Real-Time Forecasting and Scenario Analysis Using a Large Mixed-Frequency Bayesian VAR*

Michael W. McCracken,[a] Michael T. Owyang,[a] and Tatevik Sekhposyan[b]
[a]Federal Reserve Bank of St. Louis
[b]Texas A&M University

We use a mixed-frequency vector autoregression to obtain intraquarter point and density forecasts as new, high-frequency information becomes available. This model, delineated in Ghysels (2016), is specified at the lowest sampling frequency; high-frequency observations are treated as different economic series occurring at the low frequency. As this type of data stacking results in a high-dimensional system, we rely on Bayesian shrinkage to mitigate parameter proliferation. We obtain high-frequency updates to forecasts by treating new data releases as conditioning information. The same framework is used for scenario analysis to obtain forecasts conditional on a hypothetical future path of the variables in the system. We show that the

*We have benefited from conversations with Bob Rasche. We are grateful to Domenico Giannone, Eric Ghysels, and participants of seminars at the Bank of Canada, Federal Reserve Board, Ohio State University, University of Houston, 2013 CIRANO-CIREQ Workshop on Data Revisions in Macroeconomic Forecasting and Policy, 2013 $(EC)^2$ Conference, 2014 SNDE Conference, 2014 Midwest Macro Meeting, 2014 EABCN Conference, 2014 IAAE Meeting, 2015 St. Louis Fed Econometrics Workshop, 2015 Vienna Workshop on High-Dimensional Time Series in Macroeconomics and Finance, 2016 Texas Camp Econometrics, 2016 BVAR Workshop at the ECB, 2016 Econometric Society European Meetings, 2018 Kick off Workshop on Big Data and Macroeconomic Forecasting at JRC for comments and suggestions. Joe McGillicudy, Hannah G. Shell, and Aaron Amburgey provided research assistance. Part of this research was carried out while Tatevik Sekhposyan was a Visiting Fellow at the Federal Reserve Bank of San Francisco, whose hospitality is greatly acknowledged. The views expressed herein do not reflect the official positions of the Federal Reserve Banks of San Francisco and St. Louis or the Federal Reserve System. Author e-mails: Michael.W.McCracken@stls.frb.org; Michael.T.Owyang@stls.frb.org; tatevik.sekhposyan@gmail.com.

methodology results in competitive point and density forecasts and illustrate the usefulness of the methodology by providing forecasts of real GDP growth given hypothetical paths of a central bank policy rate.

JEL Codes: C22, C52, C53.

1. Introduction

Economic forecasting typically requires managing mixed-frequency data. Across any particular quarter, policymakers and professional forecasters analyze monthly, weekly, and sometimes even daily indicators of economic activity. For instance, the Federal Reserve Bank of Atlanta maintains the GDPNow series on its website, while the Federal Reserve Bank of New York regularly publishes a Nowcasting Report. Both of these forecast series use high-frequency data releases to provide an updated view on the performance of the economy—often viewed through the lens of real GDP growth, which is released at a lower, quarterly frequency.

A variety of statistical methods enable the integration of high-frequency variables into forecasting models that predict lower-frequency variables such as real GDP growth. Corrado and Green (1988) and Parigi and Schlitzer (1995) use linear bridge equations to map monthly data into quarterly frequency when modeling GDP for the United States and Italy, respectively. Giannone, Reichlin, and Small (2008) and Blasques et al. (2016) use dynamic factor models to generate forecasts of U.S. GDP using monthly, weekly, and daily data releases. Zadrozny (1990) and Mittnik and Zadrozny (2005) employ a Kalman-filtering-based maximum-likelihood estimation method to modeling mixed-frequency data in a single vector autoregression (VAR) by treating the low-frequency release as a missing-value problem: GDP exists at a monthly frequency, but we only observe it once every three months. Eraker et al. (2015), Schorfheide and Song (2015), and Brave, Butters, and Justiniano (2019) also treat the low-frequency release as a missing-value problem but use Bayesian methods. Foroni and Marcellino (2013) offer a review of these and other mixed-frequency models.

In this paper, we investigate an alternative approach to forecasting delineated in Ghysels (2016), who addresses the forecaster's mixed-frequency problem using a mixed-frequency vector autoregression estimated at the lowest common data frequency. As an example, suppose we want to forecast real GDP growth using the three monthly nonfarm payroll employment data releases that occur during the quarter. In this case, the VAR would be based on a four-dimensional vector formed by the three monthly and one quarterly series. Bacchiocchi et al. (2016) and Ghysels, Hill, and Motegi (2016) have recently used such a model to investigate the role mixed frequencies play in topics ranging from tests of Granger causality to the construction of impulse response functions, respectively. In other fields—particularly engineering—theoretical aspects of this modeling approach (sometimes referred to as "blocking") have been explored (see Bittanti, Colaneri, and De Nicolao 1988 and, more recently, Chen et al. 2012). As noted in Ghysels (2016), the model can also be interpreted as a multivariate version of the univariate, unrestricted mixed-frequency data-sampling (MIDAS) model discussed in Foroni, Marcellino, and Schumacher (2015).

By stacking the series this way, one can use a conventional, conditional forecasting framework (e.g., Waggoner and Zha 1999) to obtain nowcasts and forecasts consistent with the high-frequency information flow throughout the quarter without computationally intensive filtering (see Ghysels 2016 and Brave, Butters, and Justiniano 2019 for related discussions). Nowcasting, forecasting, and scenario analysis can be treated in the same framework using well-established Bayesian methods in a way that should be particularly appealing for both policymakers and other practitioners.

In this paper, we provide empirical evidence on the ability of this mixed-frequency VAR to forecast in real time as we move within a quarter and additional higher-frequency data become available. Our analysis focuses on the nowcasting problem (i.e., forecasting the current quarter), although we also provide results for a longer, four-quarter-ahead horizon. In our application, we use real-time vintage data on 12 monthly frequency predictors and quarterly GDP, yielding a heavily parameterized, 37-dimensional VAR.[1] To handle the high dimensionality, we estimate the model using Bayesian

[1] Prior to 1992, we forecast real GNP growth. Throughout the remainder we simply reference our target variable as GDP without also referencing GNP.

techniques, where we allow data-driven shrinkage to resolve the bias-variance tradeoff in the forecasting problem. Specifically, we use estimation procedures for reduced-form VARs developed by Giannone, Lenza, and Primiceri (2015; henceforth GLP), to obtain the posterior distribution of the parameters and the predictive densities.

To understand the nature of the model, it is instructive to compare it with the one posited by Carriero, Clark, and Marcellino (2015; henceforth CCM). In their model, CCM construct nowcasts of real GDP growth based on monthly indicators using the unrestricted MIDAS approach delineated in Foroni, Marcellino, and Schumacher (2015). Among other experiments, they use four separate linear scalar models to nowcast current-quarter real GDP growth from information available as of the first week in each month—i.e., after the release of the Employment Situation Report. Their direct multistep (DMS) approach to GDP forecasting has the advantage of not having to form a complete model for all variables in the system as one would with a VAR. Bhansali (1997) and Schorfheide (2005) provide theoretical results showing that DMS approaches to forecasting can be more robust to model misspecification than iterated multistep (IMS) forecasts generated by a fully specified VAR. On the other hand, Marcellino, Stock, and Watson (2006) show that IMS-based models often lead to more accurate predictions relative to DMS-based models. The model we use can be viewed as the VAR-based IMS analogue to the DMS system of equations that CCM use to forecast. While working with a fully specified VAR marginally complicates the empirics, the benefit is a more general framework that can be used for multivariate forecasting and scenario-based conditional forecasting of the type often used by central banks.[2]

Given similarities between our model and CCM's, we compare our nowcasting approach to theirs, noting some advantages to our approach. We use the same monthly predictors and data release schedule; consider the same intraquarter forecast origins (as well

[2]In related work, Kuzin, Marcellino, and Schumacher (2011) compare a MIDAS technique to a mixed-frequency VAR for euro-area data, where the latter is estimated using a frequentist state-space model. Our innovation is that we use a stacked VAR and estimate with Bayesian methods, which arguably provides a more flexible setup. Most notably, we use the conditional forecasting framework jointly for both nowcasting/forecasting and scenario analysis.

as others) and set of models (small and large) that are comparable with their specification. We then directly compare our IMS approach to current-quarter point and density nowcasts of real GDP growth to CCM's DMS approach. We also compare these nowcasting results to forecasts from the Survey of Professional Forecasters (SPF), forecasts from the Blue Chip Economic Indicators (BCEI) survey, and predictions from a quarterly frequency AR(2). In addition, we use the mixed-frequency Bayesian VAR (MF-BVAR) to produce point and density forecasts of GDP growth and several of our high-frequency variables at a longer, four-quarter-ahead horizon and compare the efficacy of these forecasts to those from the SPF and the BCEI.

We find that the MF-BVAR model provides competitive nowcasts and four-quarter-ahead forecasts. The nowcasts, in general, are comparable to the ones obtained using our version of the model proposed in CCM and can be as accurate as those from the SPF and BCEI at certain forecast origins. The four-quarter-ahead forecasts, on the other hand, are as good as or better than those in the SPF for variables such as industrial production, housing starts, and Treasury yields. A small-scale version of the MF-BVAR is also competitive for four-quarter-ahead real GDP growth forecasts relative to the SPF. In general, it is often the case that a lower-dimensional MF-BVAR, consisting of fewer stacked monthly series, provides more accurate point forecasts than those from a larger MF-BVAR with many stacked monthly series. Regardless of which MF-BVAR is being used, the intraquarterly arrival of additional information improves the accuracy of both point and density forecasts.[3]

We also provide two examples of conditional, scenario-based forecasting—both of which are designed to highlight the mixed-frequency nature of the model. In each experiment, we emulate a situation in which a central bank is forming forecasts of a low-frequency variable (i.e., real GDP growth), based on a hypothetical path for a high-frequency variable (i.e., a policy rate). In our first experiment, for a fixed hypothetical path of the policy rate, we show

[3]Bańbura, Giannone, and Reichlin (2010), among others, show that the medium-scale VAR performs similar to the large-scale one in terms of point forecasts accuracy, thus our finding is not surprising in the context of the literature.

how the arrival of other high-frequency observables can lead to significant revisions to the forecast as we move across forecast origins within a quarter. In the second experiment, we instead highlight the intraquarterly timing of the hypothetical path for the policy rate given a fixed forecast origin. For this experiment, we find that conditional forecasts of real GDP growth are monotonically higher the later the policy changes are made within a calendar quarter. This experiment could be potentially informative for a central bank considering the timing of policy actions. In each experiment, we compare our results to those from a model that relies solely on low-frequency aggregates, with the goal of emphasizing that such a model lacks the agility to adapt to intraquarterly movements in high-frequency observables.

The remainder of the paper is organized as follows. Section 2 lays out the specification of the model, discusses the data, and delineates the estimation methodology. Section 3 describes the construction of point and density forecasts; section 4 discusses the forecasting results and compares them to various alternatives. We conclude with section 5.

2. The Setup

The model we use is a quarterly reduced-form VAR, where a monthly variable is represented by three quarterly variables, each corresponding to an intraquarter month.[4] We estimate the model using a procedure suggested in GLP using Minnesota priors reconfigured for stationary data. In GLP, the amount of shrinkage is chosen to maximize the marginal data density. In what follows, we describe the data and the real-time properties of the various series included in our VAR specification. We then provide more details on the stacked VAR and outline the GLP estimation procedure.

2.1 Data

Our choice of predictors aims to facilitate a close comparison of our forecasting results with other mixed-frequency models in the

[4]While our application uses monthly/quarterly data, extension to any mixed frequency is straightforward.

literature, particularly the ones in CCM. CCM includes variables that have proven useful for forecasting U.S. real GDP growth (or GNP for earlier portions of our sample) and are followed by markets and policymakers. The complete set of our monthly predictors, their transformations, and mnemonics are as follows: the S&P 500 composite index (log-change, "stprice"), the 3-month Treasury bill rate ("tbill"), the 10-year Treasury bond yield ("tbond"), the Institute for Supply Management manufacturing index ("ISM"), the ISM supplier deliveries index ("supdel"), the ISM new orders index ("orders"), total nonfarm payroll employment (log-change, "emp"), average weekly hours of production and supervisory workers (log-change, "hours"), real retail sales (the nominal series is deflated using the consumer price index; log-change, "RS"), industrial production (log-change, "IP"), housing starts (change, "starts"), and finally initial unemployment claims ("claims").[5] The transformations are chosen to induce stationarity in the series. If no transformation is listed, the variable is used in levels. In addition, growth rates have been annualized.

Because the model treats higher-frequency series as multiple quarterly frequency series, the number of estimated parameters grows faster than in a standard VAR of equal lag order. In order to mitigate parameter proliferation, we consider data only in monthly and quarterly frequencies. Financial variables are summarized at a monthly frequency, constructed as averages of daily observations. For similar reasons, we do not use weekly releases of initial claims; instead, we choose the four-week moving average. The 12 monthly series, along with the quarterly GDP series (log-change), imply that our large mixed-frequency VAR has $12 \times 3 + 1 = 37$ dimensions. To allow more direct comparison to CCM, we also investigate the performance of a small mixed-frequency VAR that only uses five of the monthly variables yielding 16 dimensions.

[5]In 2001, the Census changed details in the construction of retail sales (RETAIL) and started releasing the new version (RSAFS). Therefore, we use RETAIL for all vintages up to 2001:06 and RSAFS for the vintages after. Morover, when RSAFS was released, the historical sample was extended back only to 1992. Therefore, when we use the RSAFS vintages, we splice the pre-1992:01 values from the last vintage of RETAIL in order for RSAFS to have data dating back to 1970.

While not necessary for estimating the model, we choose to organize our data based on the approximate release calendar. Because the longest publication lag in our data set is one month, we have realizations for all the information associated with the previous month by the end of the next month. Within each month, the data are ordered as follows: (i) the monthly averages of the S&P 500 composite index, 3-month Treasury bill rate, and 10-year Treasury bond yield computed on the first day of the following month; (ii) the ISM manufacturing, supplier deliveries, and new orders indexes; (iii) total nonfarm payroll employment and average weekly hours; (iv) real retail sales; (v) industrial production; (vi) housing starts; and (vii) the four-week moving average of initial unemployment claims. GDP growth is observed following initial unemployment claims.

In all of our forecasting exercises, we use real-time monthly vintage data starting in January 1985 and ending in April 2017. The monthly sequence of 388 real-time vintages of our predictors and GDP are gathered from Haver Analytics, the ALFRED database hosted by the Federal Reserve Bank of St. Louis, and the Real-Time Data Set for Macroeconomists hosted by the Federal Reserve Bank of Philadelphia. Each vintage consists of observables dating back to the first quarter of 1970. Accordingly, the target quarters for nowcasting and forecasting span from 1985:Q1 to 2017:Q1.

2.2 Model

The forecasting model is based on a standard reduced-form VAR, estimated at the lowest sampling frequency. We treat multiple releases at the highest frequency as separate observations modeled in a blocked linear form. For illustrative purposes, consider the case of one quarterly variable (e.g., real GDP) and one monthly variable (e.g., payroll employment) released in each of the three intraquarter months. Let $x_{t-\tau}$ represent the high-frequency (monthly) variable and y_t denote the low-frequency (quarterly) variable, for quarters $t = 1, \ldots T$ and months within the quarter $\tau = \{0, 1/3, 2/3\}$. In this setup, each τ represents an intraquarter month, hence $x_{t-2/3}$, $x_{t-1/3}$, and x_t index data that are observed during the second and third calendar months of quarter $t-1$ and first calendar month of quarter t, respectively.

Define the vector of data releases as $Y_t = [x_{t-2/3}, x_{t-1/3}, x_t, y_t]'$. The reduced-form VAR is

$$Y_t = C + B_1 Y_{t-1} + \cdots + B_p Y_{t-p} + \Sigma^{1/2}\varepsilon_t, \tag{1}$$

where B_l are $n \times n$ parameter matrices, p is the lag order of the VAR, C is an $n \times 1$ vector of intercepts, $\varepsilon_t \sim N(\mathbf{0},I)$, and Σ is the variance of the reduced-form shocks $\Sigma^{1/2}\varepsilon_t$. Let $X_t = [Y'_{t-1}, \ldots, Y'_{t-p}, 1]'$. We can then write (1) as

$$Y_t = DX_t + \Sigma^{1/2}\varepsilon_t, \tag{2}$$

where $D = [B\ C]$ and $B = [B_1 \ldots B_p]$. It will further be useful to define $\beta = vec\,(D)$. In our empirical section, we use a lag order of $p = 1$ as selected recursively by BIC as well as a desire to align the lag structure with that in CCM.

In general, the model is flexible enough to be generalized to include Q quarterly variables and M monthly variables. In such an environment one would treat y_t as a $(Q \times 1)$ vector and $x_{t-\tau}$ as a $(M \times 1)$ vector, producing a VAR of dimension $n = Q + 3M$.

One notable feature of the model is that monthly series have a nonstandard lag structure. For instance, an observation for July will depend explicitly on the lagged data pertaining to February, March, and April monthly series and, in addition, depend on the May and June values through the contemporaneous correlation across these monthly series. In general, this makes determining whether or not the system is stationary more complicated than is typical. See Chen et al. (2012) and Ghysels (2016) for more detail. As a practical matter, because all of our data have been transformed to remove unit roots, we maintain throughout that the VAR is stationary and has a finite-order lag polynomial.

2.3 Priors and Estimation

The dimension of the VAR increases quickly: An additional monthly predictor adds three variables to the VAR. To handle the parameter proliferation problem, we estimate the model using Bayesian methods and utilize a hierarchical shrinkage prior. Bańbura, Giannone, and Reichlin (2010) show that Bayesian VARs can forecast well, even with 100+ variables, when the shrinkage is chosen appropriately such that the prior tightness increases with the model size. Thus, given the size of the model, a careful choice of prior hyperparameters is important. We use the procedure outlined in GLP to choose hyperparameters that maximize the marginal data density for each estimation sample.

We employ a normal-inverse-Wishart prior distribution for the VAR parameters. More formally, let the priors be defined as

$$\Sigma \sim IW(\Psi, d),$$
$$\beta|\Sigma \sim N(0, \Sigma \otimes \Omega),$$

where the degrees-of-freedom parameter of the inverse-Wishart distribution $d = n + 2$, the minimum value that guarantees the existence of the prior mean for Σ. Ψ is a diagonal matrix where each element of the diagonal is set to the residual variance of an AR(1) process for the respective variable in the VAR.[6] Ω is a $k \times k$ matrix, $k = np + 1$, parameterized such that the prior covariance of the regression coefficients takes the following form:

$$cov\left((B_s)_{ij}, (B_r)_{hm} |\Sigma\right) = \begin{cases} \lambda^2 \frac{1}{s^2} \frac{\Sigma_{ih}}{\psi_j/(d-n-1)} & if \quad m = j \text{ and } r = s \\ 0 & otherwise \end{cases}.$$

Thus, while the coefficients $B_1, \ldots, B_p$ are assumed to be independent of each other, coefficients associated with the same variable are allowed to be contemporaneously correlated across different equations. In general, the prior imposes a tighter variance on the distant lags; however, given that in our specification $p = 1$, this feature of the prior is not relevant. The hyperparameter λ governs the overall tightness of the prior by controlling the scale of the variances and covariances of the VAR coefficients. The prior variance on the constant term C is diffuse.

This prior is standard, taken directly from GLP, but reconfigured for a stationary process. We rely on their procedure and accompanying codes to estimate the VAR such that the hyperparameters (in

[6]Limited robustness analysis shows that treating diagonal elements of the scale matrix Ψ as hyperparameters does not generate meaningful gains; thus, we resort to GLP's default implementation of the Minnesota prior.

our case λ) impose an optimal amount of shrinkage consistent with the marginal data density criterion. GLP simulate the posterior distribution of λ based on a standard Metropolis algorithm, under the assumption that the prior for λ follows a gamma distribution with a mode of 0.2 and standard deviation of 0.4, values consistent with those in Sims and Zha (1998). Given the posterior of the hyperparameter λ, we obtain the posterior distribution of the VAR parameters (β and Σ) by drawing from normal-inverse-Wishart posterior distributions implied by the conjugate prior. Our reported results are based on 5,000 draws.

The prior used here is conjugate and has a Kronecker covariance structure that imposes symmetric treatment of variables across equations. The prior covariance matrices for the lag coefficients pertaining to the k-th variable in different equations are proportional. This yields a computational advantage, which is essential when evaluating models' out-of-sample forecasting ability. On the other hand, one might want to treat the lags of different frequency variables differently. Our method could be adapted for more general priors (e.g., Carriero, Clark, and Marcellino 2019 and Chan 2019). We leave this for future research.

The model is estimated every month using a new vintage of real-time data. When constructing the forecasts, we use the posterior distribution of the parameters obtained from the past vintage of monthly data. We hold the posterior constant within a month. For example, analysis over the first calendar quarter would unfold as follows. On January 1, we estimate the model using the vintage of data ending on December 31. The last row of this vintage will have a ragged edge because fourth-quarter data continue to be released during January. We drop the ragged edge and estimate the model. Draws from this posterior are used throughout January. To ensure the real-time nature of the exercise, as we move across data releases within January, new vintages of monthly data supplant those available in the December 31 vintage and are used as predictors when forming the forecasts. A similar pattern of updating the posterior and constructing forecasts consistent with the real-time nature of the data continues as we move across February and March.[7]

[7]We estimate the model once per month rather than once per quarter. In those months for which we have a ragged edge, the ragged edge is dropped. An

Figure 1. Data Releases and Forecast Timing

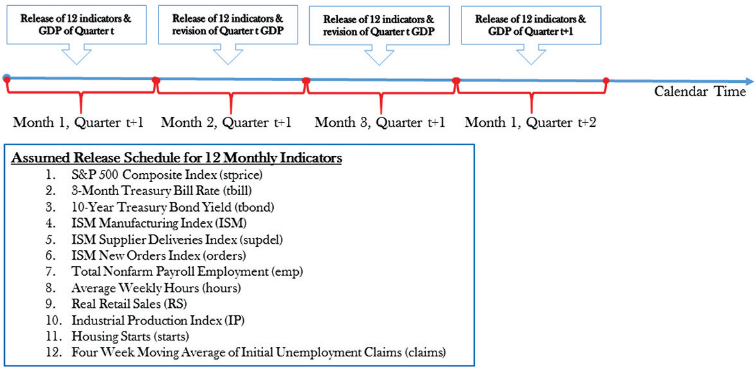

Note: The figure outlines the timing of the data releases and forecasts over the course of the quarter.

3. Forecasting

We formulate the intraquarter forecasting problem in a conditional forecasting framework. We treat the quarter-t information set as complete when the first (advance) release of GDP is released. This happens at the end of the first calendar month in quarter $t + 1$. At this point, we construct unconditional h-step-ahead point and density forecasts. As we move across the quarter, high-frequency data become available. We update the h-step-ahead unconditional forecasts using existing results on conditional forecasting as presented in Waggoner and Zha (1999). The approximate release schedule of the high-frequency data, which determines the sequence by which we update the forecasts, is delineated in section 2.1 and depicted in figure 1.

As figure 1 shows, our results are based on forecast updates across four months. In the context of forecasting current, first-quarter GDP,

---

alternative approach would be to introduce a Tanner and Wong (1987) step as in Waggoner and Zha (1999) (see their algorithm 1). We have performed a limited set of robustness experiments related to this step and, other than adding runtime, found that it had limited effect on our results. Since we want to keep our modeling framework simple to ease its use in the context of policymaking, we do not incorporate this feature into the benchmark setup.

we begin with a two-quarter-ahead forecast formed conditional on all previous-quarter information up to and including the first data release of January. We then update that forecast until obtaining end-of-January claims. When we observe the previous-quarter GDP, we now form the one-quarter-ahead forecast using a complete set of previous-quarter data. As we move into February, we then update this one-quarter-ahead forecast sequentially across each data release. This continues into March and finally ends with end-of-April claims. All together we obtain a sequence of 49 intraquarter forecasts for first-quarter GDP based on the 12 monthly releases in each month and the advance release of previous-quarter GDP.

3.1 Point Forecasts and Predictive Densities

Our forecasting model is estimated using Bayesian techniques, providing a natural characterization of uncertainty. The time-t predictive density of Y at horizon h, $p(Y_{t+h}|\mathbf{Y_t})$, is

$$p(Y_{t+h}|\mathbf{Y_t}) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} p(Y_{t+h}|\mathbf{Y_t}, D, \Sigma)p(D|\Sigma, \mathbf{Y_t})p(\Sigma|\mathbf{Y_t})dDd\Sigma, \tag{3}$$

where $\mathbf{Y_t} = [Y_1, \ldots, Y_t]'$ represents the history of observables up to time t and $p(D|\Sigma, \mathbf{Y_t})$ and $p(\Sigma|\mathbf{Y_t})$ are the time-t posterior distributions of the parameters in D and Σ, respectively.

For each saved draw, $D^{(i)}$ and $\Sigma^{(i)}$, from their respective posteriors, we obtain a forecast draw $\hat{Y}^{(i)}_{t+h|t}$ from the conditional predictive density $p(Y_{t+h}|\mathbf{Y_t}, D^{(i)}, \Sigma^{(i)})$. Collecting these draws across Markov chain Monte Carlo (MCMC) iterations yields the predictive density that accounts for the uncertainty in the estimated parameters (including the uncertainty associated with the setup of the prior tightness) and the uncertainty from the unobservable future shocks. Then, based on arguments made in Gneiting (2011), we obtain the point forecast as the mean of the predictive density.

The posterior distributions $p(D|\Sigma, \mathbf{Y_t})$ and $p(\Sigma|\mathbf{Y_t})$ are readily available as a result of the reduced-form VAR estimation in GLP framework. Next, we consider the simulation from the conditional predictive density $p(Y_{t+h}|\mathbf{Y_t}, D^{(i)}, \Sigma^{(i)})$, keeping in mind that this distribution changes with each new high-frequency release.

3.2 Conditional Predictive Density Simulation

In this section, we describe how we obtain a forecast draw, $\hat{Y}^{(i)}_{t+h|t} \sim p(Y_{t+h}|\mathbf{Y_t}, D^{(i)}, \Sigma^{(i)})$, conditional on the i-th draw of the VAR parameters obtained from the MCMC. In what follows, we suppress the superscript i denoting the MCMC iteration for notational simplicity. As we intimated above, our forecasting procedure has two components, each of which is based on the composition of the information set at the time the forecast is constructed. Specifically, at the time the last-quarter GDP value is released and the information set is deemed complete, the forecast is constructed as an *unconditional* forecast. As the quarter progresses, the forecast is constructed as a *conditional* forecast, where the already-released intraquarter data are treated as restrictions in the forecasting model. In what follows, we delineate our unified approach for producing both h-period-ahead conditional and unconditional forecasts based on well-established results in Waggoner and Zha (1999), implemented with computational simplifications proposed in Jarociński (2010).

Because our forecasting approach is an iterative (rather than direct) one, we demonstrate how to draw $\hat{Y}_{t+h|t}$ for a general h-steps-ahead horizon, assuming we have already obtained draws for the forecasts $\hat{Y}_{t+h-1|t}, \ldots, \hat{Y}_{t+1|t}$. Let $\mu = (I - \Sigma_{l=1}^{p} B_l)^{-1}C$ represent the mean of Y_t implied by the VAR; then, $Z_t = Y_t - \mu$ is the demeaned vector of period-t observables and $\hat{Z}_{t+h|t} = \hat{Y}_{t+h|t} - \mu$ are the demeaned forecast draws. Let $\hat{\mathbf{Z}}_{t+h-1|t} = [\hat{Z}'_{t+h-1|t}, \hat{Z}'_{t+h-2|t}, \ldots, \hat{Z}'_{t+h-p|t}]'$. When $h - j \leq 0$, $\hat{Z}_{t+h-j|t} = Z_{t+h-j}$, i.e., it represents observed data.

We are interested in the joint distribution of the one- to h-step-ahead forecasts of Y obtained at time t. Let $\mathbf{Y} = [Y'_{t+1}, \ldots Y'_{t+h}]'$, $\hat{\mathbf{Y}}^u = [\hat{Y}^{u'}_{t+1|t}, \ldots \hat{Y}^{u'}_{t+h|t}]'$, and $\hat{\mathbf{Z}}^u = [Z_t, \hat{Z}^{u'}_{t+1|t} \ldots \hat{Z}^{u'}_{t+h-1|t}]'$, where superscript u denotes unconditional forecasts. Define $\Phi_j \Sigma^{1/2}$ as the matrix of orthogonalized impulse responses after $j = 1, \ldots, h$ periods and let

$$R_{nh\times nh} = \begin{pmatrix} \Sigma^{1/2} & 0 & & 0 & 0 \\ \Phi_1\Sigma^{1/2} & \Sigma^{1/2} & & 0 & 0 \\ & & \cdots & \Sigma^{1/2} & 0 \\ \Phi_{h-1}\Sigma^{1/2} & \Phi_{h-2}\Sigma^{1/2} & & \Phi_1\Sigma^{1/2} & \Sigma^{1/2} \end{pmatrix},$$

where $\Sigma^{1/2}$ can be obtained as a Cholesky factor of Σ.

Suppose we know a subset m of the future values of $\mathbf{Y}$. Let ω be a $(m \times nh)$ selection matrix, such that $\omega\mathbf{Y}$ results in an $m \times 1$ vector, consisting of the known future elements of $\mathbf{Y}$. Let $r = \omega(\mathbf{Y} - \hat{\mathbf{Y}}^u)$. Then, the known future values can be thought of as linear restrictions on future values of the error term $\varepsilon = [\varepsilon'_{t+1}, \ldots, \varepsilon'_{t+h}]'$. These restrictions take the form $\tilde{R}\varepsilon = r$, where $\tilde{R} = \omega R$. For example, suppose we have one monthly and one quarterly variable and hence, for our model, $n = 4$. Now suppose we are forming a two-quarter-ahead forecast ($h = 2$) of GDP growth and we have observed the first two monthly releases in the current quarter. $\tilde{R}$ is the (2×8) matrix formed by stacking the first and second rows of R.

We can then draw the h-period-ahead conditional forecast $\hat{\mathbf{Y}}^c = [\hat{Y}^{c'}_{t+1|t}, \ldots, \hat{Y}^{c'}_{t+h|t}]'$ from the following conditional predictive distribution:

$$\hat{\mathbf{Y}}^c \sim N\left(\psi_{t+h|t}, \Psi_{t+h|t}\right), \tag{4}$$

where $\psi_{t+h|t}$ captures the conditional mean, while $\Psi_{t+h|t}$ is the conditional variance of the predictive density determined as

$$\psi_{t+h|t} = \iota_h \otimes \mu + (I_h \otimes B)\hat{\mathbf{Z}}^u + \tilde{R}'(\tilde{R}\tilde{R}')^{-1}r, \tag{5}$$

$$\Psi_{t+h|t} = R\left(I - \tilde{R}'(\tilde{R}\tilde{R}')^{-1}\tilde{R}\right)R', \tag{6}$$

and ι_h is an h-dimensioned vector of ones. Equations (5) and (6) follow from equations (6) and (9) of Jarociński (2010).

The first two terms of (5) are invariant to whether we have observed any of the intraquarter data—i.e., whether the forecast is conditional or unconditional. The third term captures the adjustment to the mean as more information, either arising from intraquarter data releases or hypothetical scenarios about future path of the variables, becomes available. In the absence of any conditioning information, model residuals are unrestricted, $\tilde{R} \equiv 0$, and this third term does not affect the mean of the predictive density. In contrast, when conditioning information exists, this third term provides the adjustment needed to revise the unconditional forecast to a conditional one. As equation (6) indicates, conditioning information affects the variance of the predictive distribution as well, thus changing the characterization of forecast uncertainty arising due to unobservable future shocks in the model's reduced-form errors. Conditioning information restricts the model residuals ($\tilde{R} \not\equiv 0$) and, in general, reduces the uncertainty of the predictive density.

We assume that the conditioning variables do not change the posterior distribution of the parameters and simulate the predictive density by relying on the computational simplifications provided in Jarociński (2010). Let USV' denote the singular value decomposition of $\tilde{R}$. Define E as the $m \times m$ diagonal matrix of m singular values of $\tilde{R}$, and let V_1 and V_2 denote matrices formed from the first m and remaining $nh - m$ columns of V, respectively. Under our assumptions, Jarociński (2010) shows that $V_1 E^{-1} U' r + V_2 \eta$, for $\eta \sim N(0, I_{(nh-m \times nh-m)})$, has the same normal distribution as that in (4) and is computationally more efficient in many cases.

3.3 Competing Models

When forecasting real GDP growth, we consider a few alternatives to our mixed-frequency Bayesian VAR. A simple and competitive alternative that has also been considered by CCM is an AR(2), which we reestimate each month with each new vintage of GDP. We use the same prior and estimation procedure proposed by GLP and used for our MF-BVAR.

We also report results associated with the DMS approach to forecasting developed in CCM. As noted previously, CCM models GDP growth directly and reports results on monthly nowcasts of GDP growth using the information available at the time of the Employment Situation Report in each month. Thus, the model size changes over the quarter: in month 1 of the quarter, the model is larger in size, while in month 2 it is the smallest. Table 1 in CCM describes the models explicitly. We extend the analysis in CCM to forecast origins other than the employment release. In particular, we consider origins associated with each data release using the schedule detailed in section 2.1. This provides us with the ability to compare our IMS-based and their DMS-based approaches to forecasting at each data release. The CCM-type models are again estimated using the GLP code described earlier, adjusted to work with autoregressive distributed lag type models. It is worth emphasizing that there will be differences between our results and those obtained by CCM. For example, our prior on quarterly and monthly series is symmetric, while CCM impose a distinct prior on series that are at the quarterly versus monthly frequencies. In addition, our prior is optimized for each estimation sample consistent with the GLP procedure, while the CCM prior is fixed. Finally, CCM also permit stochastic volatility in their model, while we abstract from such in our version.

In addition, we consider both small and large versions of our model and that of CCM. These models are defined relative to those described in CCM. Our large model uses all 12 monthly predictors and 1 quarterly predictor to form a 37-dimensional VAR. These variables are described in section 2.1 and figure 1. Our small model only considers one quarterly and five monthly predictors including the ISM manufacturing index, payroll employment, industrial production, retail sales, and housing starts, and thus forms a 16-dimensional VAR.

Finally, we compare the accuracy of our small and large models with the mean of the responses from the SPF and BCEI. We do this for GDP growth forecasts and for forecasts of some of the monthly variables in our system. For the GDP growth forecasts, we are able to compare forecasts directly because annualized quarterly GDP growth is a component of our model and is also part of both surveys.

For the monthly variables, we only compare the forecasting performance of the models to the SPF. The BCEI contains forecasts for the monthly variables. However, given our focus on GDP growth forecasting as well as the proprietary nature of the BCEI survey, we have decided not to report that comparison here. Comparing the model forecasts of monthly variables to those from the SPF is complicated by the fact that the survey provides forecasts of quarterly aggregates. For example, the SPF forecasts of the 3-month and 10-year Treasury yields are quarterly averages of the daily series, while SPF forecasts of retail sales, industrial production, and housing starts are quarterly averages of the monthly series. Our forecast of the 3-month and 10-year quarterly yield is formed by taking the average of the three monthly averages within the target quarter. Our forecast of the quarterly level of housing starts is formed by cumulating the forecasts of changes from the current level of the series and then averaging the three relevant months in the target quarter. Finally, our forecasts of the quarterly level of industrial production and retail sales are formed by extrapolating from the current level of the series based on the forecasted monthly growth rates at the one- through four-quarter-ahead horizons—and then averaging the three relevant months of the target quarter.[8] The remaining monthly series are either not in the SPF or, in the case of employment, have only been part of the survey for a brief time and hence are not included in the evaluation exercise.

3.4 Evaluation

We evaluate point forecasts using root mean square error (RMSE). We calculate RMSEs after each intraquarter data release, obtaining a term structure of RMSEs as we move across the quarter. We consider two out-of-sample evaluation periods. The first evaluation period is the same as that in CCM, where we use real-time data to obtain nowcasts for the advance release of GDP growth between 1985:Q1 and 2011:Q3. We also consider a comparably sized out-of-sample period ranging from 1992:Q2 to 2017:Q1, which has the advantage that it allows us to focus exclusively on forecasting real GDP growth rather than a mixture of GNP and GDP growth depending on the vintage.

We then evaluate the accuracy of our density forecasts based on the continuous ranked probability score (CRPS). Relative to other scoring functions, such as the log-scores, the CRPS is less sensitive to outliers and puts higher weight on draws from the predictive distribution that are close to but not equal to the outcome (see Gneiting and Raftery 2007 and Gneiting and Ranjan 2011 for further discussion). Similar to the RMSE, the CRPS is defined such that the lower the value, the better the score, and is given by

$$\begin{aligned} CRPS_t(y_{t+h}) &= \int_{-\infty}^{\infty} \left(P(z) - 1\{y_{t+h} \leq z\}\right)^2 dz \\ &= E_p \left| \hat{y}_{t+h|t} - y_{t+h} \right| - \frac{1}{2} E_p \left| \hat{y}_{t+h|t} - \hat{y}'_{t+h|t} \right|, \end{aligned} \tag{7}$$

[8] Industrial production rebases itself seven times across all of our vintages. In order to ensure that the levels at the forecast origin align with those in the target quarter, we unwind the new base year back to the base year at the forecast origin.

where $P(.)$ denotes the cumulative distribution function (CDF) associated with the predictive density $p(y_{t+h}|\mathbf{Y}_t)$, $1\{y_{t+h} \leq z\}$ denotes the indicator function, taking value 1 if the outcome $y_{t+h} \leq z$ and 0 otherwise, and $\hat{y}_{t+h|t}$ and $\hat{y}'_{t+h|t}$ are independent random draws from the conditional predictive density $p(y_{t+h}|\mathbf{Y}_t)$. We compute the CRPS using the empirical CDF-based approximation given in equation (9) of Krueger et al. (2017). As for the case of RMSEs, we obtain average CRPS values for each intraquarter data release and report the term structure of CRPS values for each out-of-sample period.

Pairwise differences in RMSEs and CRPS values across models are evaluated using a standard normal approximation to t-type tests of predictive ability akin to that developed in Diebold and Mariano (1995). Newey and West (1987) standard error estimates are used with lag orders equal to the forecast horizon plus one. All point and density forecasts of real GDP growth are evaluated using the advance release. Point and density forecasts of the monthly series are evaluated using the same vintage of data used to evaluate the real GDP growth forecasts.

4. Forecasting Results

In this section, we delineate some of the advantages of the blocked, mixed-frequency VAR in the context of forecasting. We begin by evaluating the real-time accuracy of our point and density forecasts relative to a handful of competitors outlined in section 3.3. We use figures to present measures of accuracy across all intraquarter horizons. For each figure, we also have a table that presents the value of the measure of accuracy, the ratio of these measures across models, and pairwise tests of equal accuracy—but only for those forecast origins associated with the Employment Situation Report. In our sequencing of data releases, this lines up with "hours." We then illustrate the usefulness of the model for producing scenario-based forecasts of low-frequency variables when the conditioning variable is observed at a higher frequency.

4.1 Current-Quarter Forecasts

In figure 2, we plot the term structure of RMSEs for real GDP growth nowcasts from each of the models as we move across the

Figure 2. RMSE Paths for GDP Growth Nowcasts

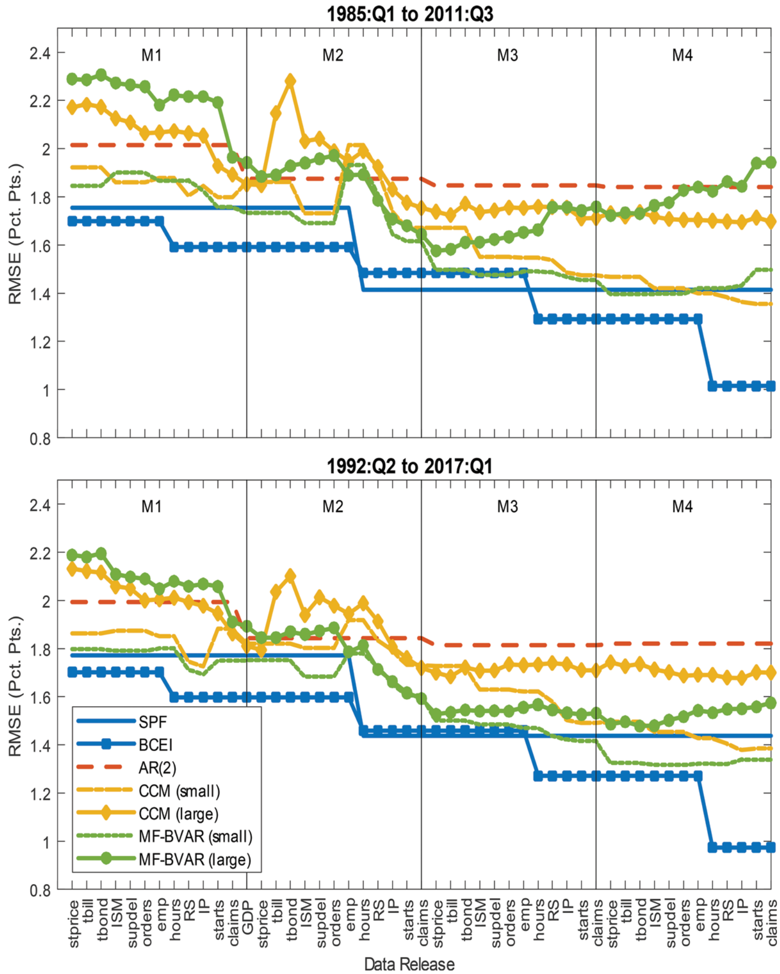

Notes: The figure depicts the RMSE paths for nowcasts of the advance release of the real GDP growth. The upper panel considers the evaluation period in CCM, while the lower panel shows the results associated with real GDP (as opposed to GNP) growth forecasting. Each tick represents a data release in the respective month.

intraquarter forecast origins. The upper panel shows the results for the subsample used by CCM in which nowcasts are generated for 1985:Q1 through 2011:Q3, while the lower panel shows the results for the pure GDP subsample—i.e., nowcasts for 1992:Q2 through 2017:Q1. For comparison with the results in CCM, we begin forecasting on the first day of the first calendar month of quarter $t + 1$ (e.g., January 1 when forecasting Q1 GDP). We then proceed across all data releases until the last high-frequency release prior to the release of advance GDP (e.g., late April when forecasting Q1 GDP).

Each panel in figure 2 contains paths associated with seven forecasts: the SPF, the BCEI, an AR(2), both small and large versions of the CCM models, and both small and large versions of our MF-BVARs. For both large models, we have 49 intraquarter updates associated with three monthly releases of the 12 monthly indicators and the quarter-t GDP advance release at the end of the first calendar month. The small models, on the other hand, are updated only at a subset of these forecast origins. More specifically, we obtain 16 updates to the forecasts, associated with three monthly releases of five monthly indicators and the quarter-t GDP advance release at the end of the first calendar month. Between updates we simply flatline the RMSE path. The AR(2) model updates its forecasts at the end of each month—i.e., after a new release or a revision to previously released GDP numbers.

The current-quarter SPF is timed to arrive after the Employment Situation Report but prior to the release of retail sales. In the first calendar month of the quarter $t + 1$, we report RMSEs for the one-quarter-ahead SPF forecasts released in the second calendar month of the previous quarter. In the case of the BCEI, we have more frequent updates since it is a monthly survey. Typically, the survey is published on the 10th of the month, while the forecasts are collected earlier in the month. Accordingly, each month we time the BCEI forecasts to arrive after the Employment Situation Report but prior to the release of retail sales. For the first three calendar months in the quarter $t + 1$ we use the BCEI nowcasts. In the first calendar month of quarter $t + 2$, prior to the GDP advance release, we use the BCEI backcast. In the first calendar month of the quarter $t + 1$, but prior to the release of the Employment Situation Report, we use the one-quarter-ahead forecast from the survey conducted a month earlier.

Looking at figure 2, perhaps the clearest observation across both evaluation samples is that the SPF and BCEI point forecasts of advance GDP growth are extremely difficult to beat. Further, in the first subsample, the best model-based forecasts tend to be those from the small-models—both CCM and the MF-BVAR—but even these are comparable to the SPF only early in the second month and at the end of the fourth month. The large models generally improve in performance across the first two calendar months of the quarter, but then either stagnate or even deteriorate as we get closer to the GDP release date. For reasons that are not obvious to us, the RMSEs of the direct multistep models both deteriorate sharply in month 2 before improving across the remaining intraquarter releases. The AR(2) performs worse than both survey forecasts, but is often competitive with the other model-based forecasts. The larger models tend to have better performance relative to the AR(2) after the second month of the calendar quarter. Across most forecast origins, the BCEI nowcasts tend to be the most accurate, with the exception that the SPF forecasts tend to be best for the first month of its release.

The models seem to perform much better relative to the SPF when we move forward into the pure GDP subsample (i.e., 1992:Q2–2017:Q1). As before, both small models are competitive with the SPF early in the second month and at the end of the fourth month but are now also competitive over a broader stretch of the quarter. While the large MF-BVAR is generally dominated by its smaller version, the large MF-BVAR has improved substantially and is competitive with the SPF in all periods except early in the first month. Relative to the first evaluation sample, the MF-BVAR models generally outperform the DMS models in terms of their point forecasts. BCEI nowcasts continue their superiority over the other forecasts though, once again, the SPF is very competitive for the first month of its release.

We assess the statistical significance of these results in table 1. The diagonals in the table have two numbers that are obtained as of the release of the Employment Situation Report: (i) the RMSE (below the slash) and (ii) the CRPS value (above the slash). For example, the third diagonal entry in the top panel, 2.01\3.41, indicates that for this sample the AR(2) has an RMSE of 2.01, while the average CRPS is 3.41. The lower off-diagonal portion of each panel in the table reports the ratio of RMSEs in the row-model to the column-model, where numbers greater than 1 indicate that

Table 1. RMSEs and Mean CRPSs of GDP Growth Nowcasts

		SPF	**BCEI**	**AR(2)**	**CCM (Small)**	**CCM (Large)**	**MF-BVAR (Small)**	**MF-BVAR (Large)**
1985:Q1 to 2011:Q3								
M1	SPF	1.75\NA						
	BCEI	0.91	1.59\NA					
	AR(2)	1.15	1.26	2.01\3.41	**1.13**	**1.21**	**1.19**	**1.26**
	CCM (Small)	1.07	**1.18**	0.93	1.88\3.03	**1.08**	**1.06**	**1.12**
	CCM (Large)	1.18	**1.30**	1.03	1.10	2.07\2.82	0.98	1.04
	MF-BVAR (Small)	1.06	**1.17**	0.93	0.99	0.90	1.87\2.86	**1.06**
	MF-BVAR (Large)	**1.27**	**1.40**	1.10	**1.18**	1.07	**1.19**	2.22\2.71
M2	SPF	1.41\NA						
	BCEI	**1.05**	1.48\NA					
	AR(2)	1.33	1.26	1.87\3.23	**1.22**	**1.21**	**1.32**	**1.44**
	CCM (Small)	**1.43**	**1.36**	1.07	2.01\2.64	0.99	**1.08**	**1.18**
	CCM (Large)	**1.41**	**1.34**	1.06	0.99	1.99\2.67	**1.09**	**1.19**
	MF-BVAR (Small)	**1.37**	**1.30**	1.03	0.96	0.97	1.93\2.45	**1.09**
	MF-BVAR (Large)	**1.34**	**1.28**	1.01	0.94	0.95	0.98	1.89\2.24
M3	SPF	1.41\NA						
	BCEI	**0.91**	1.29\NA					
	AR(2)	**1.31**	**1.43**	1.85\3.22	**1.47**	**1.28**	**1.63**	**1.64**
	CCM (Small)	1.09	**1.20**	0.84	1.55\2.19	**0.87**	**1.11**	**1.12**
	CCM (Large)	**1.24**	**1.36**	0.95	1.14	1.76\2.51	**1.27**	**1.28**
	MF-BVAR (Small)	1.05	1.15	0.81	0.96	0.85	1.49\1.97	1.01
	MF-BVAR (Large)	**1.18**	**1.29**	0.90	1.07	0.95	1.12	1.66\1.96
M4	SPF	1.41\NA						
	BCEI	**0.72**	1.01\NA					
	AR(2)	**1.30**	**1.81**	1.84\3.23	**1.52**	**1.29**	**1.73**	**1.64**
	CCM (Small)	0.99	**1.38**	**0.76**	1.40\2.12	**0.84**	**1.14**	1.07
	CCM (Large)	**1.20**	**1.68**	0.92	**1.22**	1.70\2.51	**1.35**	**1.27**
	MF-BVAR (Small)	1.01	**1.40**	0.77	1.02	0.84	1.42\1.86	0.94
	MF-BVAR (Large)	**1.29**	**1.80**	0.99	**1.30**	1.07	**1.28**	1.82\1.07

(*continued*)

Table 1. (Continued)

		SPF	**BCEI**	**AR(2)**	**CCM (Small)**	**CCM (Large)**	**MF-BVAR (Small)**	**MF-BVAR (Large)**
1992:Q2 to 2017:Q1								
M1	SPF	1.77\NA						
	BCEI	0.90	1.60\NA					
	AR(2)	1.12	1.25	1.99\3.21	**1.15**	**1.18**	**1.20**	**1.25**
	CCM (Small)	1.04	**1.16**	0.93	1.85\2.79	1.02	**1.04**	**1.08**
	CCM (Large)	1.13	**1.26**	1.01	1.09	2.01\2.73	1.02	**1.06**
	MF-BVAR (Small)	1.02	**1.13**	0.90	0.97	0.90	1.80\2.67	1.04
	MF-BVAR (Large)	1.17	**1.30**	1.04	1.12	1.3	1.15	2.08\2.57
M2	SPF	1.44\NA						
	BCEI	1.02	1.46\NA					
	AR(2)	1.28	1.26	1.84\3.04	**1.22**	**1.19**	**1.33**	**1.41**
	CCM (Small)	**1.34**	**1.32**	1.04	1.92\2.50	0.98	**1.10**	**1.16**
	CCM (Large)	**1.38**	**1.36**	1.08	1.04	1.99\2.55	**1.12**	**1.18**
	MF-BVAR (Small)	**1.24**	**1.22**	0.97	0.93	0.90	1.78\2.29	**1.06**
	MF-BVAR (Large)	**1.26**	**1.24**	0.98	0.94	0.91	1.02	1.81\2.16
M3	SPF	1.44\NA						
	BCEI	**0.88**	1.27\NA					
	AR(2)	1.26	**1.43**	1.82\3.03	**1.41**	**1.27**	**1.59**	**1.63**
	CCM (Small)	1.13	**1.28**	0.89	1.62\2.15	**0.90**	**1.13**	**1.15**
	CCM (Large)	**1.21**	**1.36**	0.96	1.07	1.73\2.39	**1.26**	**1.28**
	MF-BVAR (Small)	1.02	1.16	0.81	**0.91**	0.85	1.47\1.90	1.02
	MF-BVAR (Large)	1.09	**1.23**	0.86	0.97	0.90	1.07	1.57\1.86
M4	SPF	1.44\NA						
	BCEI	**0.68**	0.97\NA					
	AR(2)	1.27	**1.87**	1.82\3.04	**1.50**	**1.27**	**1.72**	**1.69**
	CCM (Small)	0.99	**1.46**	**0.78**	1.43\2.03	**0.85**	**1.15**	**1.13**
	CCM (Large)	1.18	**1.74**	0.93	1.19	1.69\2.40	**1.36**	**1.34**
	MF-BVAR (Small)	0.92	**1.36**	**0.73**	0.93	**0.78**	1.32\1.76	0.98
	MF-BVAR (Large)	1.07	**1.57**	0.84	1.07	0.91	**1.16**	1.53\1.79

Notes: The table shows the RMSEs (below the slash) and mean CRPSs (above the slash) of each model on the diagonal. The lower off-diagonal portion of each panel in the table reports the ratio of RMSEs in the row-model to the column-model. The upper off-diagonal portion of each panel reports the ratio of average CRPS in the row-model to the column-model. The results are as of the "hours" release in each month. Off-diagonal numbers greater than 1 indicate that the column-model is nominally more accurate. Ratios in bold are statistically different from 1 at the 5 percent significance level.

the column-model is nominally more accurate. Ratios in bold are statistically different from 1 at the 5 percent significance level. The majority of the statistically significant pairwise RMSE comparisons arise when comparing model-based forecasts to either the SPF or the BCEI. Across models, however, there are few differences that are statistically significant. The few that arise tend to do so either early or later in the quarter and simply reinforce what we observed in the figure: the small models tend to be more accurate than the large models.

In figure 3, we provide the same type of term structure but applied to mean CRPS values for each estimated model (and hence not the SPF or the BCEI) across both evaluation samples. In broad terms, the CRPS paths of all models decline as we move across forecast origins regardless of evaluation sample. Again, the exception is that both of the DMS models experience a sharp deterioration in mean CRPS values as we move into month 2. Interestingly, while the large MF-BVAR did not generally perform the best among the models in terms of RMSEs, it typically has the lowest CRPS values across all intraquarter forecast origins.

In table 1, the upper off-diagonal portion of each panel reports each row-model's mean CRPS as of the "hours" release relative to that of the column-model. Again, numbers greater than 1 indicate that the column-model is nominally more accurate. We find many more instances of statistically significant differences across the model-based forecasts. Clearly, while the AR(2) performs reasonably well in terms of RMSEs, the CRPS values are typically much higher than those of the other models. In addition, it is often the case that the CRPS values from the MF-BVAR models are significantly better than those from the DMS models used in CCM.[9]

[9]In unreported results, we have investigated the role in which the additional real and financial variables affect the forecasting performance of the large MF-BVAR relative to the small MF-BVAR. We considered two alternative MF-BVAR models: one in which we add the remaining real variables to the small model and one in which we add the financial variables to the small model. In both cases we find a decline in the accuracy of the point forecasts, relative to the small model, though adding the real series is arguably worse. In contrast, in both cases we find a very modest improvement in the accuracy of the density forecasts relative to the small model.

Figure 3. CRPS Paths for GDP Growth Nowcasts

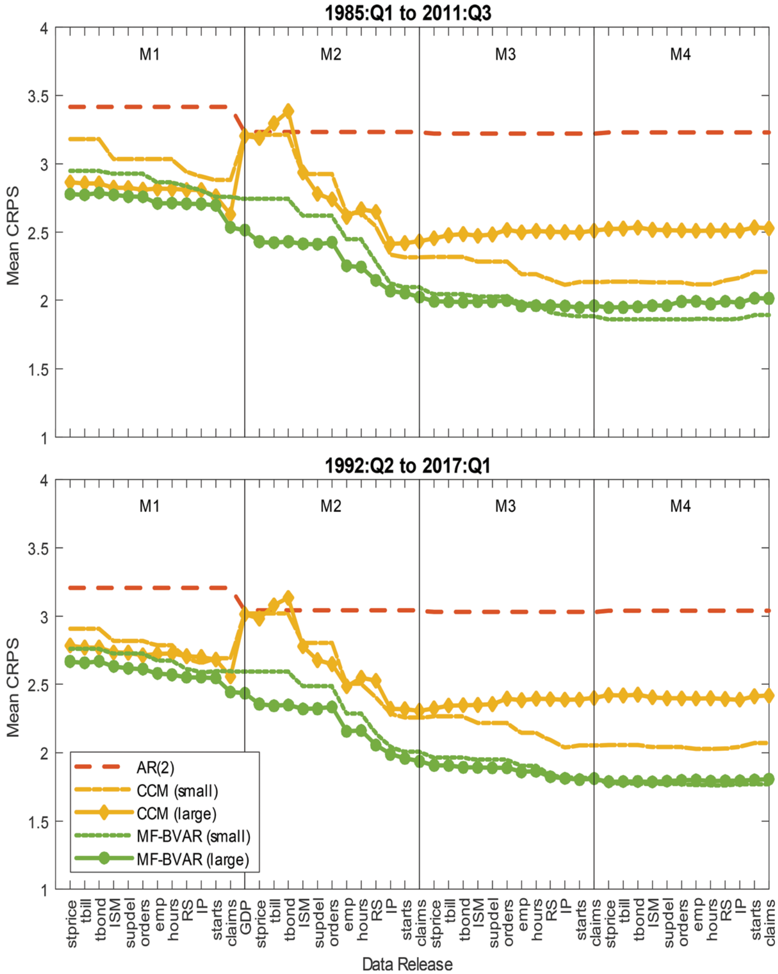

Notes: The figure depicts the average CRPS paths for nowcasts of the advance release of the GDP growth. The upper panel considers the evaluation period in CCM, while the lower panel shows the results associated with real GDP (as opposed to GNP) growth forecasting. Each tick represents a data release in the respective month.

Figure 4. RMSE Paths for Four-Quarter-Ahead Forecasts

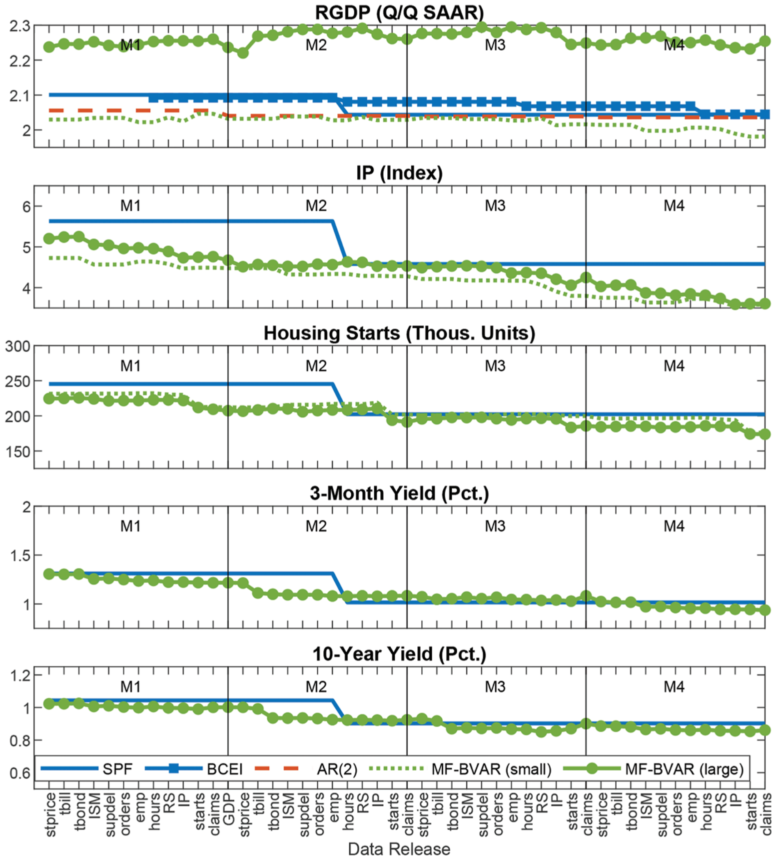

Notes: The figure depicts the RMSE paths for four-quarter-ahead forecasts for the evaluation period of 1992:Q2–2017:Q1. The quarterly AR(2) model is included for real GDP growth only. Each tick represents a data release in the respective month.

4.2 Four-Quarter-Ahead Forecasts

In figures 4 and 5, we provide the term structure of RMSE and CRPS values for four-quarter-ahead forecasts. In contrast to the nowcasting results, we now use the MF-BVARs to forecast a wider range of series consisting of GDP growth, industrial production, housing starts, and both the 3-month and 10-year Treasury yields.

Figure 5. CRPS Paths for Four-Quarter-Ahead Forecasts

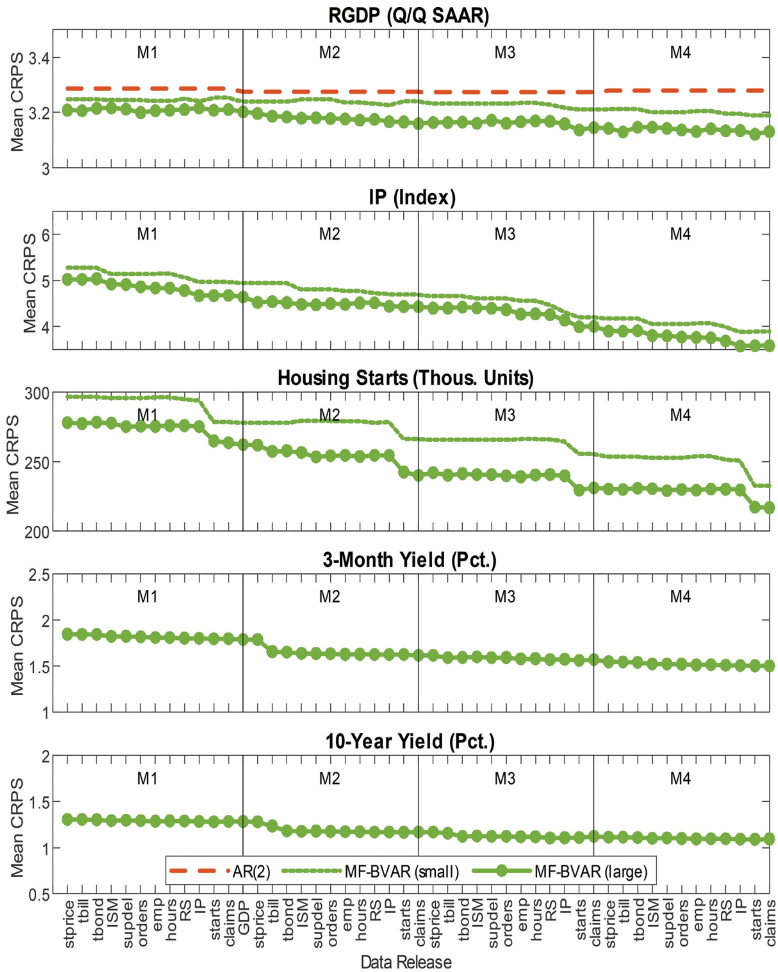

Notes: The figure depicts the average CRPS paths for four-quarter-ahead forecasts for the evaluation period of 1992:Q2–2017:Q1. The quarterly AR(2) model is included for real GDP growth only. Each tick represents a data release in the respective month.

For comparison, we consider SPF forecasts of all five series. In addition, we include BCEI and AR(2) forecasts of GDP growth.[10] When applicable, the RMSE paths include forecasts from the MF-BVARs, SPF, BCEI, and the AR(2); the CRPS paths only include forecasts from the AR(2) and the MF-BVARs. Because CCM do not consider longer horizons forecasts, we do not include the DMS forecast in the comparison.

The top panel of figure 4 shows that the SPF and BCEI point forecasts for real GDP growth exhibit similar accuracy at the four-quarter-ahead horizon. The AR(2) outperforms the surveys in month 1, and early in month 2, but they eventually equalize in performance. The small MF-BVAR performs similar to, and sometimes a bit better than, the AR(2), especially in month 4. The large MF-BVAR is considerably worse than all of the alternatives at the four-quarter-ahead horizon. In contrast, the lower four panels of figure 4 show that the models outperform the SPF for variables other than GDP growth—in particular, industrial production and housing starts—at some intraquarter forecast origins. Four-quarter-ahead forecasts of industrial production and housing starts are as good as or better than the SPF at almost all forecast origins other than those late in the second calendar month. Relative to the SPF, model-based forecasts of yields are a bit better for the 10-year than the 3-month but, in both cases, they track the RMSEs of the SPF fairly closely.

In figure 5, we plot the mean CRPS values as in figure 3, but for the four-quarter-ahead horizon. As was the case before, the values generally decline as we move across the intraquarter forecast origins. This is most evident for industrial production, but also for housing starts. Across all forecast origins and each of the three relevant series (i.e., GDP, industrial production, and housing starts), the large MF-BVAR has a lower CRPS value than that of the small model.

As we did for the GDP nowcasts in table 1, table 2 reports the nominal and relative measures of accuracy for four-quarter-ahead

[10]Note that, for this exercise, we do not consider the BCEI forecasts available prior to the Employment Situation Report in the first calendar month of quarter $t + 1$. Given our timing conventions, there are few observations associated with the requisite five-quarter-ahead forecasts.

Table 2. RMSEs and Mean CRPSs of Four-Quarter-Ahead Forecasts of GDP Growth for the 1992:Q2 to 2017:Q1 Sample

		SPF	BCEI	AR(2)	MF-BVAR (Small)	MF-BVAR (Large)
M1	SPF	2.10\NA				
	BCEI	1.00	2.09\NA			
	AR(2)	0.98	0.98	2.06\3.29	**1.01**	1.02
	MF-BVAR (Small)	0.96	0.97	0.98	2.02\3.24	1.01
	MF-BVAR (Large)	1.07	1.08	**1.10**	**1.11**	2.25\3.21
M2	SPF	2.04\NA				
	BCEI	1.02	2.08\NA			
	AR(2)	1.00	0.98	2.04\3.28	**1.01**	1.03
	MF-BVAR (Small)	0.99	0.97	0.99	2.03\3.24	1.02
	MF-BVAR (Large)	1.12	1.10	1.12	**1.12**	2.28\3.17
M3	SPF	2.04\NA				
	BCEI	1.01	2.07\NA			
	AR(2)	1.00	0.99	2.04\3.28	**1.01**	**1.03**
	MF-BVAR (Small)	0.99	0.98	0.99	2.03\3.24	1.02
	MF-BVAR (Large)	1.12	1.11	**1.12**	**1.13**	2.29\3.17
M4	SPF	2.04\NA				
	BCEI	1.00	2.04\NA			
	AR(2)	1.00	1.00	2.04\3.28	**1.02**	**1.04**
	MF-BVAR (Small)	0.98	0.98	0.99	2.01\3.20	1.02
	MF-BVAR (Large)	1.10	1.10	**1.11**	**1.12**	2.26\3.14

Notes: The table shows the RMSEs (below the slash) and mean CRPSs (above the slash) of each model on the diagonal. The lower off-diagonal portion of each panel in the table reports the ratio of RMSEs in the row-model to the column-model. The upper off-diagonal portion of each panel reports the ratio of average CRPS in the row-model to the column-model. The results are as of the "hours" release in each month. Off-diagonal numbers greater than 1 indicate that the column-model is nominally more accurate. Ratios in bold are statistically different from 1 at the 5 percent significance level.

GDP growth forecasts. Values in bold are ratios for which a pairwise test of equal accuracy indicates statistical significance at the 5 percent level. In general, the results are statistically insignificant; however, in accordance with figures 4 and 5, the small MF-BVAR is significantly more accurate than the large MF-BVAR in terms of point forecasting, while simultaneously being less accurate in terms of density forecasting.[11]

Table 3 reports the accuracy measures and tests of equal predictive accuracy from table 2 but for forecasts of industrial production, housing starts, and 3-month and 10-year yields. In many instances, especially for the 3-month and 10-year yields, differences in accuracy are statistically insignificant. Significant differences do arise when comparing the MF-BVAR density forecasts of industrial production and housing starts. In addition, while not uniform across forecast origins, both MF-BVARs provide statistically significant improvements over the SPF for point forecasts of industrial production.

4.3 Scenario Forecasting

The previous subsections indicate that point forecasts from the surveys are difficult to beat in terms of RMSEs. Even so, there is one thing our model can do that the surveys cannot—produce scenario forecasts designed to guide hypothetical policies. While perhaps not immediately obvious, the block structure of the MF-BVAR permits scenarios other lower-frequency models cannot. In this section, we provide two examples of high-frequency, policy-oriented scenarios and compare their implementation using the MF-BVAR to that of a quarter BVAR with quarterly averaged monthly variables.

In both examples, we consider a central bank that uses a high-frequency interest rate to conduct monetary policy. The goal of the policy is to influence a low-frequency series such as GDP. In a purely

[11] To get a feel for the differential performance of the small and large MF-BVARs for real GDP growth at the four-quarter-horizon, we again considered two alternative MF-BVARs: one in which we add the remaining real variables to the small model and one in which we add the financial variables to the small model. In unreported results, we find that including the financial variables leads to a sharp deterioration of the point forecasts but has little impact on the density forecasts. Adding the real variables has the opposite effect: a sharp improvement of the density forecasts with little impact on the accuracy of the point forecasts.

Table 3. RMSEs and Mean CRPSs of Four-Quarter-Ahead Forecasts of Select Monthly Variables for the 1992:Q2 to 2017:Q1 Sample

		Industrial Production			Housing Starts		
		SPF	**MF-BVAR (Small)**	**MF-BVAR (Large)**	**SPF**	**MF-BVAR (Small)**	**MF-BVAR (Large)**
M1	SPF	5.63\NA			245.39\NA		
	MF-BVAR (Small)	**0.82**	4.64\5.15	**1.06**	**0.95**	232.13\296.33	**1.07**
	MF-BVAR (Large)	**0.88**	1.07	4.96\4.84	0.91	0.96	222.67\275.90
M2	SPF	4.58\NA			202.24\NA		
	MF-BVAR (Small)	0.95	4.33\4.78	1.06	1.07	217.03\279.15	**1.10**
	MF-BVAR (Large)	1.01	1.07	4.63\4.51	1.03	0.96	208.29\253.81
M3	SPF	4.58\NA			202.24\NA		
	MF-BVAR (Small)	0.91	4.17\4.55	1.07	1.00	202.63\266.56	**1.11**
	MF-BVAR (Large)	0.95	1.05	4.37\4.27	0.97	0.97	196.00\240.45
M4	SPF	4.58\NA			202.24\NA		
	MF-BVAR (Small)	**0.81**	3.73\4.07	**1.08**	0.98	197.42\253.90	**1.10**
	MF-BVAR (Large)	**0.83**	1.02	3.81\3.75	0.92	0.94	185.80\230.32

		3-Month Yield		10-Year Yield	
		SPF	**MF-BVAR (Large)**	**SPF**	**MF-BVAR (Large)**
M1	SPF	1.31\NA		1.04\NA	
	MF-BVAR (Large)	0.95	1.24\1.81	0.96	1.01\1.29
M2	SPF	1.02\NA		0.90\NA	
	MF-BVAR (Large)	1.06	1.08\1.63	1.02	0.92\1.17
M3	SPF	1.02\NA		0.90\NA	
	MF-BVAR (Large)	1.03	1.04\1.58	0.96	0.87\1.12
M4	SPF	1.02\NA		0.90\NA	
	MF-BVAR (Large)	0.94	0.96\1.51	0.96	0.86\1.10

Notes: The table shows the RMSEs (below the slash) and mean CRPSs (above the slash) of each model on the diagonal. The lower off-diagonal portion of each panel in the table reports the ratio of RMSEs in the row-model to the column-model. The upper off-diagonal portion of each panel reports the ratio of average CRPS in the row-model to the column-model. The results are as of the “hours” release in each month. Off-diagonal numbers greater than 1 indicate that the column-model is nominally more accurate. Ratios in bold are statistically different from 1 at the 5 percent significance level.

quarterly model, the high-frequency policy rate would likely be averaged across all three months of the quarter, leaving the timing of the rate change obscured.[12] Instead, the MF-BVAR can explicitly account for the timing of the policy rate change within the quarter as well as capture the impact of any intraquarterly data that may have been released.

The specifics of both experiments, and their associated scenarios, are adapted from the pattern of federal funds rate changes made by the Federal Open Market Committee (FOMC) throughout 2017 but, to maintain consistency across sections of the paper, we use the three-month Treasury yield as the policy rate. For example, suppose that on the last business day of January 2017, a scenario forecast is made that assumes the three-month Treasury yield remains constant throughout February but rises 25 basis points (bps) in March (e.g., at the March FOMC meeting). It then remains unchanged until June, at which time it rises another 25 bps. It is then assumed to remain constant throughout much of the year before rising another 25 bps in December and stays constant until the end of the year.

In the MF-BVAR, implementing this scenario is straightforward. We map the month-to-month changes in the policy rate directly into specific variables in the large MF-BVAR: the first, second, and third three-month Treasury yields. For the quarterly BVAR, we first form quarterly averages of the scenario and then form conditional forecasts using this low-frequency approximation to the high-frequency scenario.

In our first experiment, we address two issues related to scenario forecasting. First, holding the forecast origin constant, are there substantive differences between the MF-BVAR and quarterly BVAR forecasts? Second, are there substantive differences among the MF-BVAR forecasts as we receive high-frequency intraquarter information?

In figure 6, we plot multiple scenario forecasts of *total* annualized real GDP growth (cumulative sum of quarter on quarter growth rates) from 2016:Q4 for quarterly horizons 1 through 4. Total—rather than quarter-to-quarter—growth is reported in order to align our forecasts with the fixed-event forecasts used by the FOMC. Each

[12] See Knotek and Zaman (2019) for an exception.

Figure 6. Policy-Rate-Based Scenario Forecasts

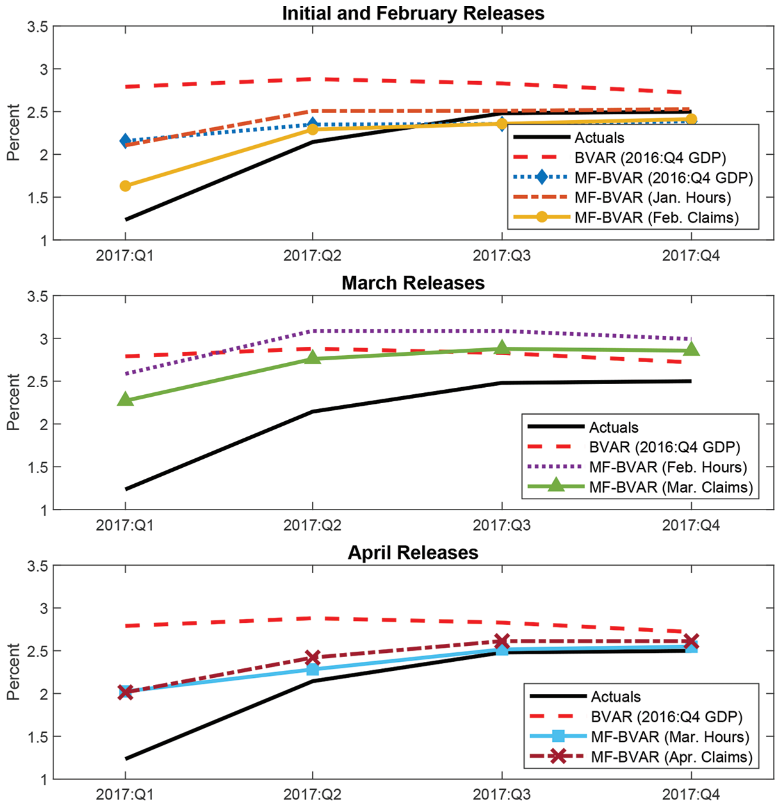

Notes: The figure depicts total real GDP growth forecasts conditional on an assumed path of an interest rate. Each figure shows the forecasts from a quarterly BVAR as well as the actual realized value. Each panel further shows forecasts from the large MF-BVAR. (.) represents the release of a variable the forecast is conditioned on.

forecast is generated using the relevant real-time vintage of data exactly as we did for the one- and four-quarter-ahead forecasting exercises in the previous subsection. Actuals are reported using data from the advance release of 2017:Q4 GDP. It is useful to keep in mind that forecasts from the MF-BVAR will evolve as we obtain intraquarter information, while those from the quarterly BVAR will not.

In the first panel, we report forecasts from both quarterly and MF-BVAR models made from the same end-of-January forecast origin. At this origin, the quarterly BVAR forecasts are uniformly higher than those from the mixed-frequency BVAR. These differences are as large as 75 bps at the one-quarter horizon but narrow substantially at the four-quarter horizon. We then update the MF-BVAR forecast twice as we move across February. The first update aligns with the Employment Situation Report, while the second aligns with the claims report (and hence we observe all February releases). The first of these remains relatively close to the path predicted by the initial MF-BVAR forecast while the latter, in particular, reduces the near-term forecast of GDP growth sharply.

The second panel reports comparably updated forecasts from the MF-BVAR moving across March. Not surprisingly, given the strong employment report in early March 2017, the first March revision shifts the forecasts upwards and is now much more in line with that of the quarterly BVAR—at least for horizons greater than one quarter. Even so, as further data arrives in March, this shift moderates.

The final panel again reports updated forecasts from the MF-BVAR but this time as we move across April. In contrast to March, reported employment growth fell sharply. As a consequence, the scenario paths both decline and are again closely in line with the realized values of total GDP growth—at least relative to those from the quarterly BVAR, which does not get updated as we move across the quarter.

In the first experiment, the hypothetical, high-frequency scenario was held fixed and we investigated the impact of intraquarterly data releases on the subsequent forecast of the low-frequency variable. In our second experiment, we do the opposite: We hold the forecast origin constant and investigate whether or not changes to the intraquarterly timing of the high-frequency scenario affect forecasts of the low-frequency variable. For example, note that in 2017, each of the changes to the policy rate came in the third calendar month of the respective quarter. With this in mind, in our second experiment we hold the forecast origin constant at the end-of-January forecast origin and compare three distinct scenarios. The first continues to maintain that the 25 bps changes are made in the third month of the quarter, but for the other two scenarios we allow the same 25 bps changes to come in the first and second months of the quarter.[13]

[13]Recall that, in the MF-BVAR model, “tbill” is the monthly average of daily values for the three-month Treasury yield. As such, the precise date for the policy

Figure 7. Policy-Rate-Timing Scenario Forecasts

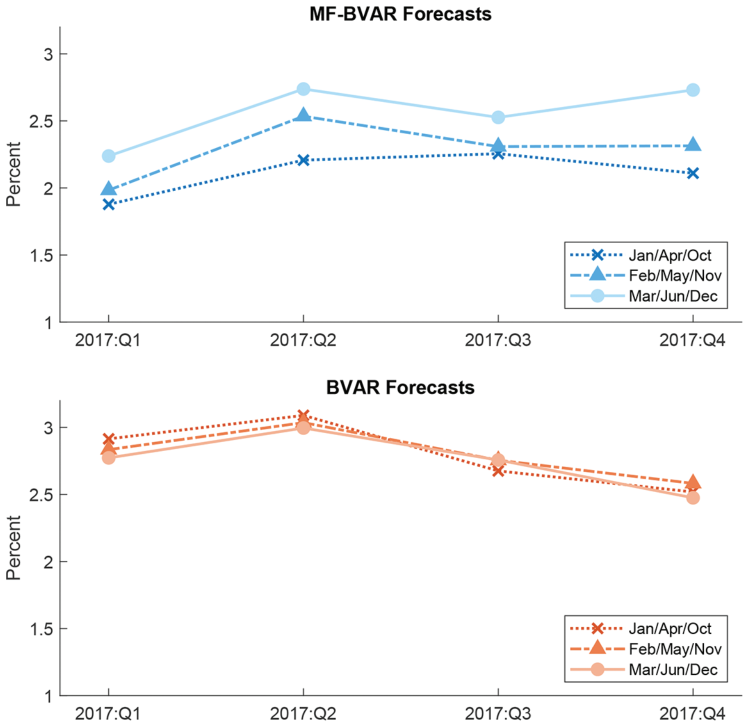

Notes: The figure depicts total real GDP growth forecasts conditional on three different assumed paths of an interest rate: (i) 25 bps increases in January, April, and October 2017, (ii) 25 bps increases in February, May, and November 2017, and (iii) 25 bps increases in March, June, and December 2017. The top panel shows the forecasts from the large MF-BVAR and the bottom panel shows the forecasts from a quarterly BVAR.

Akin to those provided in figure 6, in figure 7 we plot one- through four-quarter-ahead forecasts of total annualized real GDP growth but based on the three distinct scenarios. In the first panel we report

change matters for the average. For the first experiment we assumed that the policy rate changed mid-month, and hence the 25 bps change in the policy rate was spread out over the current and subsequent month. For the second experiment we assume that the policy rate was changed on the first business day of the month, and hence the monthly average of the three-month Treasury yield increases the full 25 bps within that month.

forecasts from the MF-BVAR, while in the second panel we report forecasts from the BVAR. In the first panel we find a clear pattern within the forecasts made by the MF-BVAR model: later policy rate changes lead to higher forecasts of growth. This is in contrast to those generated by the BVAR model in the second panel. Here we find little evidence that the timing of the policy rate changes matters for the forecasts.

While these are only two examples of scenario forecasts, both are realistic, policy-oriented scenarios that suggest some advantages to using the MF-BVAR when modeling mixed frequencies. First, the mixed frequency allows the user to implement detailed high-frequency scenarios that low-frequency models cannot. In addition, since the model is readily revised as high-frequency data is released, we are able to track how the scenario forecasts evolve within a quarter and, by implication, evolve as new data are observed between policy meetings.

5. Conclusion

In this paper, we investigate the usefulness of a particular type of mixed-frequency VAR, delineated by Ghysels (2016), in the context of real-time forecasting. In this model, multiple high-frequency intraquarter observations are treated as distinct quarterly observations, and a standard VAR is formed based on these series. Because this leads to a high-dimensional VAR, we estimate the parameters using standard shrinkage-based Bayesian methods. In addition, since the model is just a Bayesian VAR, existing methods developed by Waggoner and Zha (1999) can be used to produce end-of-quarter forecasts as well as intraquarter forecasts that account for high-frequency data releases.

In terms of both point and density forecasts, we find that the iterated multistep approach to mixed-frequency nowcasting of real GDP growth performs as well as direct multistep variants developed in Carriero, Clark, and Marcellino (2015) and, depending on the specific version of our model, can be as accurate as the SPF and BCEI at certain very short horizons. One advantage of the MF-BVAR model is that the same model can be used for both near- and longer-horizon forecasting. As such, we also compare the forecasting performance of our model to the survey forecasts at the four-quarter horizon. At this longer horizon we find that the small-scale MF-BVAR model is as good as or better than the SPF and BCEI when forecasting real GDP growth. In addition, the MF-BVAR generally outperforms the SPF when forecasting monthly variables such as industrial production, housing starts, and, to a lesser extent, 3-month and 10-year Treasury yields. Finally, we discuss the usefulness of the model for central bank-type scenario forecasting when the scenario is delineated using high-frequency observables, but the object of interest is only observed at the lower frequency.

References

Bacchiocchi, E., A. Bastianin, A. Missale, and E. Rossi. 2016. "Structural Analysis with Mixed Frequencies: Monetary Policy, Uncertainty and Gross Capital Flows." Manuscript.

Bańbura, M., D. Giannone, and L. Reichlin. 2010. "Large Bayesian Vector Auto Regressions." *Journal of Applied Econometrics* 25 (1): 71–92.

Bhansali, R. J. 1997. "Direct Autoregressive Predictors for Multistep Prediction: Order Selection and Performance Relative to Plug-In Predictors." *Statistica Sinica* 7 (2): 425–49.

Bittanti, S., P. Colaneri, and G. De Nicolao. 1988. "The Difference Periodic Riccati Equation for the Periodic Prediction Problem." *IEEE Transactions on Automatic Control* 33 (8): 706–12.

Blasques, F., S. J. Koopman, M. Mallee, and Z. Zhang. 2016. "Weighted Maximum Likelihood for Dynamic Factor Analysis and Forecasting with Mixed Frequency Data." *Journal of Econometrics* 193 (2): 405–17.

Brave, S., R. A. Butters, and A. Justiniano. 2019. "Forecasting Economic Activity with Mixed Frequency Bayesian VARs." *International Journal of Forecasting* 35 (4): 1692–1707.

Carriero, A., T. E. Clark, and M. Marcellino. 2015. "Realtime Nowcasting with a Bayesian Mixed Frequency Model with Stochastic Volatility." *Journal of Royal Statistical Society: Series A* 178 (4): 837–62.

———. 2019. "Large Vector Autoregressions with Stochastic Volatility and Non-Conjugate Priors." *Journal of Econometrics* 212 (1): 137–54.

Chan, J. 2019. "Asymmetric Conjugate Priors for Large Bayesian VARs." Mimeo.

Chen, W., B. D. O. Anderson, M. Deistler, and A. Filler. 2012. "Properties of Blocked Linear Systems." *Automatica* 48 (10): 2520–25.

Corrado, C., and M. Green. 1988. "Reducing Uncertainty in Short-term Projections: Linkage of Monthly and Quarterly Models." *Journal of Forecasting* 7 (2): 77–102.

Diebold, F. X., and R. S. Mariano. 1995. "Comparing Predictive Accuracy." *Journal of Business and Economic Statistics* 13 (3): 253–63.

Eraker, B., C. W. Chiu, A. T. Foerster, T. B. Kim, and H. D. Seoane. 2015. "Bayesian Mixed Frequency VARs." *Journal of Financial Econometrics* 13 (3): 698–721.

Foroni, C., and M. Marcellino. 2013. "A Survey of Econometric Methods for Mixed-Frequency Data." Working Paper No. 2013/06, Norges Bank.

Foroni, C., M. Marcellino, and C. Schumacher. 2015. "U-MIDAS: MIDAS Regressions with Unrestricted Lag Polynomials." *Journal of the Royal Statistical Society: Series A* 178 (1): 57–82.

Ghysels, E. 2016. "Macroeconomics and the Reality of Mixed Frequency Data." *Journal of Econometrics* 193 (2): 294–314.

Ghysels, E., J. B. Hill, and K. Motegi. 2016. "Testing for Granger Causality with Mixed Frequency Data." *Journal of Econometrics* 192 (1): 207–30.

Giannone, D., M. Lenza, and G. E. Primiceri. 2015. "Prior Selection for Vector Autoregressions." *Review of Economics and Statistics* 97 (2): 412–35.

Giannone, D., L. Reichlin, and D. Small. 2008. "Nowcasting: The Real-time Informational Content of Macroeconomic Data." *Journal of Monetary Economics* 55 (4): 665–76.

Gneiting, T. 2011. "Making and Evaluating Point Forecasts." *Journal of the American Statistical Association* 106 (494): 746–62.

Gneiting, T., and A. E. Raftery. 2007. "Strictly Proper Scoring Rules, Prediction, and Estimation." *Journal of the American Statistical Association* 102 (477): 359–78.

Gneiting, T., and R. Ranjan. 2011. "Comparing Density Forecasts Using Threshold and Quantile Weighted Proper Scoring Rules." *Journal of Business and Economic Statistics* 29 (3): 411–22.

Jarociński, M. 2010. "Conditional Forecasts and Uncertainty about Forecast Revisions in Vector Autoregressions." *Economics Letters* 108 (3): 257–59.

Knotek, E., and S. Zaman. 2019. "Financial Nowcasts and Their Usefulness in Macroeconomic Forecasting." *International Journal of Forecasting* 35 (4): 1708–24.

Krueger, F., S. Lerch, T. L. Thorarinsdottir, and T. Gneiting. 2017. "Probabilistic Forecasting and Comparative Model Assessment Based on Markov Chain Monte Carlo Output." Manuscript.

Kuzin, V., M. Marcellino, and C. Schumacher. 2011. "MIDAS vs. Mixed-frequency VAR: Nowcasting GDP in the Euro Area." *International Journal of Forecasting* 27 (2): 529–42.

Marcellino, M., J. H. Stock, and M. W. Watson. 2006. "A Comparison of Direct and Iterated Multistep AR Methods for Forecasting Macroeconomic Time Series." *Journal of Econometrics* 135 (1–2): 499–26.

Mittnik, S., and P. A. Zadrozny. 2005. "Forecasting Quarterly German GDP at Monthly Intervals using Monthly Ifo Business Conditions Data." In *Ifo Survey Data in Business Cycle and Monetary Policy Analysis*, ed. J.-E. Sturm and T. Wollmershäuser, 19–48. Heidelberg: Physica-Verlag.

Newey, W. K., and K. D. West. 1987. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica* 55 (3): 703–8.

Parigi, G., and G. Schlitzer. 1995. "Quarterly Forecasts of the Italian Business Cycle by Means of Monthly Economic Indicators." *Journal of Forecasting* 14 (2): 117–41.

Schorfheide, F. 2005. "VAR Forecasting under Misspecification." *Journal of Econometrics* 128 (1): 99–136.

Schorfheide, F., and D. Song. 2015. "Real-Time Forecasting with a Mixed-Frequency VAR." *Journal of Business and Economic Statistics* 33 (3): 366–80.

Sims, C. A., and T. Zha. 1998. "Bayesian Methods for Dynamic Multivariate Models." *International Economic Review* 39 (4): 949–68.

Tanner, M. A., and W. H. Wong. 1987. "The Calculation of Posterior Distributions by Data Augmentation." *Journal of the American Statistical Association* 82 (398): 528–40.

Waggoner, D. F., and T. Zha. 1999. "Conditional Forecasts in Dynamic Multivariate Models." *Review of Economics and Statistics* 81 (4): 639–51.

Zadrozny, P. A. 1990. "Forecasting U.S. GNP at Monthly Intervals with an Estimated Bivariate Time Series Model." *Economic Review* (Federal Reserve Bank of Atlanta) (November/December): 2–15.