAR. The bars show how many of the realized observations fall into each bin. If the density forecast is accurate, than the bars should be equally distributed across the bins. Solid line is the 20 percent line.

Figure C-5: PIT Histograms for 4-Variable QF-VAR

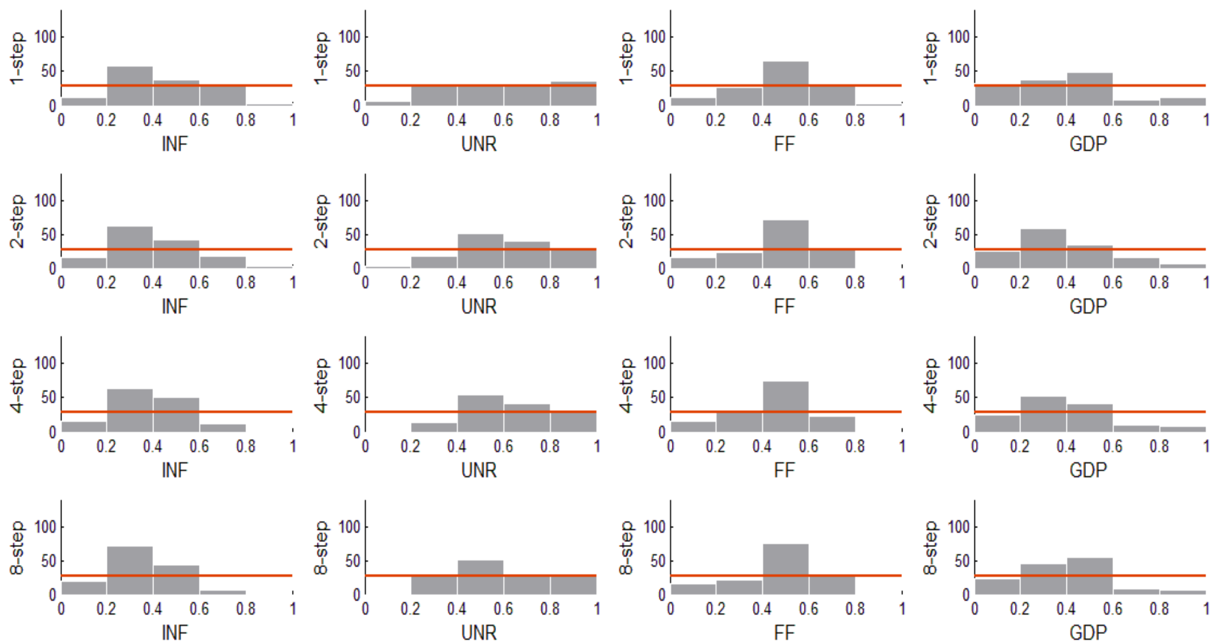

Notes: Probability integral transforms for forecasts of inflation (INF), unemployment rate (UNR), federal fund rate (FF), and GDP growth (GDP). The bars represent the frequency of PITs falling in each bin. The solid line marks 20 percent.

Figure C-6: PIT Histograms for QF-AR

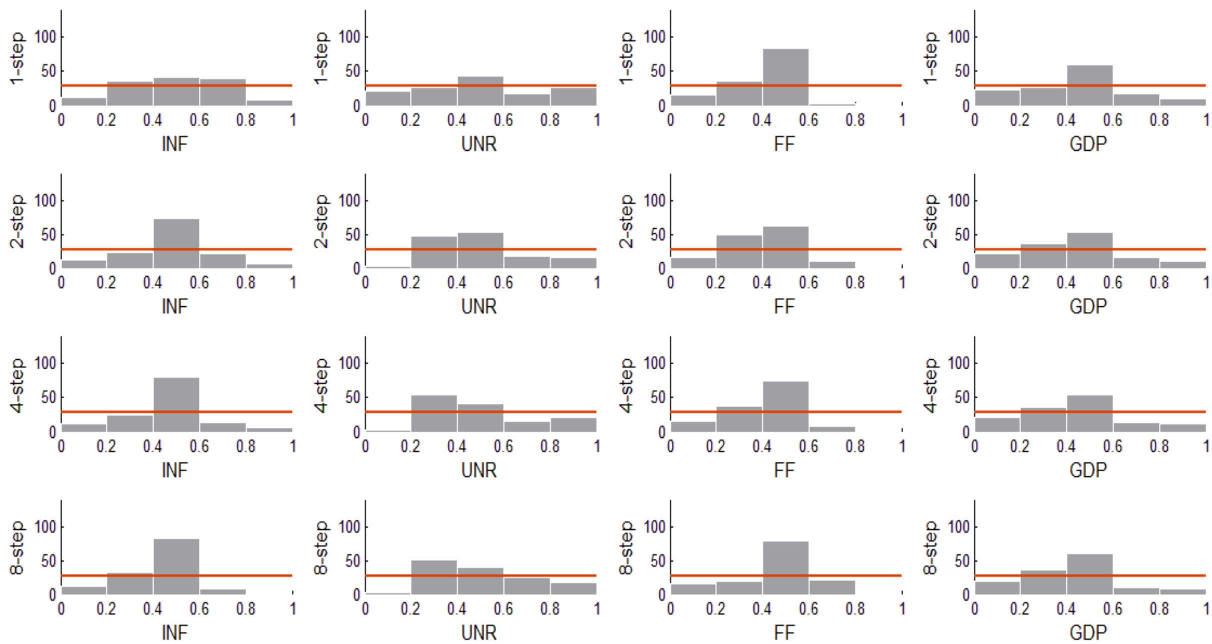

Notes: Probability integral transforms for forecasts of inflation (INF), unemployment rate (UNR), federal fund rate (FF), and GDP growth (GDP). The bars represent the frequency of PITs falling in each bin. The solid line marks 20 percent.

Figure C-7: GDP Growth Forecasts of 11-Variable MF-VAR

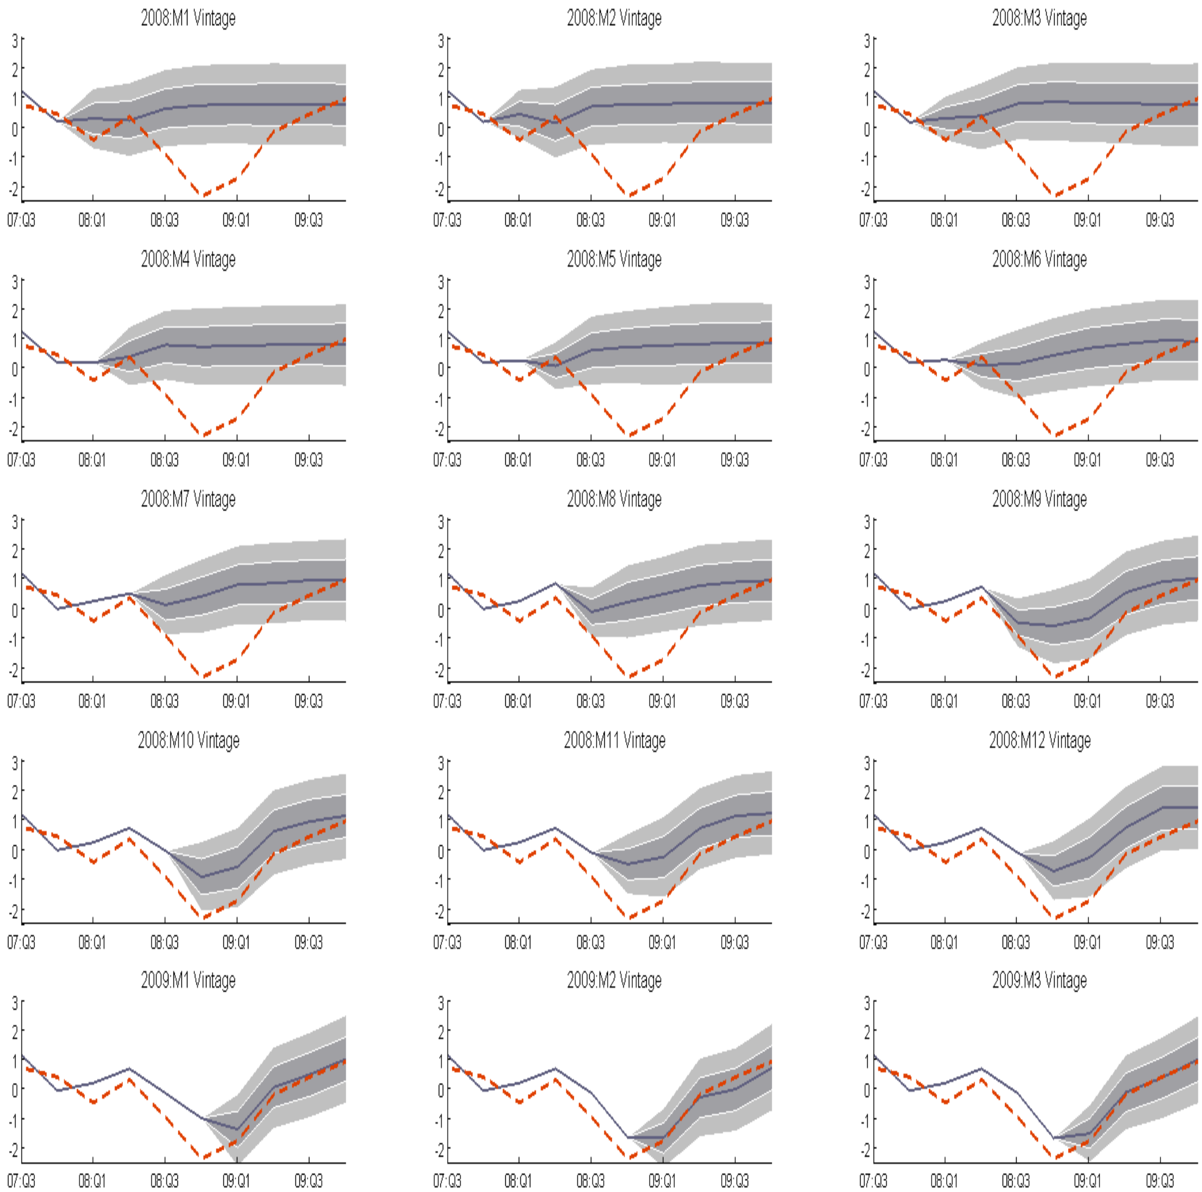

Notes: Actual values are from the $T_* = 2011 : M7$ data vintage and are denoted as the red dashed line. Starting from the leftmost column, we show the results of "+0 Months", "+1 Months", and "+2 Months" subgroups. The title in each subplot indicates the data vintage that are used in the estimation.

Figure C-8: Inflation Forecasts of 11-Variable MF-VAR

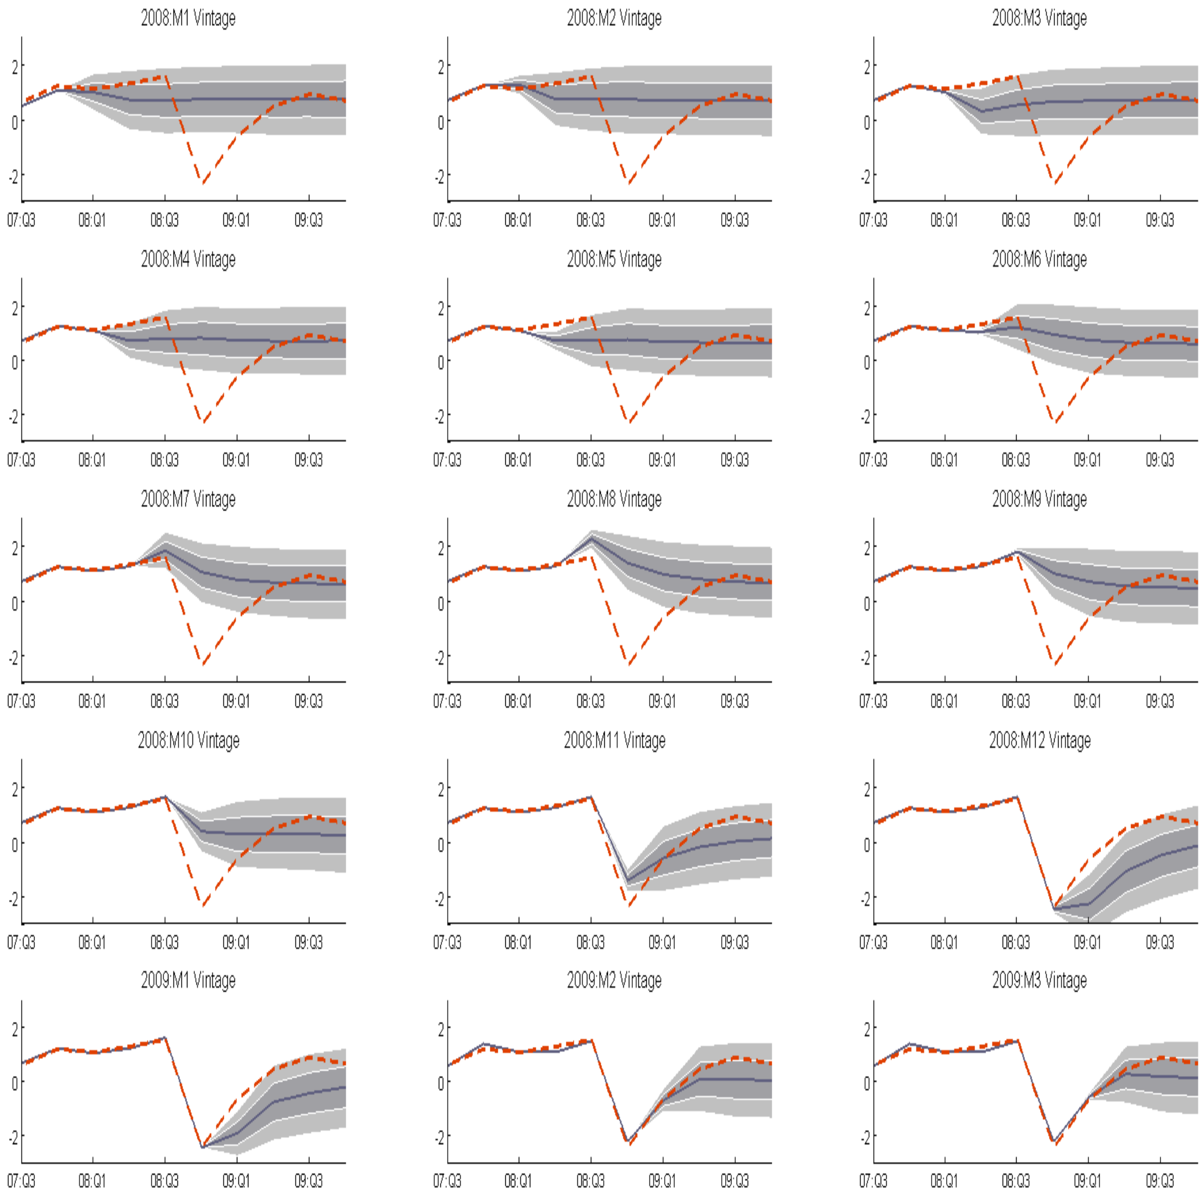

Notes: Actual values are from the $T_* = 2011 : M7$ data vintage and are denoted as the red dashed line. Starting from the leftmost column, we show the results of "+0 Months", "+1 Months", and "+2 Months" subgroups. The title in each subplot indicates the data vintage that are used in the estimation.

Figure C-9: Unemployment Forecasts of 11-Variable MF-VAR

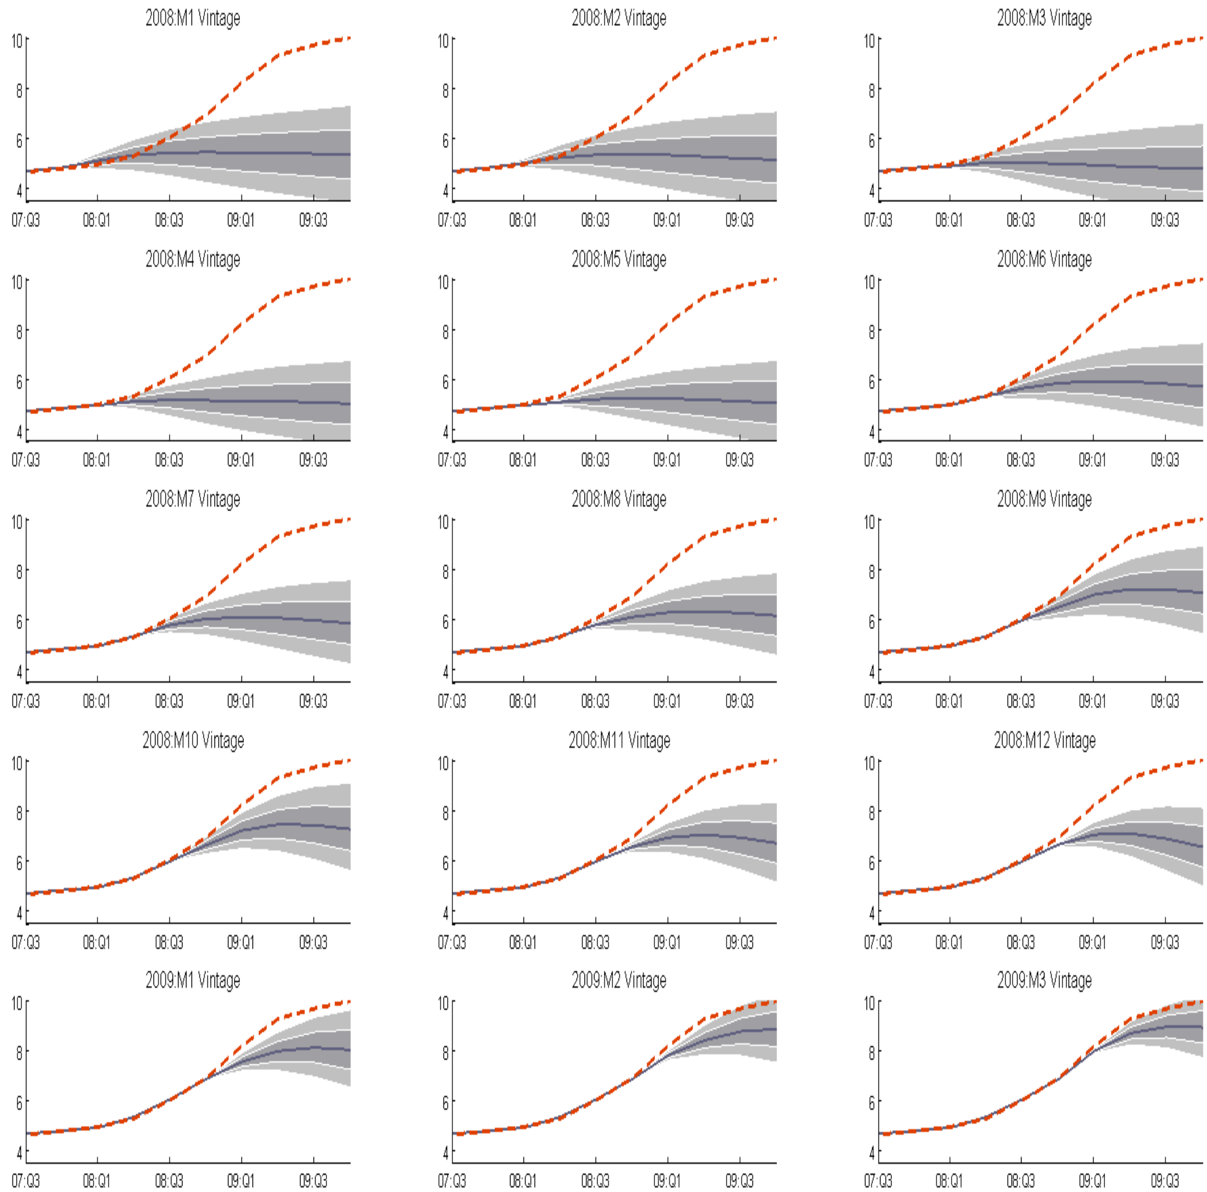

Notes: Actual values are from the $T_* = 2011 : M7$ data vintage and are denoted as the red dashed line. Starting from the leftmost column, we show the results of "+0 Months", "+1 Months", and "+2 Months" subgroups. The title in each subplot indicates the data vintage that are used in the estimation.

Figure C-10: Federal Funds Rate Forecasts of 11-Variable MF-VAR

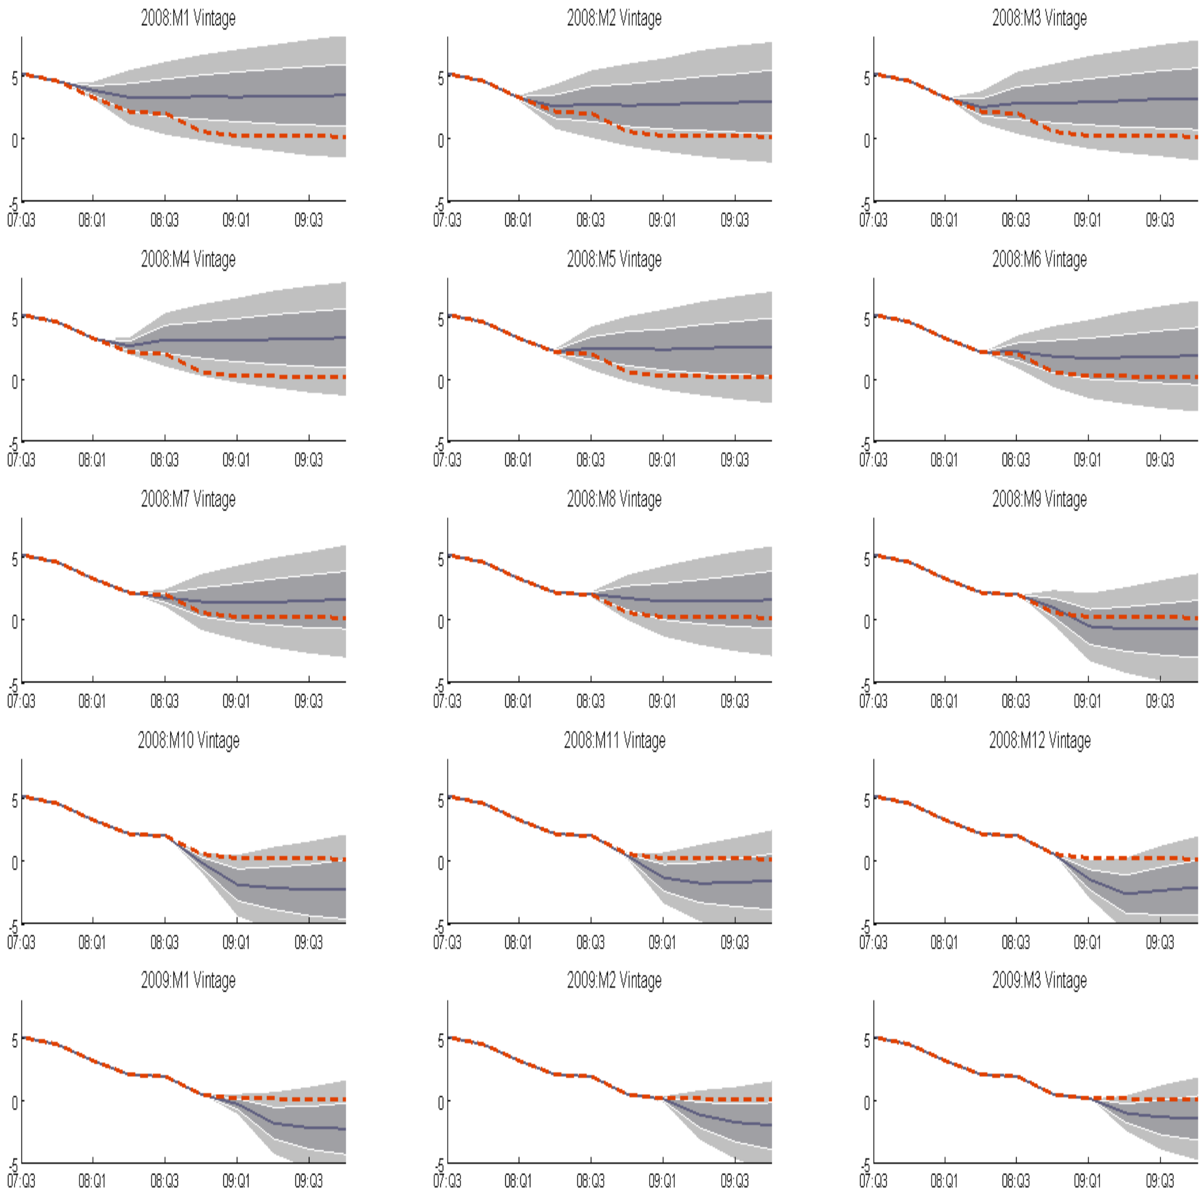

Notes: Actual values are from the $T_* = 2011 : M7$ data vintage and are denoted as the red dashed line. Starting from the leftmost column, we show the results of "+0 Months", "+1 Months", and "+2 Months" subgroups. The title in each subplot indicates the data vintage that are used in the estimation.