Federal Reserve Bank of Minneapolis
Research Department Staff Report 493

December 2013

A 14-Variable Mixed-Frequency VAR Model*

Kenneth Beauchemin

Federal Reserve Bank of Minneapolis

ABSTRACT

This paper describes recent modifications to the mixed-frequency model vector autoregression (MF-VAR) constructed by Schorfheide and Song (2012). The changes to the model are restricted solely to the set of variables included in the model; all other aspects of the model remain unchanged. Forecast evaluations are conducted to gauge the accuracy of the revised model to standard benchmarks and the original model.

Keywords: Bayesian Vector Autoregression; Forecasting
JEL Classification: C11, C32, C53

---

*Email: ken.beauchemin@mpls.frb.org. I thank Sam Schulhofer-Wohl for helpful comments and suggestions.

1 Introduction

This paper describes recent modifications to the mixed-frequency vector autoregression (MF-VAR) constructed by Schorfheide and Song (2012). The original model has been in use at the Federal Reserve Bank of Minneapolis for more than two years. A noteworthy feature of the model is that it combines data measured at both monthly and quarterly frequencies. The primary advantage of the mixed-frequency approach is that it can use more timely monthly data to help forecast quarterly variables—primarily GDP and associated national income and product account concepts—that are available on a less timely basis. The algorithm used to solve the model uses all available monthly information to construct forecasts of the quarterly variables. The changes we make to the original model are restricted solely to the set of variables that increase their number from 11 to 14; all other aspects of the model are unchanged.

Setting the mixed-frequency feature aside, the model is a descendant of the statistical approach to forecasting developed by Doan, Litterman, and Sims (1984). Specifically, it is a vector autoregression (VAR) set in a Bayesian framework that allows the introduction of extra-sample information based on prior beliefs of macroeconomic time series behavior. The "prior" information helps to counter the problems of forecast degradation due to overfitting. The precise structure of the prior information scheme is a refinement of Doan et al. (1984) and is primarily based on work by Sims and Zha (1998).

The paper is written to be brief but self-contained for those familiar with the Bayesian VAR approach to forecasting. Those interested in further details and a fuller understanding should consult the original paper by Schorfheide and Song (2012) and associated references. The following section provides an overview of the VAR model and the Bayesian framework. Section 3 discusses variable selection and the motivation for the updated list of variables. Section 4 presents a series of forecast evaluations that document forecast accuracy to standard benchmarks, including a direct comparison with the original specification. Section 5 concludes.

2 Model and Prior Specification

Let $y_t = (y_{1t}\ y_{2t}\ \dots\ y_{nt})'$ be an $n \times 1$ data vector of n random variables. The model is an n-variable VAR(p) we first write as

$$y_t = \Phi_1 y_{t-1} + \cdots + \Phi_p y_{t-p} + \Phi_c + u_t, \qquad u_t \sim \text{i.i.d. } N(0, \Sigma) \tag{1}$$

for $t = 1, \cdots, T$. In this expression, $\Phi_1, \dots, \Phi_p$ are $n \times n$ matrices of VAR coefficients, $\Phi_c = (c_1, c_2, \dots, c_n)'$ is an n-dimensional vector of constants, and $\Sigma = E u_t u_t'$. Each equation in the VAR model contains $k = np + 1$ regressors. For notational convenience, the VAR system (1) can be written more compactly by grouping the coefficient matrices into the $n \times k$ matrix $\Phi = [\Phi_1 \cdots\ \Phi_p\ \ \Phi_c]$ and defining the $k \times 1$ vector $x_t = \left(y_{t-1}'\ \cdots\ y_{t-p}'\ 1\right)'$. Then,

$$y_t = \Phi x_t + u_t\ . \tag{2}$$

Furthermore, if one allows the slight abuse of notation where the matrix Φ is formed by stacking the Φ_i's vertically rather than horizontally as in (2), the VAR can be written even more compactly as

$$Y = X\Phi + U, \tag{3}$$

where

$$Y = \begin{bmatrix} y_1' \\ \vdots \\ y_T' \end{bmatrix}, \qquad X = \begin{bmatrix} x_1' \\ \vdots \\ x_T' \end{bmatrix}, \qquad x_t' = \begin{bmatrix} y_{t-1}' & \cdots & y_{t-p}' & 1 \end{bmatrix}, \qquad U = \begin{bmatrix} u_1' \\ \vdots \\ u_T' \end{bmatrix}.$$

In (3), X is a $T \times k$ matrix, and Y and U are $T \times n$ matrices.

VAR models are richly parameterized; for the 14-variable, 6-lag model presented here, each equation has $k = 85$ regressors. With only a limited data history, the estimated model is subject to forecast degradation due to overfitting. Approaching the problem by limiting the number of regressors exposes the model to misspecification bias, which also compromises forecasting accuracy. To balance these risks, Bayesian methods are typically applied to shrink the coefficient estimates for $\Phi_1, \ldots, \Phi_p$ and Σ to their "prior" means in a reasonably large model.

Since the u_t are i.i.d, the VAR system can be used to construct the likelihood function, which gives the probability of observing the sequence of random variables $Y_{1:T} = \{y_1, \ldots, y_T\}$ conditional on the parameters (Φ, Σ) and the p initial observations:

$$p\left(Y_{1:T} \mid \Phi, \Sigma, Y_{1-p:0}\right) = \Pi_{t=1}^{T} p\left(Y_t \mid \Phi, \Sigma, Y_{1-p:t-1}\right). \tag{4}$$

To be succinct, we use the following shorthand for the likelihood function: $p\left(Y \mid \Phi, \Sigma\right) \equiv p\left(Y_{1:T} \mid \Phi, \Sigma, Y_{1-p:0}\right)$. Bayes' rule implies that

$$p\left(\Phi, \Sigma \mid Y\right) = \frac{p\left(Y \mid \Phi, \Sigma\right) p\left(\Phi, \Sigma\right)}{p\left(Y\right)} \tag{5}$$

where $p\left(\Phi, \Sigma\right)$ is the prior distribution (a subjective assessment of the probabilities on (Φ, Σ) before the data on Y are observed), and $p\left(\Phi, \Sigma \mid Y\right)$ is the posterior distribution (the assessment of probabilities (Φ, Σ) once Y has been observed). The posterior distribution is the primary object of importance in Bayesian inference and prediction.

In the current version of the model, we continue to use an updated version of the "Minnesota" prior originally introduced by Doan et al. (1984). Based on work by Sims and Zha (1998), Schorfheide and Song (2012) apply a prior distribution in the form of a multivariate normal inverted Wishart ($MNIW$). Among other advantages, the $MNIW$ prior produces a posterior distribution that is also $MNIW$ (i.e., it is the "natural conjugate" prior).[1] The prior is implemented with the mixed estimation method proposed by Litterman (1986), whereby the observed data set is augmented with dummy observations.[2]

Prior beliefs regarding the variances are determined by a parsimonious vector of hyperparameters $\lambda > 0$. The prior is parameterized so that as $\lambda \to 0$, the prior becomes noninformative (or "flat"), essentially producing ordinary least squares (OLS) estimates for the posterior means, and as $\lambda \to \infty$, the prior is said to be "dogmatic" and the prior means become the posterior

[1] The precise specification of the prior distribution is described in Del Negro and Schorfheide (2011). In brief, the vectorized coefficient matrices Φ are distributed as a multivariate normal given covariance matrix Σ, and the covariance matrix is distributed as an inverted Wishart, which is the multivariate generalization of the inverted gamma distribution.

[2] See Del Negro and Schorfheide (2013).

means. Considering the complexity produced by the large number of means and covariances controlled by the hyperparameters, purely Bayesian methods are of little use in selecting their values. Instead, we continue to select hyperparameter values that jointly maximize the marginal likelihood of the data.[3] Giannone, Lenza, and Primiceri (2012) have shown this empirical Bayesian procedure to produce better forecasting accuracy against a number of alternatives. An attractive feature of the marginal density associated with the $MNIW$ prior is that it can be obtained in closed form.[4]

The basic idea behind the Minnesota prior is that macroeconomic times series behavior is fairly well described as a collection of random walks correlated only through the innovations. In terms of the VAR system (1),

$$\Phi_1 = I_n, \quad \Phi_1 = \cdots = \Phi_p = 0. \tag{6}$$

We optimize the marginal likelihood over four hyperparameters. The first determines the degree of belief in the unit root behavior expressed by (6). The Minnesota prior also expresses the belief that the quantitative importance of a variable fades as the lag lengthens. Assumptions on the decay rate of prior variances with distant lags receive tighter prior variances, implying that they are more likely to be zero than shorter lagged coefficients. A hyperparameter controls the overall tightness of the decay scheme. Another prior that has been shown useful for forecast accuracy is the sum-of-coefficients prior (Doan et al., 1984). This prior allows for inexact differencing in the sense that a variable's own lags sum to one:

$$\Phi_1 + \cdots + \Phi_p = I_n \ . \tag{7}$$

In term of time series behavior, when lagged values of a time series are at a particular level, the same level is likely to be a good forecast for that variable. In other words, time series are assumed

[3] Given the data Y, the marginal data density $p(Y)$ is the missing factor of proportionality $p(Y)$ that forces (5) to hold with equality. Specifically,

$$p(\Phi, \Sigma \mid Y) = [p(Y \mid \Phi, \Sigma)\, p(\Phi, \Sigma)]\, p(Y)^{-1},$$

which follows from Bayes' rule. In Bayesian econometrics it is the principal measure of model fit to the data and is used extensively in model comparison.

[4] See Bauwens, Lubrano, and Richard (1999) for a derivation.

to display "persistence"; another hyperparameter is reserved to express the degree of belief in persistence. A last hyperparameter governs one's belief in "copersistence": when all lagged variables are (separately) at particular levels, then all variables tend to persist simultaneously at those levels.[5]

To operationalize the mixed-frequency feature of the model, the estimation procedure exploits the Kalman filter and its ability to handle missing observations (i.e., the unobserved monthly values of quarterly variables). The VAR (1) is written as a first-order state vector equation and is augmented by a measurement equation ensuring that the inferred monthly observations of the quarterly variables average to the observed quarterly values. To estimate the model, a two-step Gibbs sampler is used, which draws from the $MNIW$ distribution of VAR parameters conditional on the inferred monthly observations and, alternatively, applies a simulation smoother to the state-space representation to draw the missing monthly observations conditional on VAR parameters.[6]

3 Variable Selection

We have two motivations for updating the set of variables. As background, every member of the Federal Reserve Board and each Federal Reserve System bank president participates in the Summary of Economic Projections (SEP) as part of the Federal Open Market Committee's policy process. The survey is compiled by the Board four times annually. The participants currently register their projections on aspects of five variables: real GDP, the unemployment rate, personal consumption expenditures (PCE) prices, the PCE index excluding food and energy prices (the core PCE index), and the federal funds rate. As a side note, participants are instructed to conditional their outlooks on what they perceive to be appropriate monetary policy.

Although the original model contains real GDP, the unemployment rate, and the federal funds rate, the consumer price index (CPI) is the consumer price variable; the corresponding core CPI index is not included. We substitute the PCE index for the CPI. We also add the core

[5] Del Negro and Schorfheide (2011) provide a clear exposition of the $MNIW$ natural conjugate prior and the dummy variables used to implement it.

[6] See Schorfheide and Song (2012) for details.

PCE index to the model to complete the list of SEP variables. Also, with both PCE indexes included, the revised model implicitly contains important information on energy prices (and food prices to a lesser extent).

Our second objective is to augment the model with variables that may offer a partial understanding of why forecasts change over time. Although our primary focus is on unconditional forecasting, theory-inspired variable selection facilitates unconditional forecasting experiments that may help "inform causal hypotheses" (Doan et al., 1984) of monetary policy and the business cycle. Variable selection also provides a basis for choosing informative overidentifying restrictions for structural analysis. For guidance, we turn to the New Keynesian class of dynamic stochastic general equilibrium (DSGE) models that have gained influence in many central bank research departments. Because unit labor costs lie at the heart of inflation dynamics in these models, we add average hourly earnings as the key labor compensation variable. Together with real GDP and aggregate labor hours, the addition of hourly earnings implies a monthly measure of unit labor costs that can be backed out of the forecast.[7] In terms of variable selection, the core of the revised MF-VAR model is similar to the smaller VAR constructed by Christiano, Eichenbaum, and Evans (2005) to identify the effects of a monetary policy shock and to help quantify their DSGE model.

Finally, we introduce the Moody's Baa corporate bond yield to the model. This addition implies a credit spread (Baa corporate yield minus 10-year Treasury yield) in addition to the existing term spread (10-year Treasury yield minus the federal funds rate) and provides a measure of financial market stress that may be useful for forecasting turning points.

Table 1 compares the new set of variables to the original one. In summary, the PCE index replaces the CPI as the main price variable, and the core PCE index is added. Average hourly earnings and the yield on Moody's Baa-rated corporate bonds are also added to the model. The set of quarterly variables (real GDP, fixed investment, and government purchases) has not changed. Table 1 also indicates how each of the variables is transformed for the MF-VAR.

[7] Although the Bureau of Labor Statistics publishes quarterly data on unit labor costs and its components in the nonfarm business sector (and others), in principle we could add productivity and hourly compensation or even unit labor costs by itself. This approach has at least two drawbacks. The first is the usual concern over forecast degradation due to overfitting—particularly when adding highly collinear variables. The second problem is computational; adding additional quarterly variables is substantially more costly in computing time than adding monthly variables.

4 Forecast Evaluation

Our main purpose in this section is to evaluate the forecasting performance of the revised model relative to standard benchmarks and the original model. We focus our attention on the five model variables that are part of the SEP (*GDPR*, *UR*, *PC*, *PCXFE*, and *RFF*), but provide results for all model variables for completeness.

Prediction in a Bayesian framework is based on the posterior predictive density (or "predictive density" for short). It provides a complete probability assessment of future values of the model variables given current and past observations. Let $Y_{T+1:T+H} = \left(y'_{T+1}, y'_{T+2,}, \ldots, y'_{T+H}\right)'$ represent an arbitrary forecast path in the set of all possible future paths. Constructing the predictive density requires us to assign a probability to each path:

$$p\left(Y_{T+1:T+H} \mid Y^T\right) = \int p\left(Y_{T+1:T+H}, \Theta \mid Y_{1-p:T}\right) d\Theta \ , \tag{8}$$

where $\Theta = (\Phi, \Sigma)$, Φ is the vector of VAR coefficients and Σ is the variance-covariance matrix of shocks. The integrand in (8) is the joint density of model parameters and future variable observations. Using the rules of probability, it can be written

$$p\left(Y_{T+1:T+H}, \Theta \mid Y_T\right) = p\left(Y_{T+1:T+H} \mid Y_{1-p:T}, \Theta\right) p\left(\Theta \mid Y_{1-p:T}\right) \ . \tag{9}$$

The two sources of forecast uncertainty are highlighted by this expression. The first term on the right-hand side of (9) describes the uncertainty on future observables given the observed data and model parameters or, equivalently, the forecast uncertainty due to future disturbances that impact the VAR. The second term is the model posterior distribution describing parameter uncertainty. Both distributions have analytical expressions under the Normal-inverted Wishart prior, so simple direct Monte Carlo sampling can be used to produce an approximation to the predictive density.

In what we follows, we provide forecast accuracy metrics for point forecasts. We generate point forecasts using the "pseudo-iterated" approach in which parameter uncertainty is integrated out. The one-step-ahead forecast is obtained using the posterior mean Φ:

$$\widehat{y}_{T+1} = \overline{\Phi}_c + \overline{\Phi}_1 y_T + \overline{\Phi}_2 y_{T-1} + \ldots + \overline{\Phi}_p y_{T-p+1}.$$

The remaining $h = 1, \ldots, H-1$ step-ahead point forecasts are then computed by recursive substitution. For example, the $H = 2$ forecast is then computed as

$$\widehat{y}_{T+2} = \overline{\Phi}_c + \overline{\Phi}_1 \widehat{y}_{T+1} + \overline{\Phi}_2 y_T + \ldots + \overline{\Phi}_p y_{Tt-p+2}$$

and so on. More generally,

$$\widehat{y}_{T+h} = \overline{\Phi}_c + \overline{\Phi}_1 \widehat{y}_{T+h-1} + \overline{\Phi}_2 \widehat{y}_{T+h-2} + \ldots + \overline{\Phi}_p \widehat{y}_{T+h-p}, \quad h = 1, \ldots, H \ , \tag{10}$$

where $\widehat{y}_{T+h} = y_{T+h-p}$ for $h \leq p$. Alternatively, using the notation in (2) we can write

$$\widehat{y}_{T+h} = \overline{\Phi}^h x_t \ .$$

The forecast evaluations are conducted on a recursive basis in which the sample period is lengthened by one observation for each forecast. The initial sample period runs from 1968M1 to 1986M12 with the 1967M7–1967M12 observations serving as the pre-sample to accommodate the six lags. A 36-step (month) ahead forecast is then computed covering the 1987M1–1989M12 period as described above. In the next recursion, the sample is updated to 1968M1–1987M1 and the point forecast computed for the 1987M2–1990M1 period. The process continues until the last forecast that accommodates a three-year interval covering the end of the data sample can be constructed. That recursion uses the 1968M1–2010M6 sample to forecast the 2010M7–2013M6 period. New hyperparameters are computed for each recursion, implying (potentially) different prior means for all forecasts.

The forecast evaluations are conducted in “pseudo real time,” meaning that we do not use vintage (or real-time) data, that is, that were available when a forecast would have been initially performed. Schorfheide and Song (2012) evaluate the original version of the model with real-time forecasts to facilitate comparisons with the Greenbook forecasts produced by the Federal Reserve Board of Governors staff. We doubt that the current revision in model variables has been extensive enough to disrupt their overall conclusions. Interested readers should consult that paper for details. The use of only the most recent vintage of data (at the time of our evaluations) allows comparisons to a wide range of other studies and also economizes on computer time, since the Kalman-smoother step of the posterior simulator procedure can be eliminated. The monthly observations for the three quarterly series are obtained (for each model) by running the full two-step Gibbs sampler over the entire sample period 1968M1–2013M6.

Because the actual quarterly variables are only observed at that frequency, we evaluate forecasts of the quarterly averages in our analysis even though the model is solved at the underlying monthly frequency. We abuse notation slightly so that $h = 1, \dots, H$ is counted in quarters rather than months. Forecasts are evaluated over the 1987Q1–2013Q2 period, which allows for a 20-year initial evaluation window. This period is also singled out as one characterized by a single monetary policy regime.

The original 11-variable specification and the revised 14-variable model are treated symmetrically with one exception. Recall that the original used the headline CPI as its consumer price variable, whereas the updated model favors the headline PCE and core PCE indexes. To facilitate comparison, we modify the original model to use the PCE index instead of the CPI.[8] Under that substitution, the revised 14-variable model nests the slightly modified version of the original 11-variable model.

Forecast evaluations are based on the mean squared forecast error (MSFE) statistic. Let T_0 denote the beginning of the evaluation period minus one period (1986Q4) and T_1 the end period (2013Q2). The MSFE is defined as

$$MSFE_{i,h} = \frac{\sum_{t=T_0}^{T_1-h} (y_{i,t+h}^{data} - \widetilde{y}_{t+h})^2}{T_1 - h - T_0 + 1}$$

for each forecast variable i and forecast horizon $h = 1, \dots, H$. Before MSFEs are computed, the simulated projections $\widehat{y}_{t+h}$ are transformed back to original units $\widetilde{y}_{t+h}$ according to the transformations indicated in Table 1. The root mean squared forecast error is given by $\sqrt{MSFE_{i,h}}$. Table 2 reports the root MSFEs for all variables in the revised model for horizons $h = 1, \dots, 8$.

[8] Our results show that the 11-variable model with the PCE index in place of the CPI tracks consumer prices (as measured) much better, with almost no performances changes noted in other varialbles.

We report MSFE statistics relative to three different benchmark model forecasts. First, we evaluate the forecasting accuracy gained by using (optimized) informative priors with a comparison of the Bayesian MF-VAR forecasts to those generated by the same system under flat priors (OLS). Next, we compare the Bayesian MF-VAR forecast to one generated by the univariate AR(6) process for each variable (denoted AR). Because overfitting is of little threat due to the parsimonious specification, the autoregressive models are estimated using OLS. This comparison helps gauge the value of the cross-correlation information contained in VAR coefficients. Finally, we compare the forecasting accuracy of the 14-variable model relative to the original 11-variable model (designated OR).

Defining $MSFE_{i,h}^{BVAR}$ as the MSFE for the Bayesian VAR with optimized hyperparameters and $MSFE_{i,h}^{m}$ where $m \in \{OLS, AR, OR\}$ as the ones corresponding to each benchmark model, the relative mean squared forecast error (RMSFE) statistic is expressed as the ratio of the former to the latter,

$$RMSFE_{i,h} = \frac{MSFE_{i,h}^{BVAR}}{MSFE_{i,h}^{m}} \ ,$$

so that values less than one imply better forecasts from the 14-variable model. Table 3 shows that using informative priors produces dramatically better forecasts than the same model using flat priors, with the largest accuracy gains for the federal funds rate and the smallest for PCE and core PCE prices. In nearly all cases, the informative priors produce a better forecast, with a single exception: average hourly earnings in the first two quarters of the horizon.

In Table 4, we report the RMSFEs generated by the Bayesian VAR model and the classically estimated AR(6) benchmark models. Overall, the results favor the Bayesian VAR, which outperforms the AR(6) models in over 60% of the cases. The notable exceptions are for the price and wage variables (PC, $PCXFE$, $EARNS$), which for the most part, do substantially worse. In these cases, the cross-covariance information embedded in these equations is counterproductive. This suggests that the symmetry property that treats the prior shrinkage of own-lags in the same way as other lags may be overly restrictive in the case of price and wage forecasting. If so, a case could possibly be made for a more flexible specification, but only by sacrificing the considerable computational conveniences of using the $MNIW$ prior.[9]

Finally, Table 5 shows how the revised 14-variable model compares with the original 11-variable model. Although the models are very close overall, the revised model tends to perform better in the two- and three-year horizons and worse in the near term. Although the revised model performs better in only 32% of the cases in the first four quarters, it does better in roughly 60% of the cases in the second two years. With the exception of the first two quarters of the horizon, the revised model predicts real GDP ($GDPR$) better than the original model, with the advantage generally growing as the forecast horizon expands. The differences in unemployment rate (UR) forecasts mimic the same pattern as real GDP, but the differences are small. Forecasts of consumer prices are substantially improved in the revised version. Forecast accuracy improves for each of the 12 forecast horizon quarters and by as much as 7.5% for $h = 4, 5$. The new model does, however, give some ground on the accuracy of federal funds rate forecasts.

5 Summary

We expand the 11-variable MF-VAR to include 14 variables. The new version of the model includes all five variables submitted to the FOMC's Summary of Economic Projections. The PCE price index was substituted for the CPI, and the core PCE price index has been added. The addition of average hourly earnings potentially adds a key component for understanding changes in inflation forecasts. Along with real GDP and aggregate hours worked, the model's output now implies a proxy for unit labor costs—the driving force of inflation dynamics in New Keynesian DSGE models. Finally, the addition of Moody's corporate Baa yield adds credit spread information to the model to help anticipate economic downturns.

Forecast evaluations show that the models are close in forecast accuracy. The revised model better forecasts consumer prices and real GDP—the latter in the medium and longer term. The revised and original models are evenly matched on the unemployment rate, but the revised model does not does as well as the original model regarding forecasts for the federal funds rate. Since the revised model improves the forecast accuracy for consumer prices, both models

[9]The $MNIW$ natural conjugate prior is one of the very few distributional assumptions that allow for direct Monte Carlo sampling for posterior simulation. Most other specifications require Monte Carlo Markov chain (MCMC) posterior simulation methods. Layering that complication on top of the multiple-frequency feature of the model would be computationally impractical.

produce price forecasts that do not outperform those from a simpler univariate autoregression.

Ongoing research explores ways to improve price forecasts and assess the contribution of introduced variables to forecast interpretation.

References

[1] Bauwens L., M. Lubrano, and J-F. Richard, 1999. *Bayesian Inference in Dynamic Econometric Models*, Oxford University Press.

[2] Christiano, L. J., M. Eichenbaum, and C. L. Evans, 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113(1), 1-35.

[3] Del Negro, M., and F. Schorfheide, 2011. "Bayesian Macroeconometrics," in *The Oxford Handbook of Bayesian Econometrics*, ed. J. Geweke, G. Koop, and H. Van Dijk, 293-389, Oxford University Press.

[4] Doan, T., R. Litterman, and C. Sims, 1984. "Forecasting and Conditional Projection Using Realistic Prior Distributions," *Econometric Reviews*, 3(1), 1-100.

[5] Giannone, D., M. Lenza, and G. Primiceri, 2012. "Prior Selection for Vector Autoregressions," NBER Working Paper 18467.

[6] Litterman, R., 1986. "Forecasting with Bayesian Vector Autoregressions—Five Years of Experience, *Journal of Business and Economic Statistics*, 4, 25-38.

[7] Schorfheide, F., and D. Song, 2012. "Real-Time Forecasting with a Mixed-Frequency VAR," Federal Reserve Bank of Minneapolis, Research Department Working Paper 701.

[8] Sims, C. A., and T. Zha, 1998. "Bayesian Methods for Dynamic Multivariate Models," *International Economic Review*, 39(4), 948-68.

Table 1. Updated Model Variables

Code	Series	Change	Transform
GDPR*	Real GDP (chained 2005 dollars)		log-level
UR	Unemployment rate (%)		level/100
PC	PCE price index	Replaces CPI	log-level
PCXFE	Core PCE price index	New	log-level
RFF	Effective federal funds rate (%)		level/100
LHRS	Index of aggregate weekly hours		log-level
EARNS	Average hourly earnings ($)	New	log-level
IP	Industrial production index		log-level
CONSR	Real personal consumption expenditures		log-level
IFIXR*	Real fixed investment		log-level
GOVR*	Real government purchases		log-level
RTCM10	10-year Treasury note yield (%)		level/100
RBAA	Moody's Baa corporate bond yield (%)	New	level/100
SP500	S&P 500 composite stock price index		log-level

*Quarterly time series.

Table 2. Root Mean Squared Forecast Errors: Revised Model

	Forecast Horizon in Quarters							
y_i	1	2	3	4	5	6	7	8
GDPR	52.49	104.95	167.85	229.15	286.34	338.98	387.30	433.00
UR	0.17	0.17	0.55	0.76	0.97	1.15	1.30	1.43
PC	0.27	0.58	0.82	1.03	1.25	1.43	1.67	1.94
PCXFE	0.14	0.30	0.46	0.64	0.82	1.01	1.23	1.47
RFF	0.49	1.04	1.37	1.65	1.93	2.15	2.33	2.50
LHRS	0.39	0.85	1.41	2.00	2.58	3.10	3.54	3.91
EARNS	0.03	0.06	0.09	0.12	0.14	0.17	0.20	0.22
IP	0.63	1.44	2.34	3.19	3.93	4.57	5.07	5.48
CONSR	37.24	72.80	110.89	149.50	187.26	221.94	254.64	286.35
IFIXR	19.88	48.86	80.23	111.89	143.61	174.24	202.50	228.43
GOVR	11.34	23.77	30.58	39.01	48.34	58.49	67.93	76.62
RTCM10	0.30	0.65	0.82	0.91	1.00	1.04	1.08	1.14
RBAA	0.27	0.52	0.68	0.76	0.82	0.83	0.84	0.89
SP500	52.26	104.71	149.29	191.16	230.40	265.33	298.72	332.99

Table 3. Forecast Comparison: Informative vs.Flat Priors

	Forecast Horizon in Quarters											
y_i	1	2	3	4	5	6	7	8	9	10	11	12
GDPR	0.769	0.760	0.775	0.803	0.824	0.836	0.854	0.867	0.873	0.875	0.871	0.867
UR	0.804	0.732	0.750	0.805	0.839	0.870	0.899	0.925	0.945	0.952	0.946	0.924
PC	0.964	0.890	0.798	0.726	0.683	0.642	0.621	0.611	0.601	0.604	0.607	0.608
PCXFE	0.914	0.820	0.753	0.715	0.684	0.653	0.635	0.625	0.618	0.616	0.617	0.618
RFF	0.719	0.590	0.574	0.585	0.561	0.558	0.560	0.564	0.579	0.596	0.619	0.650
LHRS	0.881	0.819	0.848	0.884	0.900	0.916	0.929	0.934	0.935	0.927	0.913	0.896
EARNS	1.085	1.012	0.903	0.842	0.800	0.762	0.734	0.708	0.685	0.663	0.644	0.628
IP	0.828	0.737	0.779	0.829	0.855	0.879	0.895	0.907	0.912	0.907	0.897	0.885
CONSR	0.917	0.912	0.949	0.949	0.958	0.959	0.958	0.949	0.934	0.917	0.899	0.882
IFIXR	0.813	0.818	0.859	0.887	0.904	0.907	0.906	0.903	0.896	0.886	0.879	0.875
GOVR	0.889	0.682	0.649	0.653	0.656	0.681	0.689	0.695	0.690	0.694	0.708	0.729
RTCM10	0.871	0.769	0.760	0.720	0.689	0.665	0.649	0.627	0.610	0.598	0.597	0.603
RBAA	0.947	0.818	0.823	0.781	0.720	0.662	0.607	0.566	0.547	0.543	0.551	0.560
SP500	0.963	0.898	0.885	0.885	0.880	0.857	0.844	0.833	0.817	0.805	0.802	0.799

	h = 1,...,12	h = 1,...,4	h = 5,...,8	h = 9,...,12
Mean	0.782	0.818	0.773	0.755
Median	0.809	0.819	0.812	0.764
Min	0.543	0.574	0.558	0.543
Max	1.085	1.085	0.959	0.952
% < 1	0.988	0.964	1.000	1.000

Table 4. Forecast Comparison: Revised MF-VAR vs. AR(6)

	Forecast Horizon in Quarters											
y_i	1	2	3	4	5	6	7	8	9	10	11	12
GDPR	0.772	0.871	0.938	0.979	1.000	1.011	1.011	1.014	1.013	1.015	1.015	1.015
UR	0.933	0.917	0.942	0.961	0.992	1.007	1.015	1.018	1.021	1.024	1.027	1.034
PC	0.991	0.983	1.002	1.017	1.063	1.126	1.180	1.250	1.323	1.377	1.443	1.509
PCXFE	1.205	1.381	1.402	1.444	1.493	1.540	1.583	1.606	1.619	1.631	1.646	1.656
RFF	1.198	1.066	0.973	0.923	0.908	0.891	0.874	0.864	0.856	0.851	0.847	0.845
LHRS	0.920	0.934	0.952	0.960	0.965	0.961	0.946	0.922	0.900	0.879	0.857	0.838
EARNS	1.461	1.425	1.328	1.263	1.207	1.147	1.105	1.074	1.053	1.039	1.028	1.018
IP	0.868	0.892	0.904	0.913	0.919	0.920	0.912	0.900	0.892	0.888	0.882	0.876
CONSR	1.047	1.088	1.103	1.089	1.087	1.076	1.062	1.051	1.044	1.040	1.036	1.034
IFIXR	0.813	0.937	0.958	0.973	0.975	0.966	0.954	0.942	0.935	0.928	0.923	0.921
GOVR	0.802	0.891	0.815	0.780	0.749	0.735	0.722	0.703	0.684	0.672	0.663	0.653
RTCM10	0.849	0.995	0.994	0.933	0.889	0.838	0.790	0.766	0.744	0.720	0.699	0.680
RBAA	0.844	0.936	0.901	0.838	0.786	0.726	0.678	0.663	0.667	0.673	0.684	0.691
SP500	0.998	1.021	1.012	1.003	0.988	0.967	0.939	0.916	0.897	0.876	0.854	0.836

	h = 1,...,12	h = 1,...,4	h = 5,...,8	h = 9,...,12
Mean	0.989	1.006	0.989	0.973
Median	0.956	0.960	0.966	0.011
Min	0.653	0.772	0.663	0.653
Max	1.656	1.461	1.606	1.656
% < 1	0.607	0.661	0.589	0.571

Table 5. Forecast Comparison: Revised MF-VAR vs. Original MF-VAR

	Forecast Horizon in Quarters											
y_i	1	2	3	4	5	6	7	8	9	10	11	12
GDPR	1.043	1.019	0.999	0.989	0.978	0.972	0.972	0.971	0.968	0.968	0.968	0.970
UR	1.019	1.040	1.033	1.016	1.005	0.994	0.992	0.992	0.995	0.997	0.999	1.003
PC	0.990	0.960	0.942	0.928	0.925	0.929	0.933	0.939	0.943	0.946	0.949	0.952
PCXFE	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
RFF	1.127	1.016	1.006	1.030	1.038	1.043	1.050	1.052	1.052	1.052	1.051	1.051
LHRS	1.112	1.099	1.064	1.034	1.013	1.000	0.994	0.991	0.989	0.990	0.992	0.994
EARNS	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
IP	1.057	1.027	1.008	0.998	0.984	0.977	0.973	0.968	0.964	0.962	0.964	0.966
CONSR	1.104	1.027	1.009	0.987	0.978	0.975	0.975	0.975	0.975	0.975	0.975	0.975
IFIXR	1.082	1.037	1.009	0.987	0.977	0.975	0.974	0.973	0.972	0.973	0.975	0.977
GOVR	1.097	1.036	1.066	1.081	1.089	1.097	1.102	1.102	1.106	1.108	1.107	1.108
RTCM10	0.981	0.995	0.990	0.996	1.001	1.010	1.023	1.018	1.016	1.013	1.013	1.017
RBAA	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
SP500	1.060	1.002	0.994	1.004	1.011	1.020	1.026	1.034	1.041	1.049	1.055	1.062

	h = 1,...,12	h = 1,...,4	h = 5,...,8	h = 9,...,12
Mean	1.010	1.025	1.000	1.004
Median	1.000	1.017	0.992	0.991
Min	0.925	0.928	0.925	0.943
Max	1.127	1.127	1.102	1.108
% < 1	0.508	0.318	0.591	0.614