Sebastian Ankargren*, Måns Unosson and Yukai Yang

A Flexible Mixed-Frequency Vector Autoregression with a Steady-State Prior

<https://doi.org/10.1515/jtse-2018-0034>

Received December 18, 2018; accepted April 24, 2020

Abstract: We propose a Bayesian vector autoregressive (VAR) model for mixed-frequency data. Our model is based on the mean-adjusted parametrization of the VAR and allows for an explicit prior on the “steady states” (unconditional means) of the included variables. Based on recent developments in the literature, we discuss extensions of the model that improve the flexibility of the modeling approach. These extensions include a hierarchical shrinkage prior for the steady-state parameters, and the use of stochastic volatility to model heteroskedasticity. We put the proposed model to use in a forecast evaluation using US data consisting of 10 monthly and three quarterly variables. The results show that the predictive ability typically benefits from using mixed-frequency data, and that improvement can be obtained for both monthly and quarterly variables. We also find that the steady-state prior generally enhances the accuracy of the forecasts, and that accounting for heteroskedasticity by means of stochastic volatility usually provides additional improvements, although not for all variables.

Keywords: VAR, state space models, nowcasting, macroeconometrics, stochastic volatility

1 Introduction

The vector autoregressive model (VAR) is a commonly used tool in applied macroeconometrics, in part because of its simplicity. Over the years, VAR models have developed in many different directions under both frequentist and Bayesian paradigms. The Bayesian approach offers the attractive ability to easily incorporate soft restrictions and shrinkage, which ameliorate the issue of overparametrization.

***Corresponding author: Sebastian Ankargren**, Department of Statistics, Uppsala University, Uppsala, Sweden, E-mail: sebastian.ankargren@statistics.uu.se. <https://orcid.org/0000-0003-4415-8734>

Måns Unosson and Yukai Yang: Department of Statistics, Uppsala University, Uppsala, Sweden. <https://orcid.org/0000-0002-2623-8549> (Y. Yang)

Within the Bayesian framework itself, a large number of papers have developed prior distributions for the parameters in VAR models. Many of these are, in one way or another, variations of the Minnesota prior proposed by Litterman (1986) (see for example the book chapters Del Negro and Schorfheide 2011; Karlsson 2013). Gains in computational power have led to further alternatives in the choice of prior distribution as intractable posteriors can efficiently be sampled using Markov Chain Monte Carlo (MCMC) methods such as the Gibbs sampler (Gelfand and Smith 1990; Kadiyala and Karlsson 1997).

A particular development in the Bayesian VAR literature is the steady-state prior proposed by Villani (2009). The prior is based on a mean-adjusted form of the VAR where the unconditional mean is explicitly parameterized. This seemingly innocuous reparametrization is justified by the fact that practitioners and analysts often have prior information regarding the steady-state (or unconditional mean) readily available. In the standard parametrization, a prior on the unconditional mean is only implicit as a function of the other parameters' priors. Because the forecast in a stationary VAR converges to the unconditional mean as the horizon increases, a prior for the steady-state parameters can help retain the long-run forecast in the direction implied by theory, even if the model is estimated during a period of divergence.¹

Another modeling feature that modern VARs often include is stochastic volatility. In many macroeconomic applications, a typical characteristic of the data is that the volatility has varied over time. By fitting VARs with constant volatility, the estimated error covariance matrix attempts to balance periods of low and high volatility and find a compromise. Consequently, the predictive distribution does not account for the current level of volatility. Seminal contributions with respect to stochastic volatility were first made by Primiceri (2005); Cogley and Sargent (2005) and numerous follow-up studies have since documented the usefulness of stochastic volatility for forecasting, see e.g. work by Clark (2011); D'Agostino, Gambetti, and Giannone (2013); Clark and Ravazzolo (2015); Carriero, Clark, and Marcellino (2016). Because of the established utility thereof, we also allow for more flexibility in our model by modeling time variation in the error covariance matrix.

VARs are often estimated on a quarterly basis, see e.g. Stock and Watson (2001); Adolfson, Lindé, and Villani (2007). The reason is simply that many

¹ A related, but conceptually different approach, is a prior for the long run developed by Giannone, Lenza, and Primiceri (2019). Their long-run prior is devised to discipline the model's long-run behavior by incorporating economic theory. The approach is a fully Bayesian approach to cointegration, and common stochastic trends play a particularly important role. In the steady-state approach we follow in this work, economic theory plays a lesser role and we assume the model to be stable with no common stochastic trends. For applications where common trends and cointegration do play a role, the long-run prior is an attractive alternative to the steady-state prior approach that we take.

variables of interest are unavailable at higher frequencies, although the majority is often sampled monthly if not even more frequently. When the data are available at different frequencies, common practice is to aggregate high-frequency variables to the lowest frequency present. Such an aggregation incurs a loss of information for variables measured throughout the quarter: the aggregated quarterly values are typically weighted sums of the constituent months and so any information carried by a within-quarter trend or pattern will be disregarded by the aggregation. From a forecasting perspective an analyst will be unconsciously forced to disregard part of the information set when constructing a forecast from within a quarter as the most recent realizations are only available for the high-frequency variables. Another reason for utilizing higher frequencies of the data is that the number of observations is increased. A VAR estimated on data collected over, say, 10 years makes use of 120 observations of the monthly variables instead of being limited to the 40 aggregated quarterly observations.

Multiple approaches to dealing with the problem of mixed frequencies are available in the literature. Mixed data sampling (MIDAS) regressions and the MIDAS VAR proposed by Ghysels, Sinko, and Valkanov (2007) and Ghysels (2016), respectively, use fractional lag polynomials to regress a low-frequency variable on lags of itself as well as high-frequency lags of other variables. This approach is predominantly frequentist, although Bayesian versions are available (Rodriguez and Puggioni 2010; Ghysels 2016). A second approach, which is the focus of this work, is to exploit the general ability of state-space modeling to handle missing observations (Harvey and Pierse 1984). Eraker et al. (2015), concerned with Bayesian estimation, used this idea to treat intra-quarterly values of quarterly variables as missing data and proposed measurement and state-transition equations for the monthly VAR. Schorfheide and Song (2015) considered forecasting using a construction along the lines of Carter and Kohn (1994) and provided empirical evidence that the mixed-frequency VAR improved forecasts of 11 US macroeconomic variables as compared to a quarterly VAR. In terms of flexible time-varying models with mixed-frequency data, Cimadomo and D'Agostino (2016) employed the mixed-frequency VAR together with time-varying parameters and stochastic volatility to cope with a change in frequency of the data. Following up on the work by Schorfheide and Song (2015), Götz and Hauzenberger (2018) recently showed that more flexible models that include stochastic volatility tend to improve forecasts also within this framework.

The main contribution of this paper is that we extend the mixed-frequency toolbox by incorporating prior information on the steady states, and by adding stochastic volatility to the model. Thus, we effectively combine the steady-state parametrization of Villani (2009) with the state-space modeling approach for mixed-frequency data of Schorfheide and Song (2015) and the common stochastic

volatility (CSV) model proposed by Carriero et al. (2016). The proposed model accommodates explicit modeling of the unconditional mean with data measured at different frequencies. In order to employ the model in a realistic forecasting situation, we use a real-time dataset consisting of 13 macroeconomic variables for the US, where 10 of the variables are sampled monthly, and the remaining three are available quarterly. We implement the steady-state prior using the standard Villani (2009) approach, and using the hierarchical structure presented by Louzis (2019). In our empirical application, we find that, for most variables, mixed-frequency data, stochastic volatility, and steady-state information improve forecasting accuracy as compared to models without any of the aforementioned features.

The structure of the paper is as follows. Section 2 describes the main methodology, Section 3 provides information about the data and details about the implementation, and Section 4 evaluates the forecasting performance. Section 5 concludes.

2 Combining Mixed Frequencies with Steady-State Beliefs

The mixed-frequency method adopted in this work is a state space-based model which follows the work by Mariano and Murasawa (2010); Schorfheide and Song (2015); Eraker et al. (2015). There are several modeling approaches available for handling mixed-frequency data, including MIDAS (Ghysels et al. 2007), bridge equations (Baffigi, Golinelli, and Parigi 2004) and factor models (Mariano and Murasawa 2003; Giannone, Reichlin, and Small 2008). We do not review these further here, but instead refer the reader to the survey by Forni and Marcellino (2013) and an early comparison conducted by Kuzin, Marcellino, and Schumacher (2011).

2.1 State-Space Representation of the Mixed-Frequency Model

To cope with mixed observed frequencies of the data, we assume the system to be evolving at the highest available frequency. This assumption frames the problem of frequency mismatch as a missing data problem. By doing so, the approach naturally lends itself to a state-space representation of the system in which the underlying monthly series of the quarterly variables become the latent states of the system. Because we have a mix of monthly and quarterly frequencies in our empirical application, we will in the following proceed with the presentation of the model from this perspective. It should, however, be noted that other compositions of frequencies are viable within the same framework.

The VAR model at the core of the analysis is specified for the high-frequency and partially missing variables. More specifically, a VAR(p) for the $n \times 1$ vector z_t is employed such that

$$\Pi(L)z_t = \Phi d_t + u_t, u_t \sim N(0, \Sigma_t), \quad (1)$$

where $\Pi(L) = (I_n - \Pi_1 L - \Pi_2 L^2 - \dots - \Pi_p L^p)$ is a p -th order invertible lag polynomial, d_t is an $m \times 1$ vector of deterministic components and Φ is an $n \times m$ matrix of parameters. The time index t is here monthly. We let the error term u_t be heteroskedastic and return to the specifics thereof in Section 2.2.

The model in (1) is a conventional VAR specification, but, in the spirit of Villani (2009), we instead employ the mean-adjusted form as

$$\Pi(L)(z_t - \Psi d_t) = u_t, \quad (2)$$

where $\Psi = [\Pi(L)]^{-1}\Phi$ it can be readily confirmed that $E(z_t|\Pi, \Psi, \Sigma) = \Psi d_t =: \mu_t$, and thus μ_t is the unconditional mean—*steady state*—of the process. The steady-state representation (2) requires an explicit prior on the steady state parameters.

However, common practice is to use (1) with a loose prior on Φ , which implicitly defines an intricate (but loose) prior on Ψ and, subsequently, μ_t . We argue that in many applications, the parametrization in (2) is more convenient as it allows for a more natural elicitation of prior beliefs. In what follows, we will extend the work of Villani (2009) such that (2) may still constitute a viable option in the presence of mixed frequencies.

Next, we partition the high-frequency underlying process z_t as, $z_t = (z'_{m,t}, z'_{q,t})'$ where $z_{m,t}$ represents the n_m monthly and $z_{q,t}$ the n_q quarterly variables. Recall that the time t here takes the highest frequency, i. e. monthly. A ubiquitous problem with macroeconomic data is that what is observed varies between months such that z_t is not always fully observed.

To distinguish between the underlying process and actual observations, we denote the latter by y_t . A consequence of all variables not being observed at every time point t is that the dimension n_t of y_t is not always equal to $n = n_m + n_q$. The observed data in y_t are generally supposed to be some linear aggregate of $Z_t = (z'_t, \dots, z'_{t-p+1})'$ such that

$$y_t = \begin{pmatrix} y_{m,t} \\ y_{q,t} \end{pmatrix} = \begin{pmatrix} I_{n_m} & 0 \\ 0 & M_{q,t} \end{pmatrix} \begin{pmatrix} I_{n_m} & 0 \\ 0 & \Lambda_q \end{pmatrix} Z_t = M_t \Lambda Z_t, \quad (3)$$

where $M_{q,t}$ and Λ_q are deterministic selection and aggregation matrices, respectively.

We let $M_{q,t}$ be the n_q identity matrix I_{n_q} if all quarterly variables are observed at time t so that $y_{q,t} = (0 \quad \Lambda_q) Z_t$. In the remaining periods, $M_{q,t}$ is an empty matrix

such that $y_t = y_{m,t}$. More complicated observational structures can easily be accommodated in the same manner; instead of being empty or a full I_n matrix, M_t can have rows that correspond to unobserved variables omitted. This idea allows for seamlessly handling missing data for a subset of the monthly variables at the end of the sample.

The aggregation matrix Λ_q represents the assumed aggregation scheme of unobserved high-frequency latent observations $z_{q,t}$ into occasionally observed low-frequency observations $y_{q,t}$. To make the presentation simpler, we can write the bottom block of ΛZ_t as

$$\begin{pmatrix} 0 & \Lambda_q \end{pmatrix} \begin{pmatrix} z_{m,t} \\ z_{q,t} \\ z_{m,t-1} \\ z_{q,t-1} \\ \vdots \\ z_{m,t-p+1} \\ z_{q,t-p+1} \end{pmatrix} = \Lambda_{qq} \begin{pmatrix} z_{q,t} \\ z_{q,t-1} \\ \vdots \\ z_{q,t-p+1} \end{pmatrix},$$

where Λ_{qq} collects the columns of Λ_q that correspond to quarterly variables in Z_t .

Schorfheide and Song (2015), working with log-levels of the data, used the intra-quarterly average $y_{q,t}^* = \frac{1}{3}(z_{q,t}^* + z_{q,t-1}^* + z_{q,t-2}^*)$, where $y_{q,t}^*$ denotes the observed quarterly log-levels and the latent monthly log-levels. Because we use log-differenced data, we instead follow Mariano and Murasawa (2003, 2010). By taking the quarterly difference of $y_{q,t}^*$ to construct our observed growth rates, we obtain

$$\begin{aligned} y_{q,t} &= y_{q,t}^* - y_{q,t-3}^* = \frac{1}{3} \left[(z_{q,t}^* - z_{q,t-3}^*) + (z_{q,t-1}^* - z_{q,t-4}^*) + (z_{q,t-2}^* - z_{q,t-5}^*) \right] \\ &= \frac{1}{3} \left[(\Delta z_{q,t}^* + \Delta z_{q,t-1}^* + \Delta z_{q,t-2}^*) + (\Delta z_{q,t-1}^* + \Delta z_{q,t-2}^* + \Delta z_{q,t-3}^*) \right. \\ &\quad \left. + (\Delta z_{q,t-2}^* + \Delta z_{q,t-3}^* + \Delta z_{q,t-4}^*) \right], \end{aligned}$$

Finally, the expression can be written as

$$y_{q,t} = \frac{1}{3} \left[\Delta z_{q,t}^* + 2 \Delta z_{q,t-1}^* + 3 \Delta z_{q,t-2}^* + 2 \Delta z_{q,t-3}^* + \Delta z_{q,t-4}^* \right]. \quad (4)$$

Because the set of weights in (4) (6) sum to three, we define our latent variable of interest to be $z_{q,t} = 3 \Delta z_{q,t}^*$, i. e. the latent month-on-month growth rate scaled to be commensurate in scale with the quarterly level.

Equations (2) and (3) form a state-space model that can be used for estimation of the model. Schorfheide and Song (2015) suggested an efficient compact

formulation of the employed state-space model that is statistically equivalent but computationally more convenient. The compact treatment is based on the observation that the set of monthly variables included in the model are observed for all time points except for a handful at the end of the sample, known as a ragged edge (Bańbura, Giannone, and Reichlin 2011). The treatment proposed by Schorfheide and Song (2015) is to let the monthly variables enter the model as exogenous for $t = 1, \dots, T_b$, where T_b denotes the final time period where the monthly variables are all observed. By this approach, the monthly variables are excluded from the state equation. The state dimension is thereby reduced from np to $n_q(p + 1)$, which substantially improves the computational efficiency.

Let

$$\tilde{y}_t = y_t - M_t \Lambda \begin{pmatrix} \Psi d_t \\ \vdots \\ \Psi d_{t-p+1} \end{pmatrix} \quad (5)$$

denote the mean-adjusted data. The state-space model is thereafter formulated in terms of $\tilde{y}_t$ and $\tilde{z}_t = z_t - \Psi d_t$, leading to the model

$$\begin{pmatrix} \tilde{y}_{m,t} \\ \tilde{y}_{q,t} \end{pmatrix} = \begin{pmatrix} 0_{n_m \times n_q} & \Pi_{mq} \\ M_{q,t} \Lambda_q & 0_{n_q \times n_q} \end{pmatrix} \begin{pmatrix} \tilde{z}_{q,t} \\ \tilde{z}_{q,t-1} \end{pmatrix} + \Pi_{mm} \tilde{Y}_{m,t-1} + \begin{pmatrix} u_{m,t} \\ 0_{n_q \times 1} \end{pmatrix} \\ \begin{pmatrix} \tilde{z}_{q,t} \\ \tilde{z}_{q,t-1} \end{pmatrix} = \begin{pmatrix} \Pi_{qq} 0_{n_q \times n_q} \\ I_{n_q p} 0_{n_q p \times n_q} \end{pmatrix} \begin{pmatrix} \tilde{z}_{q,t-1} \\ \tilde{z}_{q,t-2} \end{pmatrix} + \Pi_{qm} \tilde{Y}_{m,t-1} + \begin{pmatrix} u_{q,t} \\ 0_{n_q p \times 1} \end{pmatrix} \quad (6)$$

where $\Pi_{i,j}$, $i, j \in \{m, q\}$ refer to the submatrices of regression parameters relating the j frequency variables to the conditional mean of the i frequency variables. The errors are the corresponding partitions of $u_t = (u'_{m,t}, u'_{q,t})'$ and are consequently correlated. Finally, $\tilde{Y}_{m,t-1}$ stacks the mean-adjusted monthly variables as $\tilde{Y}_{m,t-1} = (\tilde{y}'_{m,t-1}, \dots, \tilde{y}'_{m,t-p})'$ and $\tilde{Z}_{q,t} = (\tilde{z}'_{q,t}, \dots, \tilde{z}'_{q,t-p})'$.

The above state-space model remains valid as long as $t \leq T_b$, implying that all of the monthly series are observed. To deal with ragged edges and unbalanced monthly data for $t = T_b + 1$, we follow Ankargren and Jonéus (2020) and adaptively add the monthly series with missing data as appropriate. Contrary to Schorfheide and Song (2015), we thereby avoid use of the full companion form altogether.

2.2 Extending the Basic Steady-State Model

The standard Bayesian VAR (BVAR) with the steady-state prior typically produces good forecasts and is for this reason used by e.g. Sveriges Riksbank as one of its main forecasting models (see Iversen et al. 2016). However, recent work in the VAR literature demonstrates that allowing for more flexibility may be beneficial.

Particularly, letting the error covariance matrix in the model vary over time by incorporating stochastic volatility often improves the predictive ability as demonstrated by e.g. Clark (2011); Clark and Ravazzolo (2015); Carriero et al. (2016). Moreover, studies such as Bańbura, Giannone, and Reichlin (2010); Giannone, Lenza, and Primiceri (2015); Koop (2013) have shown that medium-sized models including 10–20 variables often outperform smaller models when forecasting. The caveat, however, when extending the size of the model under the use of the steady-state prior is that the researcher must set a prior mean and variance for the unconditional mean for each variable in the model. For key variables such as inflation, gross domestic product (GDP) growth and unemployment this task is relatively effortless, but it can be more challenging when the previous literature does not offer any guidance on reasonable prior specifications. To simplify the process of specifying the steady-state prior, Louzis (2019) developed a hierarchical prior for the steady-state prior that effectively relieves the researcher from eliciting the prior variances of the steady-state parameters. Instead, only prior means are required. Providing a sensible prior for the unconditional mean is generally much simpler than quantifying the uncertainty of one's specification. We next briefly describe the stochastic volatility and hierarchical steady-state prior specifications that we extend our basic model with.

2.2.1 Stochastic Volatility

The stochastic volatility model we employ is the CSV model of Carriero et al. (2016), which is a parsimonious and simple approach for letting the error covariance matrix in the model vary over time. The state equation describing the high-frequency VAR is under the CSV variance specification given by

$$\Pi(L)(z_t - \Psi d_t) = \sqrt{f_t} A^{-1} e_t, \quad e_t \sim N(0, I) \quad (7)$$

where A^{-1} is a lower triangular matrix and f_t is the latent univariate volatility series evolving according to

$$\log(f_t) = \phi \log(f_{t-1}) + v_t, \quad v_t \sim N(0, \sigma^2). \quad (8)$$

The log-volatility $\log(f_t)$ thus evolves as an AR(1) process without intercept with parameters (ϕ, σ^2) . The time-varying error covariance matrix implied by the preceding model is $\Sigma_t = f_t \Sigma$, where $\Sigma = A^{-1}(A^{-1})'$. Consequently, the CSV prior assumes a fixed covariance structure where the volatility factor provides a time-varying scaling of the constant error covariance Σ .

2.2.2 Hierarchical Steady-State Priors

The appealing feature of the steady-state prior is that it allows the researcher to use readily available information about long-run steady-state levels of the included variables. For the reasons discussed earlier, Louzis (2019) proposed a hierarchical steady-state prior using the normal-gamma construction used by e.g. Griffin and Brown (2010); Huber and Feldkircher (2019). The reason for such an approach is that the benefits of the steady-state prior are larger when we have accurate and relatively informative priors for the steady states. The normal-gamma prior employs a hierarchical specification that provides sufficiently heavy tails to allow for a large degree of shrinkage to the prior mean when appropriate, and more flexibility otherwise. In effect, the researcher only has to provide a prior mean for each steady-state parameter as the associated variances are instead obtained from the hyperparameters higher up in the hierarchy.

To be more precise, the hierarchical steady-state prior is based on the normal-gamma prior proposed by Griffin and Brown (2010) that employs a hierarchical specification given by

$$\begin{aligned}\psi_j | \omega_{\psi,j} &\sim N(\mu_{\psi,j}, \omega_{\psi,j}) \\ \omega_{\psi,j} &\sim G(\phi_\psi, 0.5\phi_\psi\lambda_\psi),\end{aligned}\tag{9}$$

where ϕ_ψ and λ_ψ are additional fixed hyperparameters, $G(a, b)$ denotes the gamma distribution with shape a and rate b and $\psi = \text{vec}(\Psi)$. The prior is therefore constructed using idiosyncratic, or local, hyperparameters $\omega_{\psi,j}$, which in turn depend on the two auxiliary hyperparameters ϕ_ψ and λ_ψ .

Griffin and Brown (2010) showed that the variance of the unconditional prior for ψ_j is negatively associated with λ_ψ , meaning that higher values of λ_ψ induce a larger degree of shrinkage towards the prior mean. The hyperparameter λ_ψ can therefore be interpreted as a global shrinkage parameter. At the same time, the excess kurtosis of the unconditional prior is negatively related to ϕ_ψ . Taken together, the implication is that if a tight prior (i. e. λ_ψ is high) is employed, the local shrinkage given by $\omega_{\psi,j}$ can still deviate notably from zero if ϕ_ψ is small due to the heavy tails of the unconditional prior distribution. This feature allows for a shrinkage profile that is in general tight, but loose when necessary.

2.3 Prior Distributions

We use a standard normal inverse Wishart prior for the VAR coefficients and error covariance (Π, Σ). Thus, we have a priori

$$\Sigma \sim \text{IW}(\underline{S}, \underline{\nu}), \text{vec}(\Pi') | \Sigma \sim \text{N}(\text{vec}(\underline{\Pi}'), \Sigma \otimes \underline{\Omega}_{\Pi}), \quad (10)$$

where $\Pi = (\Pi_1, \dots, \Pi_p)$. The main diagonal of the prior covariance matrix for the regression parameters, $\underline{\Omega}_{\Pi}$, is set in the Minnesota-style fashion

$$\omega_{\Pi, ii} = \frac{\lambda_1^2}{(l^2 s_r)^2} \text{ for lag } l \text{ of variable } r, i = (l-1)p + r \quad (11)$$

where λ_1 is the overall tightness and λ_2 determines the lag decay rate; the inclusion of s_r adjusts for differences in measurement scale of the variables. For a more thorough exposition of the normal inverse Wishart prior, the reader is referred to Karlsson (2013).

While Σ describes the fixed covariance structure, the time-varying volatility in the model is governed by the latent volatility f_t . For the two parameters associated with its evolution, (ϕ, σ^2) , we use a normal distribution truncated to the stationary region for ϕ , and an inverse gamma prior for σ^2 :

$$\begin{aligned} \phi &\sim \text{N}(\underline{\mu}_{\phi}, \underline{\Omega}_{\phi}; |\phi| < 1) \\ \sigma^2 &\sim \text{IG}(\underline{d} \cdot \underline{\sigma}^2, \underline{d}). \end{aligned} \quad (12)$$

As discussed in Section 2.2, the priors for the steady-state parameters are normal conditional on the local shrinkage parameters. Instead of fixing the top-level hyperparameters ϕ_{ψ} and λ_{ψ} , Huber and Feldkircher (2019) proceeded with an additional hierarchy by specifying priors for ϕ_{ψ} and λ_{ψ} . We follow their suggestion and obtain the following hierarchical prior specification for the steady-state parameters:

$$\begin{aligned} \psi_j | \omega_{\psi, j} &\sim \text{N}(\underline{\mu}_{\psi, j}, \omega_{\psi, j}) \\ \omega_{\psi, j} | \phi_{\psi}, \lambda_{\psi} &\sim \text{G}(\phi_{\psi}, 0.5\phi_{\psi}\lambda_{\psi}) \\ \phi_{\psi} &\sim \text{Exp}(1) \\ \lambda_{\psi} &\sim \text{G}(c_0, c_1). \end{aligned} \quad (13)$$

2.4 Posterior Sampling

To estimate the model and produce forecasts, we employ Markov Chain Monte Carlo (MCMC). The MCMC algorithm consists of multiple Gibbs sampling steps, which we describe next. We relegate some of the details to Appendix A.

2.4.1 Sampling the Latent Monthly Variables

To sample from the posterior distribution of the latent monthly variables, $p(Z|\Pi, \Sigma, \psi, f, Y, d)$, we use a simulation smoother along the lines of Durbin and Koopman (2012). To increase the computational efficiency, we implement it using the compact formulation for the balanced part of the sample as suggested by Schorfheide and Song (2015). For the unbalanced ragged edge, we instead leverage the adaptive procedure developed by Ankargren and Jonéus (2020). The simulation smoothing step is conducted based on the mean-adjusted data $\tilde{y}_t$ to produce a draw of $\tilde{z}_t$. We thereafter construct the unadjusted high-frequency series by adding the deterministic component $z_t = \tilde{z}_t + \Psi d_t$.

2.4.2 Sampling the Regression and Covariance Parameters

Given Z , ψ and f , the VAR can be transformed into a homoskedastic VAR without intercept based on $\tilde{z}_t = z_t/\sqrt{f_t}$ and $\tilde{Z}_{t-1} = (\tilde{z}'_{t-1}, \dots, \tilde{z}'_{t-p})'$:

$$\tilde{z}_t = \Pi \tilde{Z}_{t-1} + A^{-1} e_t. \quad (14)$$

By standard results (Kadiyala and Karlsson 1993, 1997), the conditional posterior distribution is also normal inverse Wishart. It is thereby possible to sample from the marginal posterior of Σ followed by the full conditional posterior of Π :

$$\begin{aligned} \Sigma | \tilde{Z} &\sim \text{IW}(\tilde{S}, T + \nu) \\ \text{vec}(\Pi') | \Sigma, \tilde{Z} &\sim \text{N}(\text{vec}(\bar{\Pi}'), \Sigma \otimes \bar{\Omega}_\Pi). \end{aligned} \quad (15)$$

The posterior moments are standard given the transformation of the model and presented in Appendix A. A draw can efficiently be made from the posterior of Π by reverting to its matrix-normal form:

$$\Pi = \text{chol}(\tilde{\Omega}_\Pi^{-1})' \left[\text{chol}(\tilde{\Omega}_\Pi^{-1}) \left(\Pi \underline{\Omega}_\Pi + \sum_{t=1}^T \tilde{z}_t \tilde{z}'_{t-1} \right) + \Xi \times \text{chol}(\Sigma)' \right], \quad (16)$$

where Ξ is an $n \times np$ matrix of numbers independently drawn from the standard normal distribution, chol is the lower triangular Cholesky decomposition and the operation $A \backslash B$ means to solve the linear system $AX = B$ for X . Because the Cholesky factor is triangular, the linear systems can be solved more efficiently using forward and back substitution.

2.4.3 Sampling the Steady-State Parameters

Prior to sampling the steady-state parameters, the associated hyperparameters are drawn from their respective conditional posterior distributions. The conditional posterior of the global shrinkage parameter λ_ψ is gamma distributed and given by

$$\lambda_\psi \sim \text{G}\left(nm\phi_\psi + c_0, 0.5\phi_\psi \sum_{j=1}^{nm} \omega_{\psi,j} + c_1\right). \quad (17)$$

The conditional posterior of ϕ_ψ is proportional to

$$\begin{aligned} p(\phi_\psi | \omega_\psi, \lambda_\psi) &\propto g(\phi_\psi | \omega_\psi, \lambda_\psi) \\ &= \prod_{j=1}^{nm} \frac{(0.5\lambda_\psi \phi_\psi)^{\phi_\psi}}{\Gamma(\phi_\psi)} \omega_{\psi,j}^{\phi_\psi-1} \exp\{-0.5\lambda_\psi \phi_\psi \omega_{\mu,j} - \phi_\psi\} \end{aligned} \quad (18)$$

and permits no representation in terms of a standard distribution. As suggested by Huber and Feldkircher (2019); Louzis (2019) we employ a random walk Metropolis-Hastings step in order to sample from the posterior distribution. The random walk operates on the log-scale and the proposal is given by

$$\log(\phi_\psi^*) = \log(\phi_\psi^{(i-1)}) + sz, \quad z \sim N(0, 1), \quad (19)$$

where s is a scaling factor. The proposed value ϕ_ψ^* is accepted with probability

$$r = \min\left\{1, \frac{g(\phi_\psi^* | \omega_\psi, \lambda_\psi)}{g(\phi_\psi | \omega_\psi, \lambda_\psi)} \frac{\phi_\psi^*}{\phi_\psi^{(i-1)}}\right\}, \quad (20)$$

where the second ratio accounts for the asymmetric proposal distribution.

Given the hyperparameters, the local shrinkage parameters $\omega_{\psi,j}$ can be sampled. The conditional posterior distribution is the generalized inverse Gaussian distribution

$$\omega_{\psi,j} | \lambda_\psi, \phi_\psi, \psi_j \sim \text{GIG}\left(\phi_\psi - 0.5, \lambda_\psi \phi_\psi, (\psi_j - \underline{\mu}_{\psi,j})^2\right) \quad j = 1, \dots, nm, \quad (21)$$

where if $y \sim \text{GIG}(a, b, c)$ then $p(y; a, b, c) \propto y^{a-1} \exp\{0.5(by + c/y)\}$. The prior covariance matrix for ψ , i. e. $\underline{\Omega}_\psi$, can thereafter be constructed as the diagonal matrix with main diagonal given by $(\omega_{\psi,1}, \dots, \omega_{\psi,nm})$.

Next, by dividing both sides of the model (7) by $\sqrt{f_t}$ we obtain a homoskedastic model given by

$$\Pi(L) \left(\frac{z_t}{\sqrt{f_t}} - \Psi \frac{d_t}{\sqrt{f_t}} \right) = A^{-1} e_t. \quad (22)$$

The posterior moments provided by Villani (2009) therefore apply directly for the preceding transformation of the model. Let

$$\begin{aligned} \check{z}_t &= \Pi(L) z_t / \sqrt{f_t}, \quad \check{d}_t' = \left(\frac{d_t'}{\sqrt{f_t}} - \frac{d_{t-1}'}{\sqrt{f_t}} \cdots \frac{d_{t-p}'}{\sqrt{f_t}} \right) \\ U &= \begin{pmatrix} I_{nm} \\ I_m \otimes \Pi_1 \\ \vdots \\ I_m \otimes \Pi_p \end{pmatrix}. \end{aligned} \quad (23)$$

The posterior distribution of ψ is

$$\psi | \check{Z}, \check{d}, \omega_\psi \sim N(\check{\mu}_\psi, \Omega_\psi) \quad (24)$$

with posterior moments

$$\begin{aligned} \bar{\Omega}_\psi^{-1} &= \underline{\Omega}_\psi^{-1} + U' \left[\left(\sum_{t=1}^T \check{d}_t \check{d}_t' \right) \otimes \Sigma^{-1} \right] U \\ \bar{\mu}_\psi &= \bar{\Omega}_\psi \left[U' \text{vec} \left(\Sigma^{-1} \sum_{t=1}^T \check{z}_t \check{d}_t' \right) + \underline{\Omega}_\psi^{-1} \underline{\psi} \right]. \end{aligned} \quad (25)$$

2.4.4 Sampling the Latent Volatility

Conditional on the other parameters in the model, we can obtain

$$\check{z}_t = A \Pi(L) (z_t - \Psi d_t) = \sqrt{f_t} e_t. \quad (26)$$

Squaring and taking the logarithm of the elements of $\check{z}_t$ yields

$$\log(\check{z}_{i,t}^2) = \log(f_t) + \log(e_{i,t}^2), \quad i = 1, \dots, n, \quad (27)$$

where $\check{z}_{i,t}$ is the i th element of $\check{z}_{i,t}$ with a similar logic for $e_{i,t}$. Coupling the preceding equation with the transition Eq. (8) defines a linear but non-normal state-space model. Kim, Shephard, and Chib (1998) proposed a sampling strategy that introduces auxiliary mixture indicators $r_{t,i}$ so that the model conditional on these indicators is normal. We use the refined ten-state mixture by Omori et al. (2007) together with the algorithm discussed by McCausland, Miller, and Pelletier (2011) and as implemented by Kastner and Frühwirth-Schnatter (2014); Kastner (2016) to sample from the posterior distribution of the latent volatility series.

The posteriors of the parameters of the volatility process are standard given f . The posterior distribution of ϕ is a truncated normal distribution whereas the posterior distribution of σ^2 is inverse gamma. We proceed by sampling (ϕ, σ^2) first, the mixture indicators $r_{t,i}$ next and, finally, the latent volatility series in order to target the correct posterior distribution as discussed by Del Negro and Primiceri (2015).

3 Data and Implementation Details

In this section, we provide information about the data used and some details regarding the implementations.

3.1 Data

Our dataset consists of 13 key macroeconomic variables for the United States. The dataset we use largely parallels that of Carriero et al. (2016); Louzis (2019) with the exception that we use consumer price index (CPI) inflation as the sole measure of

Table 1: Summary of the real-time dataset.

Series	Transformation	Frequency	Real time	$\mu_{\psi,j}$	$\sqrt{\omega_{\psi,j}}$
Nonfarm payrolls ^a	1200 $\Delta\ln$	Monthly	Yes	3	0.5
Hours ^{ab}	X13, 1200 $\Delta\ln$	Monthly	≥ 2011	3	0.5
Unemployment rate ^a	None	Monthly	Yes	6	1
Federal funds rate ^a	None	Monthly	Yes	5	0.7
Bond spread ^b	Monthly ave.	Monthly	Yes	1	1
Stock market index ^c	1200 $\Delta\ln$	Monthly	No	0	2
Personal consumption ^a	1200 $\Delta\ln$	Monthly	Yes	3	0.7
Industrial production ^a	1200 $\Delta\ln$	Monthly	Yes	3	0.7
Capacity utilization ^a	None	Monthly	Yes	80	0.7
CPI inflation ^a	1200 $\Delta\ln$	Monthly	Yes	2	0.5
Nonresidential inv. ^a	400 $\Delta\ln$	Quarterly	Yes	3	1.5
Residential inv. ^a	400 $\Delta\ln$	Quarterly	Yes	3	1.5
GDP growth ^a	400 $\Delta\ln$	Quarterly	Yes	2	0.5

Sources: ^aALFRED, Federal Reserve Bank of St. Louis.

^bFRED, Federal Reserve Bank of St. Louis.

^cFRED-MD, McCracken and Ng (2016).

Notes: 1. Real-time data for Hours is available in ALFRED from 2011 and onwards; data from FRED is used prior to 2011. Hours is seasonally adjusted using X-13ARIMA-SEATS using the seasonal package in R (Sax and Edelbuettel 2018).

2. A list of the IDs of the variables is available in Appendix C.

Figure 1: Publication Delays. The color of each box represents the number of months since the last available observation. The delay is computed for the 10th day of the corresponding month; a zero-month delay implies that the observation for the preceding month is available.

inflation. The data consist of 10 monthly and three quarterly variables and ranges over the period 1980M01–2018M12. Most of the included variables are available with real-time vintages in the ALFRED database. For variables not available in ALFRED, we turn to FRED and FRED-MD (McCracken and Ng 2016). A summary of the data is provided in Table 1.

We follow Louzis (2019) and transform the raw series to growth levels. For our monthly variables, we use month-on-month growth rates, whereas the three quarterly variables are computed as quarter-on-quarter rates. All growth rates are annualized. The final two columns of Table 1, $\mu_{\psi,j}$ and $\sqrt{\omega_{\psi,j}}$, display the prior means and prior standard deviations of the unconditional means of the variables. The values are drawn from Louzis (2019), but are also in line with e.g. Clark (2011); Österholm (2012).

We use real-time data where available throughout the forecasting exercise. To obtain a realistic pattern of available observations, we first consider the information set available on the 10th day of every month. Figure 1 displays the publication pattern during 2005–2018 and shows the number of months that has passed since the last available publication.

Figure 1 shows a characteristic pattern for real-time forecasting of macroeconomic data. Data for financial and select real and nominal variables are already available for the previous month, whereas the previous month’s outcomes for some of the monthly variables are unknown. The pattern of availability displayed shows that consumption and inflation are available with a one-month delay at every month except for a handful of occasions. Similarly, non-farm

employment, hours, unemployment and the federal funds rate are typically available with a zero-month delay with the exception of a few months. In the final dataset that we use in our forecasting exercise, we make adjustments to the publication delays in order to obtain a more uniform dataset. The adjustments change the publication structure in the vintages so that the aforementioned variables have the same delay in all vintages, i. e. consumption and inflation are always observed with a delay of one month, whereas non-farm employment, hours, unemployment and the federal funds rate are always observed without any delay. Consequently, at every month that we make our forecasts observations are available for the preceding month for six of the monthly variables, whereas four still lack data.

3.2 Implementation Details

The mixed-frequency models that we estimate use $p = 12$ lags following e.g. Bańbura et al. (2010). The overall tightness in the prior distribution for the regression parameters is set to $\lambda_1 = 0.2$ and the lag decay used is $\lambda_2 = 1$. We use 15,000 draws in the MCMC procedure and discard the first 5000.

For the hierarchical steady-state prior, we let $c_0 = c_1 = 0.01$ in line with Huber and Feldkircher (2019); Louzis (2019). To set the scale of the proposal distribution for ϕ_ψ , we employ the adaptive scaling procedure discussed by Roberts and Rosenthal (2009). We use a batch size of 100 and check every 100 iterations if the fraction of acceptances within the most recent batch exceeds 0.44. If it does, we increase s by $\delta(k) = \min(0.01, k^{-1/2})$, where k denotes the batch number. If the fraction of acceptances was less than 0.44, s is instead decreased by $\delta(k)$.

For the parameters of the log-volatility process, we let the prior mean and standard deviation for ϕ be $\underline{\mu}_\phi = 0.9$ and $\sqrt{\underline{\Omega}_\phi} = 0.1$, respectively. The prior mean and degrees of freedom of σ^2 are $\underline{g}^2 = 0.01$ and $\underline{d} = 4$.

4 Empirical Application: Real-Time Forecasting of Key US Variables

In this section, we assess the forecasting ability of the model that we propose. The assessment is carried out by studying the out-of-sample predictive accuracy of the model based on the real-time dataset for the US that was discussed in Section 3.

4.1 Forecasting Setup

The quarterly steady-state Bayesian VAR model has been used in several previous studies, see for example Adolfson et al. (2007); Österholm (2008); Villani (2009); Clark (2011); Ankargren, Bjellerup, and Shahnazarian (2017). The model is employed both for policy purposes and for forecasting and is implemented in the Matlab toolbox BEAR developed at the European Central Bank (Dieppe, Legrand, and van Roye 2016). Our empirical application targets this audience, and our main interest lies in seeing whether the components we add to the model—mixed frequencies, stochastic volatility and hierarchical steady states—improve upon the benchmark model of Villani (2009) estimated on single-frequency data. The forecasting results are also compared to models using Minnesota-style normal inverse Wishart priors, i. e. without use of the steady-state component. A summary of the models that we include in the forecast evaluation is presented in Table 2.

The benchmark model is the steady-state model estimated on single-frequency data. Depending on whether it serves as benchmark for quarterly or monthly variables, we include either the full set of variables (aggregated to the quarterly frequency) or the 10 monthly variables. The quarterly VAR uses $p = 4$, whereas for the monthly VAR $p = 12$.

We use a recursive forecasting scheme to evaluate the forecasting performance of the considered models. Beginning in January 2005, we estimate the models and

Table 2: List of models.

Model	Description
Benchmark	Single-frequency model with the steady-state prior and a normal inverse Wishart prior for (Π, Σ) , constant error covariance. Includes all 13 variables aggregated to the quarterly frequency or the 10 monthly variables depending on context.
Minn-IW	Normal-inverse Wishart prior, constant error covariance
Minn-CSV	Normal-inverse Wishart prior with common stochastic volatility
SS-IW	Steady-state prior with a normal inverse Wishart prior for (Π, Σ) , constant error covariance
SS-CSV	Steady-state prior with a normal inverse Wishart prior for (Π, Σ) with common stochastic volatility
SSNG-IW	Hierarchical normal-gamma steady-state prior with a normal inverse Wishart prior for (Π, Σ) , constant error covariance
SSNG-CSV	Hierarchical normal-gamma steady-state prior with a normal inverse Wishart prior for (Π, Σ) with common stochastic volatility

make forecasts and then recursively add months to the set of data used for estimation. The benchmark models use the balanced data, whereas the mixed-frequency models automatically handle the ragged edges.

The forecasting ability of the models is evaluated with respect to both point and density forecasts. For point forecasts, we consider the root mean squared errors (RMSE). For density forecasts, we compute univariate and multivariate log predictive density scores (LPDS). We do so by fitting a normal density to the draws from the predictive distribution following e.g. Adolfson et al. (2007); Carriero et al. (2015). That is, we compute

$$LPDS_{h,t}^{(m,s)} = n_s \ln(2\pi) + \ln \left| V_{t+h|t}^{(m,s)} \right| + \left(y_{t+h}^{(s)} - \tilde{y}_{t+h|t}^{(m,s)} \right)' \left(V_{t+h|t}^{(m,s)} \right)^{-1} \left(y_{t+h}^{(s)} - \tilde{y}_{t+h|t}^{(m,s)} \right), \quad (28)$$

where m denotes the model, s denotes the set of variables the LPDS is computed for, n_s is the dimension of s , h is the forecast horizon, and $\tilde{y}$ and V are the mean and covariance of the draws from the relevant predictive distribution. For fixed (m, s, h) , we compute the summary LPDS by averaging over the evaluation period. We calculate the LPDS jointly but separately for the monthly and quarterly variables, and univariately for all variables.

What vintage the forecasts should be evaluated with respect to is an important question when using real-time data. Two alternatives are commonly employed in the literature. The first, as used by e.g. Romer and Romer (2000) and Clark (2011), is to use the second available vintage. This choice can be justified by acknowledging that revisions that occur after longer periods of time may be unforeseeable and more structural in nature by relating to e.g. definitions, methods of measurement, etc. The second available estimate therefore provides a less noisy estimate than the initial available value, yet is produced in the same environment as the forecaster is active. The second common approach for evaluation, as followed by e.g. Schorfheide and Song (2015), is to use the most recent vintage. For whatever reason revisions may have taken place, the currently available data provide the best estimates of e.g. inflation and output in previous years. We follow the latter approach and use the most recent vintage for evaluating the forecasts, but for transparency provide the main results of the evaluation using the second available vintage in Appendix D.

4.2 In-Sample Estimation

As a preliminary analysis, we begin by estimating the mixed-frequency VAR model using the steady-state (SS) and steady-state normal-gamma (SSNG) priors to see whether the obtained steady-state posteriors differ. Because the long-term forecasts are largely determined by the steady-state posterior, seeing

Figure 2: Steady States. Kernel density estimates of the posterior distribution of the steady states (unconditional means).

whether differences are present is of direct importance for forecasts beyond the immediate short term. Figure 2 displays kernel density estimates of the posterior distributions from the mixed-frequency model with CSV. As a point of reference, the figure includes the prior distribution detailed in Table 1.

As expected, the posteriors in Figure 2 are for the most part similar. The modes of the posteriors are close to perfectly aligned for variables such as bond spread, inflation, residential investment and GDP. For others—e.g., hours, the federal funds rate and industrial production—the SSNG posteriors deviate more from both the priors and the SS posteriors.

While the steady states are of central importance for the levels of the forecasts, the precision thereof is highly influenced by the CSV factor. Figure 3 displays the mean of $\sqrt{f}_t$ together with 90% bands for the SS-CSV and SSNG-CSV models.

Figure 3: Volatility. The lines display the posterior means of $\sqrt{f_t}$ and the bands how the 90 % posterior intervals.

Figure 3 shows that there is little difference between the estimated volatility factors in the two steady-state models. Peaks of volatility are aligned and reach the same levels, while the level of the factor in the SSNG model is slightly higher in normal times. Both display the entrance into the Great Moderation in the beginning of the 1980s with heightened volatility again around the recent financial crisis. The interpretation of the level of the factor is that the time-invariant elements in the error covariance matrix Σ have been scaled by f_t , which roughly amounts to an amplification by a factor of 4–6 during the recent financial crisis and a compression of around 0.5–0.75 in recent years. This feature has a direct effect on the width of the predictive distribution.

4.3 Forecast Evaluation

In this section, we present the main results of the forecast evaluation. For space considerations, the presentation includes the results from the joint evaluations as well as the univariate results for the three quarterly variables and the three monthly variables that are typically of primary interest: the inflation, federal funds and unemployment rates.

4.3.1 Joint forecasting results

Table 3 presents the results from the LPDS computed jointly. We compute the LPDS separately for the set of quarterly and monthly variables, respectively. The forecast

Table 3: Relative joint log predictive density scores.

Model	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$	$h = 8$
Relative joint LPDS, Quarterly									
Minn-IW	-0.11	-0.33 [*]	0.01	0.24	0.33	0.52	0.52	0.45	0.41
SS-IW	-0.17	-0.48 [*]	-0.22	-0.09	-0.08	0.05	0.02	-0.07	-0.13
SSNG-IW	-0.14	-0.42 [*]	-0.15	-0.02	-0.04	0.07	0.02	-0.09	-0.12
Minn-CSV	-0.36	-1.01 ^{**}	-0.76 [*]	-0.57 [*]	-0.49	-0.22	-0.05	-0.06	-0.05
SS-CSV	-0.42	-1.07 ^{**}	-0.89[*]	-0.77[*]	-0.74[*]	-0.52	-0.39	-0.39	-0.39
SSNG-CSV	-0.43	-1.07 ^{**}	-0.86 [*]	-0.73 [*]	-0.69 [*]	-0.44	-0.32	-0.36	-0.36
Relative joint LPDS, Monthly									
Minn-IW	-1.74 ^{**}	-1.49 ^{**}	-1.21 ^{**}	-1.04 [*]	-1.14 ^{**}	-1.03 ^{**}	-1.02 ^{**}	-0.88 ^{**}	
SS-IW	-1.85 ^{**}	-1.44 ^{**}	-1.12 ^{**}	-1.04 ^{**}	-1.14 ^{**}	-1.04 ^{**}	-1.05 ^{**}	-0.95 ^{**}	
SSNG-IW	-1.83 ^{**}	-1.47 ^{**}	-1.19 ^{**}	-1.03 [*]	-1.22 ^{**}	-1.12 ^{**}	-1.08 ^{**}	-0.95 ^{**}	
Minn-CSV	-1.96 [*]	-2.93 [*]	-3.01 [*]	-3.03 [*]	-2.98 [*]	-2.77 [*]	-2.62 [*]	-2.29 [*]	
SS-CSV	-2.17[*]	-3.01 [*]	-3.00 [*]	-3.07 [*]	-3.13[*]	-3.01[*]	-2.98[*]	-2.65[*]	
SSNG-CSV	-2.07 [*]	-3.07[*]	-3.03[*]	-3.09[*]	-3.13 [*]	-2.97 [*]	-2.86 [*]	-2.53 [*]	

Note: The forecast horizons h refer to quarters and months, respectively, for the two sets of variables. The scores in the table display the score of the model in the first column minus the score of the benchmark model, whereby negative entries indicate that the mixed-frequency model is superior. Bold entries show the minimum in each column. The benchmark model for the quarterly set of variables is a VAR(4) including all 13 variables aggregated to the quarterly frequency. For the monthly LPDS, the benchmark model is a VAR(12) including the 10 monthly variables. For both cases, the steady-state prior with a constant error covariance matrix is used. Two stars (**) indicate that the Diebold-Mariano test of equal predictive ability is significant at the 1 percent level, whereas a single star (*) indicates significance at the 10 percent level. The test employs the modifications proposed by Harvey, Leybourne, and Newbold 1997.

horizons h in the table correspond to the frequency of the respective set of variables.

Across all horizons and sets of variables, SS-CSV and SSNG-CSV dominate with only one exception in which Minn-CSV does slightly better than SS-CSV. For the quarterly sets of variables, SS-CSV outperforms the other models for $h > 0$ with the SSNG-CSV model ranking first for the nowcast. Minn-CSV ranks higher than the constant volatility models for the initial horizons, but for the long-term forecasts the added value of the steady-state prior outweighs the improvements obtained from stochastic volatilities. However, given a model, stochastic volatility appears to be useful as it improves the joint forecasting performance of quarterly variables across the board when comparing the constant volatility models to their heteroskedastic counterparts. Within the two groups of models with constant and stochastic volatility, we see that the steady-state models forecast better than Minn-IW and Minn-CSV, respectively, throughout all horizons. Therefore, the table shows that steady-state information and flexible modeling of the volatility structure help to improve the quarterly forecasts.

For the performance of the monthly forecasts, the picture is largely the same. The three models with stochastic volatility outperform the constant models for all horizons and SSNG-CSV produces the most accurate density forecasts for $h = 2, 3, 4$. For the remaining horizons, SS-CSV picks up the lead. Among the constant volatility models, the ranking is no longer uniform across horizons.

With respect to the joint log predictive scores, we can therefore conclude the following. First, there are gains in utilizing prior information on the steady states. Second, further improvements can be obtained by allowing for stochastic volatility. Third, with a handful of exceptions for the quarterly forecasts made by Minn-IW and SS-IW, the relative LPDS is negative throughout, indicating that the mixed-frequency models produce better density forecasts than the single-frequency benchmarks. The three points are in line with the previous literature and can be seen as a synthesis of the conclusions made by Villani (2009); Clark (2011); Schorfheide and Song (2015); Carriero et al. (2016); Louzis (2019).

Table 4: GDP: forecast evaluation.

Model	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$	$h = 8$
Relative LPDS (model in first column – benchmark)									
Minn-IW	–0.23 [*]	–0.09	0.11	0.18	0.14	0.18	0.14	0.11	0.08
SS-IW	–0.23 [*]	–0.13	0.05	0.06	0.00	0.01	–0.02	–0.05	–0.10[*]
SSNG-IW	–0.23 [*]	–0.10	0.10	0.15	0.10	0.11	0.04	0.01	–0.04
Minn-CSV	–0.27	–0.24 [*]	0.01	0.15	0.16	0.20	0.19	0.16	0.09
SS-CSV	–0.27	–0.26[*]	–0.02	0.10	0.08	0.12	0.10	0.06	0.00
SSNG-CSV	–0.27	–0.25 [*]	0.00	0.13	0.14	0.17	0.14	0.10	0.03
Relative RMSE (model in first column/benchmark)									
Minn-IW	0.90 ^{""}	0.95 [*]	1.05	1.11	1.07	1.10	1.07	1.06	1.04
SS-IW	0.90^{""}	0.94[*]	1.03	1.05	1.00	1.01	0.98	0.96[*]	0.95[*]
SSNG-IW	0.90 ^{""}	0.94 [*]	1.05	1.10	1.05	1.07	1.03	1.01	0.99
Minn-CSV	0.92	0.95	1.03	1.13	1.10	1.12	1.08	1.06	1.03
SS-CSV	0.92	0.94	1.01	1.08	1.02	1.04	1.01	0.97	0.95
SSNG-CSV	0.92	0.94	1.02	1.12	1.08	1.09	1.05	1.01	0.99

Note: The forecast horizon h denotes quarters. Negative LPDS entries indicate that the mixed-frequency model is superior in terms of density forecasting and values of the RMSE below 1 indicate better point forecasts. Bold entries show the minimum per column. The benchmark model is a VAR(4) including all 13 variables aggregated to the quarterly frequency using the steady-state prior with a constant error covariance matrix. Two stars (") indicate that the Diebold–Mariano test of equal predictive ability is significant at the 1 percent level, whereas a single star (*) indicates significance at the 10 percent level. The test employs the modifications proposed by Harvey et al. (1997).

Table 5: Residential investment: forecast evaluation.

Model	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$	$h = 8$
Relative LPDS (model in first column – benchmark)									
Minn-IW	0.07	0.00	0.22	0.16	0.14	0.12	0.09	0.14	0.18
SS-IW	–0.01	–0.08	0.09	0.02	–0.01	–0.03	–0.08	–0.04	–0.01
SSNG-IW	0.03	–0.03	0.12	0.00	–0.09	–0.18	–0.25*	–0.22*	–0.19*
Minn-CSV	–0.10	–0.49*	–0.38*	–0.42*	–0.44*	–0.38*	–0.36*	–0.32*	–0.28*
SS-CSV	–0.17	–0.54*	–0.46*	–0.53*	–0.56**	–0.53*	–0.53*	–0.48*	–0.46*
SSNG-CSV	–0.17	–0.54*	–0.43*	–0.51*	–0.55**	–0.53*	–0.56*	–0.52*	–0.51*
Relative RMSE (model in first column/benchmark)									
Minn-IW	0.92**	0.96**	1.03	1.02	1.01	1.03	1.03	1.06	1.06
SS-IW	0.90**	0.92**	0.99	0.97	0.97	0.98	0.99*	1.01	1.01
SSNG-IW	0.92**	0.94**	1.00	0.98	0.96	0.97	0.96	0.98	0.98
Minn-CSV	0.88**	0.90**	0.96	0.94	0.94*	0.96	0.95*	0.98	1.00
SS-CSV	0.87**	0.90**	0.95	0.91*	0.92*	0.93	0.92*	0.95	0.96
SSNG-CSV	0.88**	0.90**	0.95	0.93	0.92*	0.94	0.92*	0.95	0.96

Note: The forecast horizon h denotes quarters. Negative LPDS entries indicate that the mixed-frequency model is superior in terms of density forecasting and values of the RMSE below 1 indicate better point forecasts. Bold entries show the minimum per column. The benchmark model is a VAR(4) including all 13 variables aggregated to the quarterly frequency using the steady-state prior with a constant error covariance matrix. Two stars (**) indicate that the Diebold-Mariano test of equal predictive ability is significant at the 1 percent level, whereas a single star (*) indicates significance at the 10 percent level. The test employs the modifications proposed by Harvey et al. (1997).

4.3.2 Quarterly Univariate Forecasting Results

Tables 4–6 present the univariate LPDS and RMSE for the three quarterly variables GDP, Residential investment and Non-residential investment.

Starting with GDP, a somewhat different pattern than what was seen for the joint LPDS emerges. For both evaluation metrics, SS-IW is generally the better forecaster beyond the short term and is only outperformed by CSV models at the first three horizons and in terms of density forecasts. Table 4 shows that the mixed-frequency models do better than the quarterly benchmark for the immediate short term when either nowcasting the current quarter or forecasting the next quarter. Beyond the first quarter forecast, the quarterly model generally produces more accurate forecasts. A similar result is found by Schorfheide and Song (2015). Use of the steady-state prior results in more accurate forecasts at every horizon, but whether or not a hierarchical prior formulation and stochastic volatility provide improvements varies. The homoskedastic steady-state models outperform the Minn-IW model at all horizons, and the stochastic volatility steady-state models consistently forecast GDP growth more accurately than Minn-CSV.

Table 6: Non-residential investment: forecast evaluation.

Model	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$	$h = 8$
Relative LPDS (model in first column – benchmark)									
Minn-IW	–0.09*	–0.38**	–0.16*	0.01	0.13	0.24	0.28	0.21	0.16
SS-IW	–0.10*	–0.42**	–0.22*	–0.09	–0.02	0.08	0.11	0.03	–0.03
SSNG-IW	–0.09*	–0.41**	–0.20*	–0.05	0.05	0.16	0.20	0.11	0.05
Minn-CSV	–0.12	–0.45**	–0.24	–0.06	0.06	0.21	0.32	0.31	0.26
SS-CSV	–0.11	–0.47**	–0.32*	–0.17	–0.06	0.06	0.16	0.13	0.09
SSNG-CSV	–0.12*	–0.46**	–0.29*	–0.14	–0.02	0.13	0.25	0.22	0.18
Relative RMSE (model in first column/benchmark)									
Minn-IW	0.97	0.83**	0.93	0.98	1.02	1.09	1.12	1.11	1.08
SS-IW	0.95	0.82**	0.92*	0.95	0.98	1.03	1.04	1.00	0.98*
SSNG-IW	0.97	0.83**	0.91*	0.95	0.99	1.05	1.08	1.06	1.03
Minn-CSV	0.93	0.83**	0.93	1.01	1.07	1.15	1.23	1.22	1.19
SS-CSV	0.94	0.81**	0.88*	0.94	0.99	1.04	1.08	1.06	1.03
SSNG-CSV	0.93	0.81**	0.90	0.97	1.03	1.10	1.17	1.16	1.13

Note: The forecast horizon h denotes quarters. Negative LPDS entries indicate that the mixed-frequency model is superior in terms of density forecasting and values of the RMSE below 1 indicate better point forecasts. Bold entries show the minimum per column. The benchmark model is a VAR(4) including all 13 variables aggregated to the quarterly frequency using the steady-state prior with a constant error covariance matrix. Two stars (**) indicate that the Diebold-Mariano test of equal predictive ability is significant at the 1 percent level, whereas a single star (*) indicates significance at the 10 percent level. The test employs the modifications proposed by Harvey et al. (1997).

For residential investment, Table 5 presents forecasting results that more closely resemble the joint results. SS-CSV and SSNG-CSV dominate for all horizons, although the difference with respect to Minn-CSV is occasionally small, particularly for the point forecasts. Nevertheless, both steady-state models with stochastic volatility perform well with better scores than all other models for every horizon and with respect to both point and density forecasts.

Finally, Table 6 shows the forecast evaluation for Non-residential investment. The pattern displayed in Table 6 is a mix of the patterns in Tables 4–5. For the nowcast, Minn-CSV provides better forecasts than the others, whereas SS-CSV generally does well and ranks first for horizons 1–5 with respect to the density forecasts. The utility of the steady-state prior is clear from Table 6: while Minn-CSV and Minn-IW start out well, the performance deteriorates more rapidly with h than what is manifested by the other models employing information about the steady states. We can again see that both SS-CSV and SSNG-CSV dominate Minn-CSV for all $h > 0$.

4.3.3 Monthly Univariate Forecasting Results

Moving to the monthly variables, Table 7 presents the forecast evaluation for inflation. The results indicate that there is little to gain from using the mixed-

Table 7: Inflation: forecast evaluation.

Model	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$	$h = 8$
Relative LPDS (model in first column – benchmark)									
Minn-IW	0.07	–0.09	–0.02	–0.01	0.01	–0.02	–0.01	–0.01	–0.02
SS-IW	0.06	–0.10	0.00	0.00	0.03	–0.00	–0.00	0.00	–0.01
SSNG-IW	0.05	–0.10	–0.01	–0.01	0.04	–0.01	–0.04	–0.02	–0.03 [*]
Minn-CSV	–0.22	–0.42	–0.31	–0.33	–0.34	–0.30	–0.29	–0.32	–0.30
SS-CSV	–0.22	–0.43	–0.32	–0.31	–0.33	–0.32	–0.31	–0.36	–0.32
SSNG-CSV	–0.22	–0.43	–0.31	–0.32	–0.33	–0.29	–0.33	–0.36	–0.34
Relative RMSE (model in first column/benchmark)									
Minn-IW	1.02	0.98	0.99	0.99	1.00	1.00	1.00	1.00	1.00
SS-IW	1.02	0.98	1.00	1.00	1.01	1.01	1.00	1.00	1.00
SSNG-IW	1.02	0.98	0.99	0.99	1.00	1.00	1.00	1.00	1.00
Minn-CSV	1.05	0.99	0.97	0.96	0.97	0.97	0.97 [*]	0.98 [*]	0.99
SS-CSV	1.04	0.98	0.97	0.97	0.97 [*]	0.97	0.97 [*]	0.97[*]	0.98
SSNG-CSV	1.04	0.99	0.97	0.96	0.97	0.97	0.97[*]	0.98 [*]	0.98

Note: The forecast horizon h denotes months. Negative LPDS entries indicate that the mixed-frequency model is superior in terms of density forecasting and values of the RMSE below 1 indicate better point forecasts. Bold entries show the minimum per column. The benchmark model is a VAR(12) including the 10 monthly variables using the steady-state prior with a constant error covariance matrix. Two stars (**) indicate that the Diebold–Mariano test of equal predictive ability is significant at the 1 percent level, whereas a single star (*) indicates significance at the 10 percent level. The test employs the modifications proposed by Harvey et al. (1997).

frequency VAR for forecasting monthly inflation as compared to a monthly VAR. The relative RMSE is close to unity and few of the Diebold–Mariano tests of equal predictive ability indicate any difference between the benchmark and the mixed-frequency models. Use of stochastic volatility improves the density forecasts somewhat, whereas the quality of the point forecasts deteriorates.

Next, the evaluation of the forecasts of the federal funds rate is displayed in Table 8. In contrast to the results for inflation, we here find large benefits from using the mixed-frequency models for forecasting the monthly federal funds rate. All three models with stochastic volatility do well with respect to both density and point forecasts, but the steady-state models have a small edge across most horizons. Contrasting these results with the results for inflation in Table 7, we now find larger improvements from using stochastic volatility. We interpret this result as an indication that the federal funds rate has been more volatile than the inflation rate relative to constant historical levels of volatility. In addition, the improved accuracy of the forecasts obtained from the mixed-frequency models

Table 8: Federal funds rate: forecast evaluation.

Model	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$	$h = 8$
Relative LPDS (model in first column – benchmark)								
Minn-IW	–0.91**	–0.50**	–0.32*	–0.24*	–0.20	–0.18	–0.17	–0.17
SS-IW	–0.93**	–0.53**	–0.35**	–0.27**	–0.23*	–0.21*	–0.21*	–0.21*
SSNG-IW	–0.92**	–0.52**	–0.35**	–0.27**	–0.23**	–0.21*	–0.20*	–0.20*
Minn-CSV	–1.45**	–1.04**	–0.80**	–0.63**	–0.52**	–0.44**	–0.38**	–0.34*
SS-CSV	–1.47**	–1.06**	–0.82**	–0.64**	–0.53**	–0.45**	–0.39**	–0.35**
SSNG-CSV	–1.46**	–1.05**	–0.81**	–0.64**	–0.52**	–0.44**	–0.37**	–0.34**
Relative RMSE (model in first column/benchmark)								
Minn-IW	0.58**	0.75*	0.85	0.90	0.93	0.94	0.94	0.94
SS-IW	0.56**	0.72*	0.81*	0.86*	0.89*	0.90*	0.90*	0.90*
SSNG-IW	0.57**	0.72*	0.81*	0.86*	0.89*	0.90*	0.91*	0.90**
Minn-CSV	0.53**	0.68*	0.76*	0.81	0.84	0.86	0.87	0.87
SS-CSV	0.51**	0.65*	0.73*	0.78*	0.82*	0.83*	0.85*	0.85
SSNG-CSV	0.52**	0.67*	0.75*	0.80*	0.83*	0.84*	0.85	0.86

Note: The forecast horizon h denotes months. Negative LPDS entries indicate that the mixed-frequency model is superior in terms of density forecasting and values of the RMSE below 1 indicate better point forecasts. Bold entries show the minimum per column. The benchmark model is a VAR(12) including the 10 monthly variables using the steady-state prior with a constant error covariance matrix. Two stars (**) indicate that the Diebold-Mariano test of equal predictive ability is significant at the 1 percent level, whereas a single star (*) indicates significance at the 10 percent level. The test employs the modifications proposed by Harvey et al. (1997).

highlight the importance of utilizing all real-time information that is available. As explained in Section 4.1, the mixed-frequency VAR automatically handles ragged edges of the data, whereas the single-frequency benchmark is estimated on the balanced data set. For some variables, e.g., the federal funds rate, this evidently makes a difference. Schorfheide and Song (2015) reached the same conclusion when forecasting quarterly growth rates of monthly variables. A likely explanation as to why this appears to matter more for the federal funds rate is because of its high persistence.

The final series we evaluate univariate forecasts for is the unemployment rate. The results are presented in Table 9. The table reveals that mixed-frequency models are useful also for forecasting unemployment. The results mirror those from the federal funds rate, displaying the importance of the ragged edge used by the mixed-frequency models. SS-IW appears to be the better forecaster in terms of point forecasts, whereas SS-CSV provides more accurate density forecasts for all horizons. Thus, adding stochastic volatility does not improve point forecasts of the unemployment rate, but the density forecasts exhibit enhancements. We interpret these results as indications that the stochastic volatility alternatives better

Table 9: Unemployment: forecast evaluation.

Model	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$	$h = 8$
Relative LPDS (model in first column – benchmark)								
Minn-IW	–0.43	–0.34	–0.25	–0.28	–0.34	–0.30	–0.32 [*]	–0.29
SS-IW	–0.46	–0.39	–0.31	–0.35	–0.41 [*]	–0.38	–0.40	–0.37
SSNG-IW	–0.43	–0.33	–0.21	–0.24	–0.29	–0.25	–0.26	–0.22
Minn-CSV	–0.48	–0.48	–0.51	–0.61	–0.74	–0.77	–0.82	–0.84
SS-CSV	–0.49	–0.50	–0.52	–0.63	–0.77	–0.80	–0.87	–0.89
SSNG-CSV	–0.47	–0.47	–0.49	–0.59	–0.72	–0.74	–0.80	–0.82
Relative RMSE (model in first column/benchmark)								
Minn-IW	0.78 ^{**}	0.83 [*]	0.87 [*]	0.87	0.86	0.88	0.89	0.90
SS-IW	0.77^{**}	0.81[*]	0.84	0.84	0.84	0.86	0.87	0.88
SSNG-IW	0.79 ^{**}	0.84 [*]	0.87 [*]	0.87	0.87	0.89	0.89	0.91
Minn-CSV	0.82 ^{**}	0.87 [*]	0.90	0.90	0.89	0.89	0.90	0.91
SS-CSV	0.81 ^{**}	0.86 [*]	0.89	0.89	0.88	0.89	0.89	0.90
SSNG-CSV	0.83 [*]	0.88 [*]	0.91	0.91	0.90	0.91	0.91	0.92

Note: The forecast horizon h denotes months. Negative LPDS entries indicate that the mixed-frequency model is superior in terms of density forecasting and values of the RMSE below 1 indicate better point forecasts. Bold entries show the minimum per column. The benchmark model is a VAR(12) including the 10 monthly variables using the steady-state prior with a constant error covariance matrix. Two stars (**) indicate that the Diebold-Mariano test of equal predictive ability is significant at the 1 percent level, whereas a single star (*) indicates significance at the 10 percent level. The test employs the modifications proposed by Harvey et al. (1997).

characterize the evolution of the unemployment rate. The result is not surprising—Clark (2011), for example, found that the effects of stochastic volatility were more pronounced for density than for point forecasts.

4.3.4 Comparison with a Model with Equation-Specific Volatilities

We next compare the results with a more complex model to see whether the CSV we have assumed is restrictive or serves as an efficient approximation. The extended model, labeled SSNG-SV, is characterized by

$$\begin{aligned}
 \Pi(L)(z_t - \Psi d_t) &= A^{-1} D_t^{1/2} e_t, \quad e_t \sim N(0, I) \\
 D_t &= \text{diag}(f_{1,t} \dots f_{n,t}) \\
 \log(f_{i,t}) &= \mu_{i,f} + \phi_i(\log(f_{i,t-1}) - \mu_{i,f}) + v_{i,t}, \quad v_{i,t} \sim N(0, \sigma_i^2),
 \end{aligned} \quad (29)$$

and A is lower triangular as before, but now with ones along the main diagonal. For more details, see Appendix B.

The results in Table 10 show that using CSV leads to both gains and losses in terms of predictive ability. The SSNG-SV model improves the density forecasts for

Table 10: Forecast evaluation: SSNG-SV and SSNG-CSV.

Variable	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$	$h = 8$
Relative LPDS (SSNG-SV – SSNG-CSV)									
GDP	–0.07	0.07	0.10	0.07	0.03	0.06	0.06	0.08	0.08
Non-residential investment	0.05	0.09	0.00	–0.00	–0.03	–0.05	–0.08	–0.08	–0.09
Residential investment	–0.43 ^{**}	0.32 [*]	0.44 [*]	0.34 [*]	0.27	0.14	0.15	0.13	0.09
Inflation	–0.15 [*]	–0.11	–0.17	–0.17	–0.08	–0.04	0.10	0.10	0.06
Federal funds rate		–1.33 ^{**}	–1.14 [*]	–1.00	–0.83	–0.71	–0.65	–0.58	–0.53
Unemployment		–0.01	0.04	0.13	0.20	0.26	0.32	0.36	0.39
Relative RMSE (SSNG-SV/SSNG-CSV)									
GDP	0.95	0.99	1.05 [*]	1.06	1.05	1.07 [*]	1.09 [*]	1.11 [*]	1.12 [*]
Non-residential investment	1.00	1.00	0.98	0.99	0.98	0.98	0.97	0.97	0.97
Residential investment	1.06 [*]	1.06 ^{**}	1.06 [*]	1.04 [*]	1.00	0.98	1.00	0.99	0.97
Inflation	0.99	0.99	1.00	1.01	1.01	1.01	1.01 [*]	1.01 [*]	1.00
Federal funds rate		0.86	0.87	0.88	0.90	0.92	0.94	0.96	0.99
Unemployment		0.94 [*]	0.93	0.94	0.94	0.95	0.97	0.98	0.98

Note: The forecast horizon h represents quarters for GDP, Non-residential investment and Residential investment, and months otherwise. Negative LPDS entries indicate that the SSNG-SV model is superior in terms of density forecasting and values of the RMSE below 1 indicate better point forecasts. Two stars (**) indicate that the Diebold–Mariano test of equal predictive ability is significant at the 1 percent level, whereas a single star (*) indicates significance at the 10 percent level. The test employs the modifications proposed by Harvey et al. (1997).

residential investment and inflation at the $h = 0$ horizon, while the gain for GDP forecasts is negligible and for non-residential investment even negative. However, for horizons $h > 0$ the model with CSV produces better density forecasts for residential investment. For the federal funds rate, use of a separate volatility factor improves the density forecast substantially. The results for the point forecasts are generally in line with those for the density forecasts, although the Diebold–Mariano test no longer indicates a difference between the two models' predictive abilities with respect to the federal funds rate. When focusing on the point forecasts, the model with CSV is no longer inferior forecasting residential investment at the $h = 0$ horizon; instead, the Diebold–Mariano test signals rejection in favor of its superiority.

The results thus neither strongly support nor strongly oppose the use of a model with CSV. Table 8 already illustrates the gain from using time-varying volatility for forecasting the federal funds rate, and Table 10 shows how a volatility model that is unique for the federal funds rate equation improves these forecasts even further. Given that the CSV model sacrifices some flexibility for improved computational speed, it is natural that some variables, whose volatility patterns differ relatively more from the average pattern captured by the common volatility

factor, would benefit from an idiosyncratic volatility factor. But, with the exception of the federal funds rate, the differences in predictive ability displayed in Table 10 are either small or in favor of the model with CSV.

4.3.5 Forecasts Evaluated against the Second Vintage

To ensure that our results are not primarily driven by our choice of data to evaluate the forecasts against, Appendix D presents the same tables as shown in the main text but with the evaluations carried out with the second available vintages. Qualitatively, the results remain. For the forecasts of GDP, the gains obtained by using mixed-frequency data are larger when the forecasts are evaluated against the second vintage. Occasional changes in rankings among the models occur across variables, but for the most part the rankings remain unaltered and the conclusions made so far are intact irrespective of the choice of evaluation vintage.

5 Conclusion

We present a VAR model that is a synthesis of recent important contributions. Our model incorporates three main features. First, the model allows for mixed-frequency data by use of a state-space formulation. We deal with the particular mixed-frequency case involving monthly and quarterly data and solve the frequency mismatch problem by postulating a monthly VAR with missing values similar to the work by Schorfheide and Song (2015). Second, we include prior beliefs about the steady states, or unconditional means, of the variables in the model by means of the steady-state prior developed by Villani (2009). We also employ the hierarchical formulation of the prior proposed by Louzis (2019), whose advantage is that it is only necessary to specify prior means of the steady-state parameters while the prior variances are, in turn, equipped with hyperpriors. Third, to allow for an error covariance matrix that varies over time we include as the final component the CSV model presented by Carriero et al. (2016).

We estimate our model and competing alternatives using US data including 10 monthly and three quarterly variables. The results show that the forecasts are generally improved by adding the three components to the benchmark VAR model. Using mixed instead of single frequencies of the data generally does not produce worse forecasts, and usually performs better. Including prior information about the steady states generally outperforms the corresponding alternatives that lack this information. The hierarchical steady-state prior is appealing as it allows for shrinkage to the prior means of the steady states, and is generally on par with or

better than the standard steady-state prior. Finally, we find that CSV mostly improves the accuracy of the forecasts as the models including heteroskedasticity generally outperform the models with constant volatility.

Appendix

A Posterior Moments

Regression and Covariance Parameters

The moments of the posterior distributions for the regression and covariance parameters are:

$$\begin{aligned}
 \hat{S} &= \underline{S} + S + (\underline{\Pi} - \hat{\Pi}) \left[\underline{\Omega}_{\Pi} + \left(\sum_{t=1}^T \tilde{Z}_{t-1} \tilde{Z}'_{t-1} \right)^{-1} \right]^{-1} (\underline{\Pi} - \hat{\Pi})', \\
 \hat{\Pi} &= \sum_{t=1}^T \tilde{z}_t \tilde{Z}'_{t-1} \left(\sum_{t=1}^T \tilde{Z}_{t-1} \tilde{Z}'_{t-1} \right)^{-1} \\
 S &= \sum_{t=1}^T (\tilde{z}_t - \hat{\Pi} \tilde{Z}_{t-1}) (\tilde{z}_t - \hat{\Pi} \tilde{Z}_{t-1})', \\
 \overline{\Omega}_{\Pi}^{-1} &= \underline{\Omega}_{\Pi}^{-1} + \sum_{t=1}^T \tilde{Z}_{t-1} \tilde{Z}'_{t-1} \\
 \overline{\Pi} &= \overline{\Omega}_{\Pi} \left(\underline{\Pi} \underline{\Omega}_{\Pi} + \sum_{t=1}^T \tilde{z}_t \tilde{Z}'_{t-1} \right),
 \end{aligned} \tag{30}$$

where $\tilde{z}_t = (z_t - \Psi d_t) / \sqrt{f_t}$ and $\tilde{Z}_t = (\tilde{z}'_{t-1} \quad \dots \quad \tilde{z}'_{t-p})'$.

Latent Volatility

Let $h_t = \log(f_t)$. The conditional posterior distribution of ϕ is

$$\begin{aligned}
 \phi | h, \sigma^2 &\sim N(\bar{\mu}_{\phi}, \bar{\Omega}_{\phi}; |\phi| < 1) \\
 \bar{\mu}_{\phi} &= \bar{\Omega}_{\phi} \left(\frac{\sum_{t=1}^T h_{t-1} h_t}{\sigma^2} + \frac{\mu_{\phi}}{\underline{\Omega}_{\phi}} \right) \\
 \bar{\Omega}_{\phi}^{-1} &= \left(\underline{\Omega}_{\phi}^{-1} + \frac{\sum_{t=1}^T h_{t-1}^2}{\sigma^2} \right).
 \end{aligned} \tag{31}$$

The conditional posterior distribution of σ^2 is

$$\begin{aligned} \sigma^2 | h, \phi &\sim \text{IG}(\tilde{d}, \tilde{\sigma}^2) \\ \tilde{d} &= \underline{d} + T \\ \tilde{\sigma}^2 &= \sum_{t=1}^T (h_t - \phi h_{t-1})^2 + \underline{d} \cdot \underline{\sigma}^2. \end{aligned} \tag{32}$$

B Details on Model with Extended Stochastic Volatility Specification

We now use an independent normal prior for the regression parameters with a prior variance structure given by

$$\sqrt{\text{Var}(\pi_{i,j}^{(l)})} = \frac{\lambda_1 \lambda_2}{l^{\lambda_3}} \frac{s_i}{s_j}, \tag{33}$$

where $\pi_{i,j}^{(l)}$ is element (i, j) of Π_l . We use a standard hyperparameter specification with $\lambda_1 = 0.2$, $\lambda_2 = 0.5$ and $\lambda_3 = 1$. For A , we follow Cogley and Sargent (2005) and use row-by-row normal priors for the lower triangular part of the matrix where the prior distributions are independent $N(0, 10)$. To speed up computations, we employ the triangularization algorithm proposed by Carriero, Clark, and Marcelino (2019). The stochastic volatilities are sampled efficiently using ancillarity-sufficiency interweaving as described by Kastner and Frühwirth-Schnatter (2014) and implemented by the `stochvol` package (Kastner 2016). As priors, we use $\mu_i \sim N(0, 1000)$, $0.5(\phi_i + 1) \sim \text{Beta}(5, 1.5)$ and $\sigma_i^2 \sim \chi^2(1)$.

Table 11: Source and ID of series used.

Series	Source	ID
Nonfarm payrolls	ALFRED	PAYEMS
Hours	FRED/ALFRED	CEU0500000034
Unemployment rate	ALFRED	UNRATE
Federal funds rate	ALFRED	FEDFUNDS
Bond spread	ALFRED	T10YFF
Stock market index	FRED-MD	S&P500
Personal consumption	ALFRED	PCE
Industrial production	ALFRED	INDPRO
Capacity utilization	ALFRED	TCU
CPI inflation	ALFRED	CPIAUCSL
Nonresidential inv.	ALFRED	PNFI
Residential inv.	ALFRED	PRFI
GDP growth	ALFRED	GDPC1

Table 12: GDP: forecast evaluation (second vintage).

Model	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$	$h = 8$
Relative LPDS (model in first column – benchmark)									
Minn-IW	–0.34**	–0.15	0.03	0.12	0.14	0.13	0.09	0.09	0.09
SS-IW	–0.31*	–0.16*	0.01	0.04	0.04	–0.00	–0.05*	–0.05**	–0.05
SSNG-IW	–0.33**	–0.15	0.03	0.11	0.12	0.08	0.02	0.01	0.00
Minn-CSV	–0.28	–0.27**	–0.06	0.10	0.17	0.19	0.17	0.17	0.16
SS-CSV	–0.24	–0.28**	–0.05	0.08	0.14	0.14	0.10	0.10	0.10
SSNG-CSV	–0.26	–0.27**	–0.05	0.10	0.17	0.19	0.15	0.14	0.11
Relative RMSE (model in first column/benchmark)									
Minn-IW	0.83**	0.92**	1.03	1.09	1.11	1.10	1.07	1.07	1.06
SS-IW	0.84**	0.92**	1.03	1.06	1.04	1.00	0.97	0.96	0.96
SSNG-IW	0.83**	0.92**	1.03	1.08	1.10	1.06	1.03	1.02	1.01
Minn-CSV	0.85**	0.90*	0.98	1.09	1.14	1.12	1.08	1.08	1.06
SS-CSV	0.87**	0.90**	0.97	1.05	1.07	1.04	1.00	0.99	0.99
SSNG-CSV	0.85**	0.90**	0.97	1.08	1.12	1.10	1.05	1.04	1.02

Note: The forecast horizon h denotes quarters. Negative LPDS entries indicate that the mixed-frequency model is superior in terms of density forecasting and values of the RMSE below 1 indicate better point forecasts. Bold entries show the minimum per column. The benchmark model is a VAR(4) including all 13 variables aggregated to the quarterly frequency using the steady-state prior with a constant error covariance matrix. Two stars (**) indicate that the Diebold-Mariano test of equal predictive ability is significant at the 1 percent level, whereas a single star (*) indicates significance at the 10 percent level. The test employs the modifications proposed by Harvey et al. (1997).

Table 13: Residential investment: forecast evaluation (second vintage).

Model	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$	$h = 8$
Relative LPDS (model in first column – benchmark)									
Minn-IW	0.01	–0.11*	0.14	0.14	0.14	0.11	0.08	0.14	0.18
SS-IW	–0.05	–0.18*	0.03	0.01	–0.01	–0.05	–0.09	–0.06	–0.01
SSNG-IW	–0.00	–0.12*	0.07	0.01	–0.06	–0.16	–0.22*	–0.20*	–0.15
Minn-CSV	–0.18	–0.53*	–0.38*	–0.36*	–0.37*	–0.31*	–0.28*	–0.23	–0.17
SS-CSV	–0.24	–0.56*	–0.44*	–0.45*	–0.46*	–0.44*	–0.43*	–0.38*	–0.33
SSNG-CSV	–0.23	–0.56*	–0.42*	–0.43*	–0.46*	–0.43*	–0.45*	–0.41*	–0.37
Relative RMSE (model in first column/benchmark)									
Minn-IW	0.93*	0.96*	1.03	1.02	1.02	1.03	1.04	1.06	1.06
SS-IW	0.90*	0.92**	0.99	0.98	0.97	0.98	0.98*	1.00	1.01
SSNG-IW	0.92*	0.95*	1.01	0.98	0.97	0.97	0.96	0.98	0.98
Minn-CSV	0.89*	0.91**	0.96	0.94	0.95*	0.96	0.96	1.00	1.01
SS-CSV	0.88**	0.91**	0.95	0.91*	0.92*	0.93	0.93*	0.96	0.97
SSNG-CSV	0.88*	0.91**	0.96	0.93	0.93*	0.94	0.93	0.97	0.97

Note: The forecast horizon h denotes quarters. Negative LPDS entries indicate that the mixed-frequency model is superior in terms of density forecasting and values of the RMSE below 1 indicate better point forecasts. Bold entries show the minimum per column. The benchmark model is a VAR(4) including all 13 variables aggregated to the quarterly frequency using the steady-state prior with a constant error covariance matrix. Two stars (**) indicate that the Diebold-Mariano test of equal predictive ability is significant at the 1 percent level, whereas a single star (*) indicates significance at the 10 percent level. The test employs the modifications proposed by Harvey et al. (1997).

Table 14: Non-residential investment: forecast evaluation (second vintage).

Model	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$	$h = 8$
Relative LPDS (model in first column – benchmark)									
Minn-IW	-0.07	-0.52*	-0.08	0.05	0.19	0.34	0.38	0.29	0.23
SS-IW	-0.11*	-0.62*	-0.21*	-0.11	-0.03	0.09	0.12	0.02	-0.04
SSNG-IW	-0.09*	-0.57*	-0.15*	-0.02	0.10	0.24	0.27	0.16	0.10
Minn-CSV	-0.16*	-0.74*	-0.37	-0.17	-0.04	0.18	0.33	0.29	0.25
SS-CSV	-0.19*	-0.81*	-0.49	-0.33	-0.22	-0.04	0.08	0.03	0.01
SSNG-CSV	-0.17*	-0.78*	-0.44	-0.27	-0.14	0.09	0.23	0.17	0.14
Relative RMSE (model in first column/benchmark)									
Minn-IW	0.99	0.84*	0.96	0.99	1.05	1.11	1.14	1.12	1.11
SS-IW	0.96	0.81*	0.94*	0.95	0.98	1.03	1.03	0.99	0.99
SSNG-IW	0.98	0.83*	0.95*	0.97	1.01	1.08	1.09	1.07	1.06
Minn-CSV	0.96	0.84*	0.96	1.01	1.09	1.18	1.24	1.23	1.21
SS-CSV	0.94	0.81*	0.92*	0.95	1.01	1.06	1.10	1.07	1.07
SSNG-CSV	0.95	0.82*	0.94	0.98	1.05	1.14	1.19	1.17	1.16

Note: The forecast horizon h denotes quarters. Negative LPDS entries indicate that the mixed-frequency model is superior in terms of density forecasting and values of the RMSE below 1 indicate better point forecasts. Bold entries show the minimum per column. The benchmark model is a VAR(4) including all 13 variables aggregated to the quarterly frequency using the steady-state prior with a constant error covariance matrix. Two stars (**) indicate that the Diebold-Mariano test of equal predictive ability is significant at the 1 percent level, whereas a single star (*) indicates significance at the 10 percent level. The test employs the modifications proposed by Harvey et al. (1997).

Table 15: Inflation: forecast evaluation (second vintage).

Model	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$	$h = 8$
Relative LPDS (model in first column – benchmark)									
Minn-IW	0.06	-0.07	-0.02	-0.01	0.01	-0.02	-0.02	-0.02	-0.03
SS-IW	0.06	-0.09	-0.00	-0.00	0.03	-0.00	-0.01	-0.00	-0.01
SSNG-IW	0.05	-0.08	-0.01	-0.01	0.03	-0.01	-0.04	-0.02*	-0.04
Minn-CSV	-0.03	-0.32	-0.27	-0.31	-0.32	-0.31	-0.31	-0.36	-0.34
SS-CSV	-0.05	-0.34	-0.28	-0.29	-0.33	-0.33	-0.34	-0.40	-0.37
SSNG-CSV	-0.04	-0.34	-0.27	-0.30	-0.32	-0.30	-0.35	-0.40	-0.38
Relative RMSE (model in first column/benchmark)									
Minn-IW	1.02	0.98	0.99	0.99	1.00	1.00	1.00	1.00	1.00
SS-IW	1.01	0.98	0.99	1.00	1.01	1.01	1.00	1.00	1.00
SSNG-IW	1.02	0.98	0.99	0.99	1.00	1.00	1.00	1.00	1.00
Minn-CSV	1.04	0.99	0.98	0.97	0.97	0.97	0.98*	0.98*	0.98
SS-CSV	1.04	0.99	0.98	0.97	0.97*	0.98	0.98*	0.98*	0.98
SSNG-CSV	1.04	0.99	0.98	0.97	0.97	0.97	0.98*	0.98*	0.98

Note: The forecast horizon h denotes months. Negative LPDS entries indicate that the mixed-frequency model is superior in terms of density forecasting and values of the RMSE below 1 indicate better point forecasts. Bold entries show the minimum per column. The benchmark model is a VAR(12) including the 10 monthly variables using the steady-state prior with a constant error covariance matrix. Two stars (**) indicate that the Diebold-Mariano test of equal predictive ability is significant at the 1 percent level, whereas a single star (*) indicates significance at the 10 percent level. The test employs the modifications proposed by Harvey et al. (1997).

Table 16: Unemployment: forecast evaluation (second vintage).

Model	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$	$h = 8$
Relative LPDS (model in first column – benchmark)								
Minn-IW	–0.57	–0.43	–0.28	–0.29	–0.35*	–0.32	–0.32*	–0.29
SS-IW	–0.60	–0.48	–0.35	–0.36	–0.43*	–0.40	–0.41	–0.38
SSNG-IW	–0.56	–0.41	–0.25	–0.25	–0.30	–0.26	–0.26	–0.21
Minn-CSV	–0.54	–0.48	–0.46	–0.54	–0.68	–0.72	–0.78	–0.80
SS-CSV	–0.56	–0.50	–0.48	–0.56	–0.71	–0.75	–0.82	–0.84
SSNG-CSV	–0.53	–0.46	–0.44	–0.51	–0.65	–0.69	–0.76	–0.78
Relative RMSE (model in first column/benchmark)								
Minn-IW	0.74**	0.81**	0.86*	0.87	0.86	0.87	0.88	0.90
SS-IW	0.73**	0.79*	0.83*	0.84	0.84	0.85	0.86	0.88
SSNG-IW	0.75**	0.81**	0.87*	0.87	0.87	0.88	0.89	0.91
Minn-CSV	0.79**	0.85*	0.90*	0.90	0.89	0.89	0.90	0.91
SS-CSV	0.78**	0.84*	0.89*	0.89	0.88	0.88	0.89	0.90
SSNG-CSV	0.79**	0.86*	0.91	0.92	0.90	0.91	0.91	0.92

Note: The forecast horizon h denotes months. Negative LPDS entries indicate that the mixed-frequency model is superior in terms of density forecasting and values of the RMSE below 1 indicate better point forecasts. Bold entries show the minimum per column. The benchmark model is a VAR(12) including the 10 monthly variables using the steady-state prior with a constant error covariance matrix. Two stars (**) indicate that the Diebold-Mariano test of equal predictive ability is significant at the 1 percent level, whereas a single star (*) indicates significance at the 10 percent level. The test employs the modifications proposed by Harvey et al. (1997).

C Data Sources

The IDs of the series used and their sources are shown in Table 11.

D Forecast Evaluation Tables (Second Vintage)

The tables in the main text present the results of the forecast evaluation when evaluated with respect to the most recent vintage. The tables in this Appendix (Table 12–16) present the same evaluations but conducted with respect to the second available vintage. Because the federal funds rate is not revised, the second vintage is the same as the most recent vintage. Therefore, the results when evaluating the forecasts against the second vintage are identical to Table 8 and therefore not reproduced again here.

E Additional Results

This Appendix presents forecast evaluation tables (Table 17–23) for the variables not discussed in the main text.

Table 17: Hours: forecast evaluation.

Model	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$	$h = 8$
Relative LPDS (model in first column – benchmark)								
Minn-IW	-0.11	-0.12	-0.05	-0.10	-0.04*	-0.11	-0.06	-0.06
SS-IW	-0.08	-0.11	-0.03	-0.10	-0.04*	-0.11	-0.06	-0.06
SSNG-IW	-0.11	-0.12	-0.04	-0.12*	-0.04	-0.11	-0.07	-0.05
Minn-CSV	0.08	-0.10	-0.01	-0.17	-0.12	-0.08	0.01	0.04
SS-CSV	0.13	-0.08	-0.01	-0.16	-0.14	-0.13	-0.04	0.01
SSNG-CSV	0.08	-0.09	-0.01	-0.19	-0.15	-0.12	-0.02	0.02
Relative RMSE (model in first column/benchmark)								
Minn-IW	0.96*	0.95*	0.98	0.96*	0.99	0.98	0.98	0.98
SS-IW	0.96*	0.96*	0.98	0.96*	0.99*	0.98	0.98	0.98
SSNG-IW	0.95*	0.95*	0.98	0.96*	0.99*	0.98	0.98	0.98
Minn-CSV	0.97	0.96	0.97	0.95*	0.97	0.99	0.99	0.99
SS-CSV	0.98	0.97	0.98	0.95*	0.98*	0.98	0.98	0.99
SSNG-CSV	0.97	0.96	0.97	0.95*	0.97	0.98	0.98	0.99

Note: The forecast horizon h denotes months. Negative LPDS entries indicate that the mixed-frequency model is superior in terms of density forecasting and values of the RMSE below 1 indicate better point forecasts. Bold entries show the minimum per column. The benchmark model is a VAR(12) including the 10 monthly variables using the steady-state prior with a constant error covariance matrix. Two stars (**) indicate that the Diebold-Mariano test of equal predictive ability is significant at the 1 percent level, whereas a single star (*) indicates significance at the 10 percent level. The test employs the modifications proposed by Harvey et al. (1997).

Table 18: Industrial production: forecast evaluation.

Model	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$	$h = 8$
Relative LPDS (model in first column – benchmark)									
Minn-IW	-0.06	-0.04	-0.16*	-0.16*	-0.07*	-0.06	-0.12	-0.04	-0.08
SS-IW	-0.05	-0.02	-0.15*	-0.16*	-0.08*	-0.06	-0.12	-0.06	-0.10
SSNG-IW	-0.03	-0.03	-0.15*	-0.15*	-0.06*	-0.06	-0.10	-0.06	-0.06
Minn-CSV	-0.36	-0.27	-0.36	-0.35*	-0.32	-0.28	-0.30	-0.23	-0.23
SS-CSV	-0.37	-0.25	-0.35	-0.34*	-0.33	-0.30	-0.33	-0.27	-0.27
SSNG-CSV	-0.38*	-0.27	-0.36	-0.35*	-0.34	-0.31	-0.33	-0.27	-0.26
Relative RMSE (model in first column/benchmark)									
Minn-IW	0.92*	0.98	0.95*	0.95*	0.98*	0.99	0.97	0.99	0.99
SS-IW	0.91*	0.98	0.95*	0.95*	0.98	0.99	0.97	0.99	0.98
SSNG-IW	0.92*	0.98	0.96*	0.95*	0.98*	0.99	0.97	0.99	0.99
Minn-CSV	0.91**	0.97	0.94*	0.94*	0.98*	0.98*	0.98	1.01	1.01
SS-CSV	0.91**	0.97	0.95*	0.94*	0.98	0.99*	0.98	1.00	1.00
SSNG-CSV	0.91**	0.97	0.94*	0.94*	0.98*	0.98*	0.98	1.00	1.00*

Note: The forecast horizon h denotes months. Negative LPDS entries indicate that the mixed-frequency model is superior in terms of density forecasting and values of the RMSE below 1 indicate better point forecasts. Bold entries show the minimum per column. The benchmark model is a VAR(12) including the 10 monthly variables using the steady-state prior with a constant error covariance matrix. Two stars (**) indicate that the Diebold-Mariano test of equal predictive ability is significant at the 1 percent level, whereas a single star (*) indicates significance at the 10 percent level. The test employs the modifications proposed by Harvey et al. (1997).

Table 19: Non-farm employment: forecast evaluation.

Model	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$	$h = 8$
Relative LPDS (model in first column – benchmark)								
Minn-IW	-0.13**	-0.24**	-0.21*	-0.18*	-0.18	-0.18	-0.15	-0.18
SS-IW	-0.12**	-0.24**	-0.22*	-0.20*	-0.20	-0.21	-0.19	-0.23
SSNG-IW	-0.14**	-0.24**	-0.21*	-0.19*	-0.18	-0.20	-0.15	-0.18
Minn-CSV	-0.39**	-0.41**	-0.40*	-0.36*	-0.31	-0.31	-0.27	-0.29
SS-CSV	-0.38**	-0.41**	-0.40*	-0.35*	-0.31	-0.31	-0.27	-0.30
SSNG-CSV	-0.39**	-0.42**	-0.40*	-0.35*	-0.31	-0.31	-0.26	-0.29
Relative RMSE (model in first column/benchmark)								
Minn-IW	0.93*	0.88*	0.91*	0.92	0.93	0.93	0.95	0.94
SS-IW	0.93*	0.88*	0.90*	0.91*	0.92	0.92	0.94	0.92
SSNG-IW	0.93*	0.88*	0.90*	0.92*	0.93	0.93	0.95	0.94
Minn-CSV	0.89**	0.87*	0.88*	0.90*	0.91	0.91	0.93	0.93
SS-CSV	0.90**	0.87*	0.88*	0.90*	0.91	0.91	0.93	0.92
SSNG-CSV	0.89**	0.87*	0.88*	0.90*	0.91	0.91	0.93	0.93

Note: The forecast horizon h denotes months. Negative LPDS entries indicate that the mixed-frequency model is superior in terms of density forecasting and values of the RMSE below 1 indicate better point forecasts. Bold entries show the minimum per column. The benchmark model is a VAR(12) including the 10 monthly variables using the steady-state prior with a constant error covariance matrix. Two stars (**) indicate that the Diebold-Mariano test of equal predictive ability is significant at the 1 percent level, whereas a single star (*) indicates significance at the 10 percent level. The test employs the modifications proposed by Harvey et al. (1997).

Table 20: Consumption: forecast evaluation.

Model	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$	$h = 8$
Relative LPDS (model in first column – benchmark)									
Minn-IW	-0.05*	-0.03	-0.00	0.00	0.00	-0.00	0.00	0.01	0.00
SS-IW	-0.07*	-0.04	-0.00	0.01	0.01	0.00	0.00	0.01	-0.00
SSNG-IW	-0.06*	-0.04	-0.01	-0.00	0.00	0.00	0.00	0.00	-0.00
Minn-CSV	-0.23*	-0.15*	-0.04	0.03	0.08	0.12	0.15	0.18	0.19
SS-CSV	-0.23*	-0.15*	-0.05	0.04	0.09	0.13	0.16	0.19	0.20
SSNG-CSV	-0.22*	-0.14*	-0.04	0.03	0.09	0.13	0.17	0.19	0.21
Relative RMSE (model in first column/benchmark)									
Minn-IW	0.98*	0.98	1.01	1.01	1.00	1.01	1.01	1.01	1.00
SS-IW	0.96*	0.97	1.01	1.01	1.01	1.01	1.00	1.00	0.99
SSNG-IW	0.98*	0.98	1.00	1.00	1.00	1.01	1.00	1.00	0.99*
Minn-CSV	1.01	0.97	1.03	0.99	1.00	1.00	1.00	1.02	1.01
SS-CSV	0.98	0.96	1.02	0.99	0.99	0.99	0.98	0.99	0.99
SSNG-CSV	1.01	0.97	1.02	0.99	1.00	0.99	1.00	1.01	1.00

Note: The forecast horizon h denotes months. Negative LPDS entries indicate that the mixed-frequency model is superior in terms of density forecasting and values of the RMSE below 1 indicate better point forecasts. Bold entries show the minimum per column. The benchmark model is a VAR(12) including the 10 monthly variables using the steady-state prior with a constant error covariance matrix. Two stars (**) indicate that the Diebold-Mariano test of equal predictive ability is significant at the 1 percent level, whereas a single star (*) indicates significance at the 10 percent level. The test employs the modifications proposed by Harvey et al. (1997).

Table 21: S&P500: forecast evaluation.

Model	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$	$h = 8$
Relative LPDS (model in first column – benchmark)								
Minn-IW	−0.10	−0.06 [*]	−0.03	−0.06	−0.07	−0.04	−0.05	−0.05
SS-IW	−0.09	−0.01	0.03	−0.01	−0.01	0.02	0.01	0.01
SSNG-IW	−0.10	−0.04 [*]	−0.01	−0.05	−0.06	−0.02	−0.03 [*]	−0.02
Minn-CSV	−0.20	−0.16	−0.11	−0.08	−0.07	−0.04	−0.05	−0.03
SS-CSV	−0.19	−0.15	−0.08	−0.06	−0.06	−0.02	−0.02	−0.00
SSNG-CSV	−0.21	−0.16	−0.11	−0.07	−0.08	−0.04	−0.03	−0.01
Relative RMSE (model in first column/benchmark)								
Minn-IW	0.94	0.98	0.98	0.97	0.96	0.98	0.98	0.98
SS-IW	0.95	1.00	1.02	1.00	1.00	1.02	1.01	1.01
SSNG-IW	0.94	0.99	0.99	0.98	0.97	0.99	0.99	0.99
Minn-CSV	0.96	0.98	0.99	0.97	0.97	0.99	0.98	0.98
SS-CSV	0.97	0.99	1.01	0.99	0.99	1.00	1.00	0.99
SSNG-CSV	0.96	0.98	0.99	0.98	0.97	0.99	0.99	0.99

Note: The forecast horizon h denotes months. Negative LPDS entries indicate that the mixed-frequency model is superior in terms of density forecasting and values of the RMSE below 1 indicate better point forecasts. Bold entries show the minimum per column. The benchmark model is a VAR(12) including the 10 monthly variables using the steady-state prior with a constant error covariance matrix. Two stars (**) indicate that the Diebold-Mariano test of equal predictive ability is significant at the 1 percent level, whereas a single star (*) indicates significance at the 10 percent level. The test employs the modifications proposed by Harvey et al. (1997).

Table 22: Bond spread: forecast evaluation.

Model	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$	$h = 8$
Relative LPDS (model in first column – benchmark)								
Minn-IW	−0.82 ^{**}	−0.36 ^{**}	−0.18 ^{**}	−0.10	−0.08	−0.08	−0.08	−0.08
SS-IW	−0.83 ^{**}	−0.37 ^{**}	−0.19 ^{**}	−0.11 [*]	−0.08	−0.08	−0.06	−0.05
SSNG-IW	−0.82 ^{**}	−0.37 ^{**}	−0.19 ^{**}	−0.12 [*]	−0.09	−0.10	−0.09	−0.09
Minn-CSV	−1.06^{**}	−0.58^{**}	−0.33[*]	−0.17	−0.05	0.01	0.06	0.08
SS-CSV	−1.05 ^{**}	−0.58 ^{**}	−0.33 [*]	−0.16	−0.04	0.03	0.08	0.10
SSNG-CSV	−1.05 ^{**}	−0.57 [*]	−0.33 [*]	−0.16	−0.04	0.02	0.07	0.09
Relative RMSE (model in first column/benchmark)								
Minn-IW	0.65 ^{**}	0.84 [*]	0.94 [*]	0.98	0.99	0.99	0.98	0.98
SS-IW	0.64 ^{**}	0.83 ^{**}	0.92 [*]	0.96	0.97	0.98	0.99	0.99
SSNG-IW	0.65 ^{**}	0.83 ^{**}	0.92 [*]	0.96	0.97	0.97	0.97	0.97
Minn-CSV	0.63 [*]	0.82 [*]	0.92	1.00	1.04	1.06	1.08	1.09
SS-CSV	0.62^{**}	0.81[*]	0.91	0.98	1.03	1.06	1.08	1.09
SSNG-CSV	0.63 [*]	0.82 [*]	0.92	0.99	1.04	1.06	1.08	1.09

Note: The forecast horizon h denotes months. Negative LPDS entries indicate that the mixed-frequency model is superior in terms of density forecasting and values of the RMSE below 1 indicate better point forecasts. Bold entries show the minimum per column. The benchmark model is a VAR(12) including the 10 monthly variables using the steady-state prior with a constant error covariance matrix. Two stars (**) indicate that the Diebold-Mariano test of equal predictive ability is significant at the 1 percent level, whereas a single star (*) indicates significance at the 10 percent level. The test employs the modifications proposed by Harvey et al. (1997).

Table 23: Capacity utilization: forecast evaluation.

Model	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$	$h = 8$
Relative LPDS (model in first column – benchmark)									
Minn-IW	1.40	0.39	0.07	–0.11	–0.16	–0.23	–0.28	–0.30	–0.34
SS-IW	1.30	0.37	0.09	–0.10	–0.19	–0.27	–0.34	–0.38	–0.44
SSNG-IW	1.37	0.42	0.15	–0.03	–0.09	–0.17	–0.24	–0.26	–0.30
Minn-CSV	3.53	0.77	–0.08	–0.43	–0.62	–0.76	–0.88	–0.93	–1.00
SS-CSV	3.39	0.70	–0.12	–0.46	–0.68	–0.84	–0.98	–1.06	–1.14
SSNG-CSV	3.48	0.74	–0.12	–0.47	–0.68	–0.84	–0.98	–1.05	–1.13
Relative RMSE (model in first column/benchmark)									
Minn-IW	0.96 ^{**}	0.95 [*]	0.94 [*]	0.93 [*]	0.93 [*]	0.94	0.94	0.95	0.95
SS-IW	0.95^{**}	0.95[*]	0.93[*]	0.92[*]	0.92[*]	0.92	0.92	0.92	0.93
SSNG-IW	0.95 ^{**}	0.95 [*]	0.94 [*]	0.94 [*]	0.94 [*]	0.94	0.94	0.95	0.95
Minn-CSV	0.98 [*]	0.98	0.97	0.97	0.98	0.99 [*]	0.99 [*]	1.00	1.01
SS-CSV	0.97 ^{**}	0.97 [*]	0.96 [*]	0.96	0.97	0.97	0.97	0.98	0.98 [*]
SSNG-CSV	0.97 [*]	0.97 [*]	0.96 [*]	0.96	0.97	0.97	0.97	0.98	0.99 [*]

Note: The forecast horizon h denotes months. Negative LPDS entries indicate that the mixed-frequency model is superior in terms of density forecasting and values of the RMSE below 1 indicate better point forecasts. Bold entries show the minimum per column. The benchmark model is a VAR(12) including the 10 monthly variables using the steady-state prior with a constant error covariance matrix. Two stars (**) indicate that the Diebold-Mariano test of equal predictive ability is significant at the 1 percent level, whereas a single star (*) indicates significance at the 10 percent level. The test employs the modifications proposed by Harvey et al. (1997).

References

Adolfson, M., J. Lindé, and M. Villani. 2007. “Forecasting Performance of an Open Economy DSGE Model.” *Econometric Reviews* 26 (2–4): 289–328.

Ankargren, S., M. Bjellerup, and H. Shahnazarian. 2017. “The Importance of the Financial System for the Real Economy.” *Empirical Economics* 53 (4): 1553–86.

Ankargren, S., and P. Jonéus. 2020. “Simulation Smoothing for Nowcasting with Large Mixed-Frequency VARs.” *Econometrics and Statistics, Advance Online Publication*, <https://doi.org/10.1016/j.ecosta.2020.05.007>.

Baffigi, A., R. Golinelli, and G. Parigi. 2004. “Bridge Models to Forecast the Euro Area GDP.” *International Journal of Forecasting* 20 (3): 447–60.

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

Bañbura, M., D. Giannone, and L. Reichlin. 2011. “Nowcasting.” In *The Oxford Handbook of Economic Forecasting*, edited by M. P. Clements, and D. F. Hendry, Oxford University Press, chapter 8.

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

- Carriero, A., T. E. Clark, and M. Marcellino. 2016. "Common Drifting Volatility in Large Bayesian VARs." *Journal of Business & Economic Statistics* 34 (3): 375–90.
- Carriero, A., T. E. Clark, and M. Marcellino. 2019. "Large Bayesian Vector Autoregressions with Stochastic Volatility and Non-Conjugate Priors." *Journal of Econometrics* 212 (1): 137–54.
- Carter, C. K., and R. Kohn. 1994. "On Gibbs Sampling for State Space Models." *Biometrika* 81 (3): 541–53.
- Cimadomo, J., and A. D'Agostino. 2016. "Combining Time Variation and Mixed Frequencies: An Analysis of Government Spending Multipliers in Italy." *Journal of Applied Econometrics* 31 (7): 1276–90.
- Clark, T. E. 2011. "Real-Time Density Forecasts From Bayesian Vector Autoregressions with Stochastic Volatility." *Journal of Business & Economic Statistics* 29 (3): 327–41.
- Clark, T. E., and F. Ravazzolo. 2015. "Macroeconomic Forecasting Performance Under Alternative Specifications of Time-Varying Volatility." *Journal of Applied Econometrics* 30 (4): 551–75.
- Cogley, T., and T. J. Sargent. 2005. "Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII US." *Review of Economic Dynamics* 8 (2): 262–302.
- D'Agostino, A., L. Gambetti, and D. Giannone. 2013. "Macroeconomic Forecasting and Structural Change." *Journal of Applied Econometrics* 28 (1): 82–101.
- Del Negro, M., and G. E. Primiceri. 2015. "Time Varying Structural Vector Autoregressions and Monetary Policy: A Corrigendum." *The Review of Economic Studies* 82 (4): 1342–5.
- Del Negro, M., and F. Schorfheide. 2011. "Bayesian Macroeconometrics." In *The Oxford Handbook of Bayesian Econometrics*, edited by J. Geweke, G. Koop, and H. van Dijk, 293–389. Oxford: Oxford University Press.
- Dieppe, A., R. Legrand, and B. van Roye. 2016. *The BEAR Toolbox*, Working Paper No. 1934, Frankfurt, Germany: European Central Bank.
- Durbin, J., and S. J. Koopman. 2012. *Time Series Analysis by State Space Methods*, 2nd edn. Oxford, UK: Oxford University Press.
- 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.
- Forni, C., and M. Marcellino. 2013. *A Survey of Econometric Methods for Mixed-Frequency Data*, Working Paper No. 6, Oslo, Norway: Norges Bank.
- Gelfand, A. E., and A. F. M. Smith. 1990. "Sampling-Based Approaches to Calculating Marginal Densities." *Journal of the American Statistical Association* 85 (410): 398–409.
- Ghysels, E. 2016. "Macroeconomics and The Reality of Mixed Frequency Data." *Journal of Econometrics* 193 (2): 294–314.
- Ghysels, E., A. Sinko, and R. Valkanov. 2007. "MIDAS Regressions: Further Results and New Directions." *Econometric Reviews* 26 (1): 53–90.
- Giannone, D., M. Lenza, and G. E. Primiceri. 2015. "Prior Selection for Vector Autoregressions." *The Review of Economics and Statistics* 97 (2): 436–51.
- Giannone, D., M. Lenza, and G. E. Primiceri. 2019. "Priors for the Long Run." *Journal of the American Statistical Association* 114 (526): 565–80.
- 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.
- Götz, T. B., and K. Hauzenberger. 2018. *Large Mixed-Frequency VARs with A Parsimonious Time-Varying Parameter Structure*, Discussion Paper No. 40. Frankfurt, Germany: Deutsche Bundesbank.
- Griffin, J. E., and P. J. Brown. 2010. "Inference with Normal-Gamma Prior Distributions in Regression Problems." *Bayesian Analysis* 5 (1): 171–88.

- Harvey, A. C., and R. G. Pierse. 1984. "Estimating Missing Observations in Economic Time Series." *Journal of the American Statistical Association* 79 (385): 125–31.
- Harvey, D., S. Leybourne, and P. Newbold. 1997. "Testing the Equality of Prediction Mean Squared Errors." *International Journal of Forecasting* 13 (2): 281–91.
- Huber, F., and M. Feldkircher. 2019. "Adaptive Shrinkage in Bayesian Vector Autoregressive Models." *Journal of Business & Economic Statistics* 37 (1): 27–39.
- Iversen, J., S. Laséen, H. Lundvall, and U. Söderström. 2016. *Real-Time Forecasting for Monetary Policy Analysis: The Case of Sveriges Riksbank*, Working Paper No. 318. Stockholm, Sweden: Sveriges Riksbank.
- Kadiyala, K. R., and S. Karlsson. 1993. "Forecasting with Generalized Bayesian Vector Auto Regressions." *Journal of Forecasting* 12 (3–4): 365–78.
- Kadiyala, K. R., and S. Karlsson. 1997. "Numerical Methods for Estimation and Inference in Bayesian VAR-Models." *Journal of Applied Econometrics* 12 (2): 99–132.
- Karlsson, S. 2013. "Forecasting with Bayesian Vector Autoregression." In *Handbook of Economic Forecasting*, Vol. 2, edited by G. Elliott, and A. Timmermann, 791–897. Elsevier B.V., chapter 15.
- Kastner, G. 2016. "Dealing with Stochastic Volatility in Time Series Using the R Package Stochvol." *Journal of Statistical Software* 69 (5): 1–30.
- Kastner, G., and S. Frühwirth-Schnatter. 2014. "Ancillarity-Sufficiency Interweaving Strategy (ASIS) for Boosting MCMC Estimation of Stochastic Volatility Models." *Computational Statistics and Data Analysis* 76: 408–23.
- Kim, S., N. Shephard, and S. Chib. 1998. "Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models." *The Review of Economic Studies* 65 (3): 361–93.
- Koop, G. M. 2013. "Forecasting with Medium and Large Bayesian VARs." *Journal of Applied Econometrics* 28 (2): 177–203.
- 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.
- Litterman, R. B. 1986. "A Statistical Approach to Economic Forecasting." *Journal of Business & Economic Statistics* 4 (1): 1–4.
- Louzis, D. P. 2019. "Steady-State Modeling and Macroeconomic Forecasting Quality." *Journal of Applied Econometrics* 34 (2): 285–314.
- Mariano, R. S., and Y. Murasawa. 2003. "A New Coincident Index of Business Cycles Based on Monthly and Quarterly Series." *Journal of Applied Econometrics* 18 (4): 427–43.
- Mariano, R. S., and Y. Murasawa. 2010. "A Coincident Index, Common Factors, and Monthly Real GDP." *Oxford Bulletin of Economics and Statistics* 72 (1): 27–46.
- McCausland, W. J., S. Miller, and D. Pelletier. 2011. "Simulation Smoothing for State-Space Models: A Computational Efficiency Analysis." *Computational Statistics & Data Analysis* 55 (1): 199–212.
- McCracken, M. W., and S. Ng. 2016. "FRED-MD: A Monthly Database for Macroeconomic Research." *Journal of Business & Economic Statistics* 34 (4): 574–89.
- Omori, Y., S. Chib, N. Shephard, and J. Nakajima. 2007. "Stochastic Volatility with Leverage: Fast and Efficient Likelihood Inference." *Journal of Econometrics* 140 (2): 425–49.
- Österholm, P. 2008. "Can Forecasting Performance be Improved by Considering the Steady State? An Application to Swedish Inflation and Interest Rate." *Journal of Forecasting* 27 (1): 41–51.
- Österholm, P. 2012. "The Limited Usefulness of Macroeconomic Bayesian VARs When Estimating the Probability of a US Recession." *Journal of Macroeconomics* 34 (1): 76–86.

- Primiceri, G. E. 2005. "Time Varying Structural Vector Autoregressions and Monetary Policy." *The Review of Economic Studies* 72 (3): 821–52.
- Roberts, G. O., and J. S. Rosenthal. 2009. "Examples of Adaptive MCMC." *Journal of Computational and Graphical Statistics* 18 (2): 349–67.
- Rodriguez, A., and G. Puggioni. 2010. "Mixed Frequency Models: Bayesian Approaches to Estimation and Prediction." *International Journal of Forecasting* 26 (2): 293–311.
- Romer, C. D., and D. H. Romer. 2000. "Federal Reserve Information and The Behavior of Interest Rates." *American Economic Review* 90 (3): 429–57.
- Sax, C., and D. Eddelbuettel. 2018. "Seasonal Adjustment by X-13ARIMA-SEATS in R." *Journal of Statistical Software* 87 (11): 1–17.
- Schorfheide, F., and D. Song. 2015. "Real-Time Forecasting with a Mixed-Frequency VAR." *Journal of Business & Economic Statistics* 33 (3): 366–80.
- Stock, J. H., and M. W. Watson. 2001. "Vector Autoregressions." *Journal of Economic Perspectives* 15 (4): 101–15.
- Villani, M. 2009. "Steady State Priors for Vector Autoregressions." *Journal of Applied Econometrics* 24 (4): 630–50.